Osvaldo Gervasi · Beniamino Murgante ·
Chiara Garau · Ana Maria A. C. Rocha ·
M. Noelia Faginas Lago · Ana C. Braga ·
Yeliz Karaca (Eds.)

LNCS 16759

Computational Science and Its Applications – ICCSA 2026 Workshops

Braga, Portugal, June 30 – July 3, 2026
Proceedings, Part II

2 Part II



Springer

Lecture Notes in Computer Science **16759**

Founding Editors

Gerhard Goos
Juris Hartmanis

Editorial Board Members

Elisa Bertino, *Purdue University, West Lafayette, IN, USA*
Wen Gao, *Peking University, Beijing, China*
Bernhard Steffen, *TU Dortmund University, Dortmund, Germany*
Moti Yung, *Columbia University, New York, NY, USA*

The series Lecture Notes in Computer Science (LNCS), including its subseries Lecture Notes in Artificial Intelligence (LNAI) and Lecture Notes in Bioinformatics (LNBI), has established itself as a medium for the publication of new developments in computer science and information technology research, teaching, and education.

LNCS enjoys close cooperation with the computer science R & D community, the series counts many renowned academics among its volume editors and paper authors, and collaborates with prestigious societies. Its mission is to serve this international community by providing an invaluable service, mainly focused on the publication of conference and workshop proceedings and postproceedings. LNCS commenced publication in 1973.

Osvaldo Gervasi · Beniamino Murgante ·
Chiara Garau · Ana Maria A. C. Rocha ·
M. Noelia Faginas Lago · Ana C. Braga ·
Yeliz Karaca
Editors

Computational Science and Its Applications – ICCSA 2026 Workshops

Braga, Portugal, June 30 – July 3, 2026
Proceedings, Part II

Springer

Editors
Osvaldo Gervasi
University of Perugia
Perugia, Perugia, Italy

Chiara Garau
University of Cagliari
Cagliari, Italy

M. Noelia Faginas Lago
University of Perugia
Perugia, Italy

Yeliz Karaca
University of Massachusetts
Worcester, MA, USA

Beniamino Murgante
University of Basilicata
Potenza, Italy

Ana Maria A. C. Rocha
University of Minho
Braga, Portugal

Ana C. Braga
University of Minho
Braga, Portugal

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-032-30517-6 ISBN 978-3-032-30518-3 (eBook)
https://doi.org/10.1007/978-3-032-30518-3

© The Editor(s) (if applicable) and The Author(s), under exclusive license
to Springer Nature Switzerland AG 2027

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

If disposing of this product, please recycle the paper.

Preface

These 9 volumes (LNCS volumes 16758–16766) comprise the peer-reviewed papers from the 57 Workshops of the 2026 International Conference on Computational Science and Its Applications (ICCSA 2026), held between June 30 and July 3, 2026, in Braga, Portugal. The peer-reviewed papers from the six main conference tracks are published in a separate set of 3 volumes (LNCS volumes 16767–16769).

The conference was held in a hybrid format, with the vast majority of participants attending in person, hosted by the University of Minho in Braga, Portugal.

We enabled virtual participation for those who did not attend the event in person due to logistical, political, or economic constraints by adopting a technological infrastructure based on open-source software (Jitsi + Riot) and a commercial Cloud infrastructure. ICCSA 2026 was organized by the University of Minho (Portugal), the University of Perugia (Italy), the University of Basilicata (Italy), the University of Cagliari (Italy), Monash University (Australia), and Kyushu Sangyo University (Japan).

With the 2026 edition, ICCSA celebrated its 26th anniversary, marking a new chapter in its longstanding history after its memorable 25th anniversary celebrated in Istanbul, Türkiye. Hosted in Braga, Portugal, the conference found a particularly meaningful setting in one of Europe's oldest and most vibrant academic cities, a place where historical depth, cultural continuity, and contemporary scientific exchange meet. Braga offers an inspiring European context for renewing the international and intercultural dialogue that has always characterized ICCSA, while giving new visibility to the role of computational science in addressing the complex challenges of cities, territories, technologies, and societies.

ICCSA 2026 marked another fruitful and thought-provoking academic event in the International Conference on Computational Science and Its Applications (ICCSA) conference series, previously held in Istanbul, Türkiye (2025), Hanoi, Vietnam (2024), Athens, Greece (2023), Málaga, Spain (2022), Cagliari, Italy (hybrid with few participants in presence in 2021 and completely online in 2020), Saint Petersburg, Russia (2019), Melbourne, Australia (2018), Trieste, Italy (2017), Beijing, China (2016), Banff, Canada (2015), Guimarães, Portugal (2014), Ho Chi Minh City, Vietnam (2013), Salvador, Brazil (2012), Santander, Spain (2011), Fukuoka, Japan (2010), Suwon, South Korea (2009), Perugia, Italy (2008), Kuala Lumpur, Malaysia (2007), Glasgow, UK (2006), Singapore (2005), Assisi, Italy (2004), Montreal, Canada (2003), and (as ICCS) Amsterdam, the Netherlands (2002) and San Francisco, USA (2001).

Computational Science constitutes the main pillar of most present research, industrial, and commercial applications, and plays a unique role in exploiting innovative ICT technologies, and the ICCSA conference series has, accordingly, provided ample opportunities to researchers and industry practitioners to discuss new ideas, to share complex problems and their solutions, and to shape new trends in Computational Science. As the conference mirrors society from a scientific point of view, this year's undoubtedly dominant theme was large language models, machine learning, and Artificial Intelligence,

and their applications across the most diverse technological, economic, and industrial fields, amongst others.

The ICCSA 2026 conference included 57 workshops, focusing on topical issues of utmost importance to science, technology, and society: from new computational approaches for earth science, to mathematical methods for image processing, new statistical and optimisation methods, several Artificial Intelligence approaches, sustainability issues, smart cities and related technologies, to name a few. In addition, ICCSA 2026 included six general tracks covering computational science and its applications: Computational Methods, Algorithms and Scientific Applications; High-Performance Computing and Networks; Geometric Modelling, Graphics and Visualisation; Advanced and Emerging Applications; Information Systems and Technologies; and Urban and Regional Planning.

For these Workshops Proceedings, we accepted 328 full papers, 23 short papers, and 3 Ph.D. Showcase papers from a total of 1184 submissions (acceptance rate 29.9%). Reviews were single blind, and submissions received 2.6 reviews each, on average. In the Main Conference Proceedings, we accepted 119 full papers, 3 short papers, and 2 Ph.D. Showcase papers from 431 submissions to the General Tracks of the Conference (acceptance rate of 29.9%). We would like to convey our sincere appreciation to the workshop chairs and co-chairs and program committee members, for their diligent work, commitment, and dedication.

The success and consistent maintenance of the ICCSA conference series in general, and of ICCSA 2026 in particular, rely upon the support of many people: authors, presenters, participants, keynote speakers, workshop chairs, session chairs, Organizing Committee members, student volunteers, Program Committee members, Advisory Committee members, International Liaison chairs, reviewers, and other individuals in various roles. Thus, we take this opportunity to wholeheartedly thank each and every one.

We also wish to thank the publisher, Springer, for accepting the proceedings, and for their kind assistance and cooperation during the editing process. We would cordially like to invite you to refer to the ICCSA website https://iccsa.org, where you can find the relevant details regarding this academic endeavour and our event.

June 2026

Chiara Garau
Ana Maria A. C. Rocha
Osvaldo Gervasi
Beniamino Murgante

Welcome Message from the Organizers

On behalf of the Local Organizing Committee of ICCSA 2026, it is our great pleasure to welcome you to the 26th International Conference on Computational Science and Its Applications, held from June 30 to July 3, 2026, at the University of Minho in Braga, Portugal.

We are deeply honoured by the trust placed in us by the ICCSA General Chairs in organizing another edition of this successful conference series. We hope that ICCSA 2026 provided an inspiring environment for scientific exchange, collaboration, and discussion among researchers from all over the world.

The conference took place at the University of Minho, one of the most prestigious higher education institutions in Portugal. Founded in 1973, the University of Minho has campuses in Braga and Guimaraes and is widely recognized for its academic excellence, research activity, and strong international outlook.

Braga, one of the oldest cities in Portugal, is renowned for its historical heritage, cultural vitality, and welcoming atmosphere.

Located in the heart of the Minho region, the city combines tradition with modernity, offering visitors renowned gastronomy, beautiful landscapes, and the unique character of *Vinho Verde* wine.

The scientific program of ICCSA 2026 included plenary lectures delivered by distinguished scientists, together with a broad range of workshops covering several areas of Computational Science and its Applications. We believe that the conference provided excellent opportunities to share ideas, discuss emerging challenges, and foster new collaborations.

In addition to the scientific sessions, participants had the opportunity to enjoy a rich social program.

Planned activities included a Welcome Reception at Largo do Paço, the Rectorate of the University of Minho, located in the former Archbishop's Palace - one of the most emblematic historical buildings in the city center; a *Vinho Verde* wine-tasting experience at Quinta d'Amares, where particiapnts could discover the art of winemaking; and a guided visit to the coastal city of Viana do Castelo, renowned for its maritime heritage, traditional architecture, and beautiful Atlantic scenery. The conference gala dinner took place at Colunata Eventos, located within the Bom Jesus do Monte sanctuary, a UNESCO World Heritage Site, where ICCSA participants enjoyed delicious recipes and a wonderful view over the city of Braga.

The success of ICCSA 2026 would not have been possible without the dedication and enthusiasm of many collaborators and volunteers, whose hard work and commitment greatly contributed to the organization of the conference.

On behalf of the Local Organizing Committee, it was our honour to warmly welcome you to the beautiful city of Braga and we hope you had an enjoyable and productive conference.

We enjoyed meeting you in Braga.

Sincerely yours,

On behalf of the Local Organizing Committee.

June 2026 Ana Maria A. C. Rocha

Organization

Honorary General Chairs

Bernady O. Apduhan	Kyushu Sangyo University, Japan
Kenneth C. J. Tan	Sardina Systems, UK

General Chairs

Ana Maria A. C. Rocha	University of Minho, Portugal
David Taniar	Monash University, Australia
Osvaldo Gervasi	University of Perugia, Italy

Program Committee Chairs

Beniamino Murgante	University of Basilicata, Italy
Chiara Garau	University of Cagliari, Italy

International Advisory Committee

Jemal Abawajy	Deakin University, Australia
Dharma P. Agarwal	University of Cincinnati, USA
Jagannath Aryal	University of Melbourne, Australia
Rajkumar Buyya	University of Melbourne, Australia
Claudia Bauzer Medeiros	University of Campinas, Brazil
Carl K. Chang	Iowa State University, USA
Elizabeth Chang	University of the Sunshine Coast, Australia
Manfred M. Fisher	Vienna University of Economics and Business, Austria
Pierre Frankhauser	University of Franche-Comté/CNRS, France
Marina L. Gavrilova	University of Calgary, Canada
Sumi Helal	University of Florida, USA & Lancaster University, UK
Bin Jiang	University of Gävle, Sweden
Yee Leung	Chinese University of Hong Kong, China
Biswajeet Pradhan	University of Technology Sydney, Australia
My Thai	University of Florida, USA

International Liaison Chairs

Giuseppe Borruso	University of Trieste, Italy
Elise De Donker	Western Michigan University, USA
Maria Noelia Faginas Lago	University of Perugia, Italy
Maria Irene Falcão	University of Minho, Portugal
Robert C. H. Hsu	Chung Hua University, Taiwan
Yeliz Karaca	UMass Chan Medical School, USA
Tae-Hoon Kim	Zhejiang University of Science and Technology, China
Vladimir Korkhov	Saint Petersburg University, Russia
Takashi Naka	Kyushu Sangyo University, Japan
Rafael D. C. Santos	National Institute for Space Research, Brazil
Maribel Yasmina Santos	University of Minho, Portugal
Anastasia Stratigea	National Technical University of Athens, Greece
Nguyen Canh Thai	Thuyloi University, Vietnam
Kh Md Nahiduzzaman	Mohammed VI Polytechnic University, Rabat, Morocco

Workshop and Session Organizing Chairs

Beniamino Murgante	University of Basilicata, Italy
Chiara Garau	University of Cagliari, Italy

Award Chair

Wenny Rahayu	La Trobe University, Australia

Publicity Committee Chairs

Elmer Dadios	De La Salle University, Philippines
Nataliia Kulabukhova	Saint Petersburg University, Russia
Daisuke Takahashi	Tsukuba University, Japan
Doan Quang Tu	Academy of Military Science and Technology, Vietnam
Shangwang Wang	Beijing University of Posts and Telecommunications, China

Local Organizing Committee Chairs

Maria Irene Falcão	University of Minho, Portugal
Ana Cristina Braga	University of Minho, Portugal
Rui José	University of Minho, Portugal
José Manuel Machado	University of Minho, Portugal
Jorge G. Rocha	University of Minho, Portugal

Technology Chair

Damiano Perri	University of Perugia, Italy

Workshops

Workshop on Advancements in Applied Machine-learning and Data Analytics (AAMDA 2026)

Workshop Organizers

Alessandro Costantini	INFN, Italy
Elisabetta Ronchieri	INFN CNAF, Italy

Workshop Program Committee Members

Luca Dell'Agnello	INFN, Italy
Barbara Martelli	INFN, Italy
Daniele Cesini	INFN CNAF, Italy
Giovanni Zurlo	INFN CNAF, Italy
Luca Giommi	INFN CNAF, Italy
Valentino Santucci	Università per Stranieri di Perugia, Italy
Francesco Santini	University of Perugia, Italy
Paolo Mengoni	Hong Kong Baptist University, China
Alessandro Bombini	Istituto Nazionale di Fisica Nucleare, Italy

Artificial Intelligence Tools In Managing Future Communities (AiMFutCOM 2026)

Workshop Organizers

Marco Vona	University of Basilicata, Italy
Francesco Scorza	University of Basilicata, Italy
Donatello Cardone	University of Basilicata, Italy
Maurizio Pollino	ENEA, Italy
Gregorio DAgostino	ENEA, Italy

Workshop Program Committee Members

Maria Danese	CNRISPC, Italy
Rachele Vanessa Gatto	Università della Basilicata, Italy
Simone Corrado	University of Basilicata, Italy
Alessandro Plaisant	University of Sassari, Italy
Sabrina Lai	University of Cagliari, Italy
Alfonso Annunziata	University of Basilicata, Italy
Jolanta Dvarioniene	Kaunas University of Technology, Lithuania

Advanced and Innovative Web Apps 2026 (AIWA 2026)

Workshop Organizers

Damiano Perri	University of Perugia, Italy
Osvaldo Gervasi	University of Perugia, Italy
Luciano Alessandro Ipsaro Palesi	University of Florence, Italy

Workshop Program Committee Members

David Berti	ART SpA, Italy
JungYoon Kim	Gachon University, South Korea
Taihoon Kim	Chonnam National University, South Korea

Advanced Mathematics and Computing Models in Complex Data-Intensive Computational Systems (AMCM 2026)

Workshop Organizers

Yeliz Karaca	UMass Chan Medical School and Massachusetts Institute of Technology, USA
Dumitru Baleanu	Lebanese American University, Lebanon
Osvaldo Gervasi	University of Perugia, Italy
Yudong Zhang	University of Leicester, UK
Majaz Moonis	UMass Chan Medical School, USA

Workshop Program Committee Members

Aziz Dursun	Virginia Tech University, USA
Akif Akgul	Hitit University, Turkey
Nizamettin Aydin	Istanbul Technical University, Turkey
Martin Bohner	Missouri University of Science and Technology, USA
Khan Muhammad	Sungkyunkwan University, South Korea

Advanced Numerical Approaches for assessment and design of no-tension Masonry Structures (ANAMS 2026)

Workshop Organizers

Antonino Iannuzzo	Universitá degli studi del Sannio, Italy
Natalia Pingaro	Politecnico di Milano, Italy
Andrea Montanino	International Centre for Numerical Methods in Engineering, Spain
Elham Mousavian	University of Edinburgh, UK
Stefano Galassi	University of Florence, Italy

Workshop Program Committee Members

Amal Gerges	University of Cagliari, Italy
Vittorio Paris	University of Bergamo, Italy
Alessandra Capolupo	Polytechnic University of Bari, Italy
Carlo Olivieri	Pegaso University, Italy
Generoso Vaiano	Pegaso University, Italy

Luigi Sibille	Norwegian University of Science and Technology, Norway
Francesco Picariello	Universitas Mercatorum, Italy
Ilaria Scarcelli	Politecnico di Bari, Italy
Fabian Bernal Orozco Barrera	National Autonomous University of Mexico, Mexico
Gabriel Barbat	Universitat Politècnica de Catalunya, Spain
Alejandro Cornejo	Universitat Politècnica de Catalunya, Spain
Mohammad Pourfouladi	Politecnico di Milano, Italy

Artificialization of land, Urban Transformation, and Territorial Resilience in a changing Climate (AUTTREC 2026)

Workshop Organizers

Celestina Fazia	Università di Enna Kore, Italy
Kh Md Nahiduzzaman	Mohammed VI Polytechnic University, Morocco
Hassan Radoine	Mohammed VI Polytechnic University, Morocco
Cezar Morar	University of Oradea, Romania
Antonio García Martínez	Universidad de Sevilla, Spain
Clara Stella Vicari Aversa	Università degli Studi Mediterranea di Reggio Calabria, Italy

Workshop Program Committee Members

Federica Sortino	Universitià degli studi di Enna Kore, Italy
Giulia Fernanda Grazia Catania	Universit di Enna Kore, Italy
Maria Francesca Faro	Ministry of Education and Merit (MIM), Italy
Aleksandar Valjarević	University of Belgrade, Serbia
Laila El Ghazouani	Mohammed VI Polytechnic University, Morocco
Adnane Labbaci	Ibn Tofail University, Morocco

Blockchain and Distributed Ledgers: Technologies and Applications (BDLTA 2026)

Workshop Organizers

Vladimir Korkhov	Saint Petersburg State University, Russia
Elena Stankova	Saint Petersburg State University, Russia
Nataliia Kulabukhova	AutoTech, Russia

Workshop Program Committee Members

Vasily Solodkov	National Research University Higher School of Economics, Moscow, Russia
Alexander Bogdanov	St. Petersburg State University, Russia
Alexander Degtyarev	Saint Petersburg State University, Russia
Dmitrii Vasiunin	Deutsche Telekom Cloud Services E.P.E., Greece
Serob Balyan	Institute for Informatics and Automation Problems, National Academy of Sciences of Armenia, Armenia
Suren Abrahamyan	Institute for Informatics and Automation Problems of NAS RA, Armenia
Ashot Gevorkyan	Institute for Informatics and Automation Problems NAS of Armenia, Armenia
Martin Vala	Pavol Jozef Šafárik University, Slovak Republic
Nodir Zaynalov	Samarkand Branch of Tashkent University of Information Technologies Named After Muhammad al Khwarizmi, Uzbekistan
Michail Panteleyev	Saint Petersburg Electrotechnical University, Russia
Alexander Degtyarev	St. Petersburg State University, Russia
Alexander Bogdanov	St. Petersburg State University, Russia

Bio and Neuro inspired Computing and Applications (BIONCA 2026)

Workshop Organizers

Nadia Nedjah	State University of Rio de Janeiro, Brazil
Luiza de Macedo Mourelle	State University de Rio de Janeiro, Brazil

Workshop Program Committee Members

Alan Oliveira	Universidade de Lisboa, Portugal
Rogerio Calazan	Instituto de Estudos do Mar Almirante Paulo Moreira - Brazilian Navy, Brazil
Marcos Farias	Instituto de Engenharia Nuclear, Brazil
Joelmir Ramos	State University de Rio de Janeiro, Brazil
Luigi Maciel Ribeiro	State University of Rio de Janeiro, Brazil

Computational and Applied Mathematics (CAM 2026)

Workshop Organizers

Maria Irene Falcao	University of Minho, Portugal
Fernando Miranda	Universidade do Minho, Portugal

Workshop Program Committee Members

Graca Tomaz	Instituto Politécnico da Guarda, Portugal
Helmuth Malonek	University of Aveiro, Portugal
Isabel Cacao	University of Aveiro, Portugal
Joao Morais	Autonomous Technological Institute of Mexico (ITAM), Mexico
Lidia Aceto	Università del Piemonte Orientale, Italy
Luís Ferrás	University of Minho, Portugal
Patricia Beites	University of Beira Interior, Portugal
Paulo Amorim	Escola de Matemática Aplicada da Fundação Getulio Vargas, Brazil
Regina de Almeida	University of Trás-os-Montes and Alto Douro, Portugal
Ricardo Severino	University of Minho, Portugal

Computational and Applied Statistics for Data & AI (CAS 2026)

Workshop Organizers

Ana Cristina Braga	University of Minho, Portugal
Ines Sena	University of Minho, Portugal
Joao Mendes	Polytechnic Institute of Bragança, Portugal

Workshop Program Committee Members

A. Manuela Goncalves	University of Minho, Portugal
Vera Afreixo	University of Aveiro, Portugal
Ivan Rodriguez Conde	University of Arkansas at Little Rock, USA
Isabel Natário	Universidade Nova de Lisboa, Portugal
Luca Parodi	University of Genoa, Italy
Flávia Pires	Instituto Politecnico de Bragança, Portugal
Beatriz Flamia Azevedo	Instituto Politecnico de Bragança, Portugal
Adriano Silva	Instituto Politecnico de Bragança VAT 600013758, Portugal

Ana Paula Amorim	Universidade do Minho, Portugal
Angela Silva	Universidade do Minho, Portugal
Arunasalam Rahunanthan	Central State University, USA
Bruna Ramos	University of Minho, Portugal
Cecília Maria Vasconcelos Costa e Castro	Universidade do Minho, Portugal
Florbela Correia	Polytechnic Institute of Viana do Castelo, Portugal
Joaquim Goncalves	Polytechnic Institute of Cávado and Ave, Portugal
Lino Costa	University of Minho, Portugal
Luís Miguel Matos	Universidade do Minho, Portugal
Marco Reis	University of Coimbra, Portugal
Maria Mourao	Polytechnic Institute of Viana do Castelo, Portugal
Martin Perez Perez	University of Vigo, Spain
Sandra Aleixo	FCiências.ID, Portugal
Paula Odete Fernandes	Instituto Politecnico de Bragança, Portugal
Iván Rodríguez Conde	Universidade de Vigo, Spain
Telmo Miguel Pires Pinto	University of Minho, Portugal

Circular Economy and Territorial Competitiveness (CETC 2026)

Workshop Organizers

Gerardo Carpentieri	University of Naples Federico II, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Floriana Zucaro	University of Naples Federico II, Italy
Zhadyra Konurbaeva	D. Serikbayev East Kazakhstan Technical University, Kazakhstan
Saule Kumargazhanova	D. Serikbayev East Kazakhstan Technical University, Kazakhstan

Workshop Program Committee Members

Rosaria Battarra	ISMed-CNR, Italy
Laura Ascione	Università degli studi di Napoli Federico II, Italy
Valerio Martinelli	Università di Napoli Federico II, Italy
Madina Yussubaliyeva	D. Serikbayeva East Kazakhstan Technical University, Kazakhstan
Samat Baigereyev	D. Serikbayev East Kazakhstan Technical University, Kazakhstan

Circular Ecological Heritage Communities: Regeneration and Evaluation Strategies for Historic Villages, Landscapes and Rural Territories (CHEC 2026)

Workshop Organizers

Mariarosaria Angrisano	Pegaso University, Italy
Antonia Gravagnuolo	CNR Institute of Heritage Science, Italy
Francesca Buglione	CNR ISPC, Italy
Piera Della Morte	Pegaso Telematic University, Italy

Workshop Program Committee Members

Giulia Datola	Politecnico di Milano, Italy
Marco Rossitti	Politecnico di Milano, Italy
Federico DellAnna	Politecnico di Torino, Italy
Grazia Neglia	Pegaso Telematic University, Italy
Vanessa Assumma	University of Bologna, Italy
Ippolita Mecca	Pegaso Telematic University, Italy
Ferdinando Verardi	Pegaso Telematic University, Italy
Eleonora Sottile	University of Cagliari, Italy
Dina Jovanovic	Politecnico di Milano, Italy
Martina Mulinacci	Politecnico di Milano, Italy
Davide Cerati	Politecnico di Milano, Italy
Daniela Oreni	Politecnico di Milano, Italy
Patrizia Riganti	University of Glasgow, UK
Carlo Olivieri	Pegaso University, Italy
Giulio Cavana	Politecnico di Torino, Italy
Maria Tartari	CNR, Italy
Luigi Fusco Girard	Pegaso Telematic University, Italy
Albina Scioti	Università Telematica Pegaso, Italy
Eleonora Righetto	Università degli Studi di Padova, Italy
Chiara D'Alpaos	University of Padua, Italy
Anthony Suppa	Politecnico di Torino, Italy
Giorgia Malavasi	Politecnico di Torino, Italy
Kaoutare Amini Alaoui	Mohamed VI Polytechnic University, Morocco
Elena Todella	University of Parma, Italy
Leonhard Seyboth	Federal University of Santa Catarina, Brazil
Jianzhuo Xu	Politecnico di Milano, Italy
Ludovica Nasca	University of Catania, Italy
Alessandro Plaisant	University of Sassari, Italy

Maria Carla Saliu	University of Cagliari, Italy
Ons Zoghlami	Sapienza University of Rome, Italy
Irena Ištoka Otković	Josip Juraj Strossmayer University of Osijek, Croatia
Daniele Cannatella	TU Delft, Italy
Tanja Congiu	University of Sassari, Italy
Juan Camilo Giraldo	Institute for Advanced Architecture of Catalonia, Spain
YiChang Liu	Bio-Architecture Formosana, Taiwan
Yasmine Selim	Ain Shams University, Egypt
Simona Tondelli	Alma Mater Studiorum University of Bologna, Italy
Efrossyni Tsakiri	National Technical University of Athens, Greece

Cyber Intelligence and Applications (CIA 2026)

Workshop Organizers

Gianni D'Angelo	University of Salerno, Italy
Francesco Palmieri	University of Salerno, Italy
Massimo Ficco	University of Salerno, Italy

Workshop Program Committee Members

Arcangelo Castiglione	University of Salerno, Italy
Ugo Fiore	University of Salerno, Italy
Michele Mastroianni	University of Foggia, Italy
Marcello Trovati	Edge Hill University, UK
Raffaele Esposito	University of Salerno, Italy
Marek Ogiela	AGH University of Science and Technology, Poland
Chang Choi	Gachon University, South, South Korea
Mauro Migliardi	University of Padua, Italy

Computational Methods, Statistics and Industrial Mathematics (CMSIM 2026)

Workshop Organizers

Maria Filomena Teodoro	CINAV, Escola Naval, Instituto Universitário Militar, Portugal

Paula Simões	Academia Militar, Portugal
Teresa A. Oliveira	University of Lisbon and Universidade Aberta, Portugal

Workshop Program Committee Members

Luis Teodoro	University of Oslo, Norway
Eliana Costa e Silva	Polytechnic of Porto, Portugal
Ana Borges	Instituto Politécnico do Porto, Portugal
Cristina Lopes	Polytechnic Institute of Porto, Portugal
Fernando Lucas Carapau	University of Évora, Portugal
Antonio S. Macias	University of A Coruña, Spain
Matteo Sacchet	University of Turin, Italy
Marina Marchisio Conte	University of Turin, Italy
Christos Kitsos	University of West Attica, Greece
M. Rosário Ramos	Universidade Aberta, University of Lisbon, Portugal
António Pacheco	Instituto Superior Técnico, Portugal
Ana Sofia Rézio	Polytechnic Institute of Lisbon, Portugal
Paula Simões	Military Academy Portuguese Army, Portugal
Ricardo Moura	NOVA University of Lisbon, Portugal
Ana Oliveira	Academia Militar, Portugal
Sílvia Chá	Military Academy, Portugal
Carlos Rodrigues	Instituto Politecnico de Bragança, Portugal

Computational Optimization and Applications (COA 2026)

Workshop Organizers

Ana Maria A. C. Rocha	University of Minho, Portugal
Humberto Rocha	University of Coimbra, Portugal

Workshop Program Committee Members

Eligius Hendrix	Universidad de Málaga, Spain
Joana Dias	University of Coimbra, Portugal
Ana Pereira	Polytechnic Institute of Bragança, Portugal
Lino Costa	University of Minho, Portugal, Portugal
Filipe Alvelos	Universidade do Minho, Portugal
Ana Paula Teixeira	Universidade do Minho, Portugal

Florbela P. Fernandes	Politecnico de Bragança, Portugal
Ana Monteiro	University of Coimbra, Portugal
Rui Pascoal	University of Coimbra, Portugal
Rommel Regis	Saint Joseph's University, USA
Telmo Miguel Pires Pinto	University of Minho, Portugal
Emerson Paiva	Minho University, Portugal

Computational Methods in Advanced Networking Technologies (CoMANT 2026)

Workshop Organizers

Eslam Farsimadan	University of Salerno, Italy
Leila Moradi	University of Salerno, Italy
Francesco Palmieri	University of Salerno, Italy

Workshop Program Committee Members

Massimo Ficco	University of Salerno, Italy
Ugo Fiore	University of Salerno, Italy
Michele Mastroianni	University of Foggia, Italy
Arcangelo Castiglione	University of Salerno, Italy
Francesco Giuseppe Zammarrelli	University of Salerno, Italy
Annalaura Miglino	University of Salerno, Italy
Mojtaba Sharifzadeh	Università di Trento, Italy
Zahra Ebadi Ansaroudi	Fondazione Bruno Kessler, Italy
Mohammad Vazir Panah	University of Salerno, Italy
Gennaro Pio Rimoli	Università degli Studi di Salerno, Italy

Computational Astrochemistry 2026 (CompAstro 2026)

Workshop Organizers

Marzio Rosi	University of Perugia, Italy
Daniela Ascenzi	Università degli Studi di Trento, Italy
Nadia Balucani	University of Perugia, Italy
Stefano Falcinelli	University of Perugia, Italy

Workshop Program Committee Members

Luca Mancini	University of Perugia, Italy
Andrea Giustini	University of Perugia, Italy
Albert Rimola	Universitat Autònoma de Barcelona, Spain
Piero Ugliengo	Università degli Studi di Torino, Italy
Dario Campisi	University of Perugia, Italy
Gianmarco Valuzzo	University of Perugia, Italy
Giacomo Pannacci	Università degli Studi di Perugia, Italy
Giacomo Giorgi	Università degli Studi di Perugia, Italy
Gabriella Di Genova	Università degli Studi di Perugia, Italy
Lisa Giani	Université Grenoble Alpes, IPAG, France
Emilia Valenca Ferreira de Aragao	University of Montpellier, CEA, CNRS, ENSCM, France
Marco Parriani	Università degli studi di Perugia, Italy
Costanza Borghesi	CNR Istituto Nanoscienze, Italy

Computational methods for Porous materials (CompPor 2026)

Workshop Organizers

Vadim Lisitsa	Institute of Petroleum Geology and Geophysics SB RAS, Russia
Evgeniy Romenski	Sobolev Institute of Mathematics, Russia

Workshop Program Committee Members

Tatyana Khachkova	Institute of Petroleum Geology and Geophysics SB RAS, Russia
Dmitry Prokhorov	Institute of Petroleum Geology and Geophysics SB RAS, Russia
Mikhail Novikov	Sobolev Institute of Mathematics SB RAS, Russia
Galina Reshetova	Institute of Computational Mathematics and Mathematical Geophysics SB RAS, Russia
Sergey Solovyev	Sobolev Institute of Mathematics SB RAS, Russia
Vladimir Cheverda	Sobolev Institute of Mathematics SB RAS, Russia

Workshop on Computational Science and HPC (CSHPC 2026)

Workshop Organizers

Elise de Doncker	Western Michigan University, USA
Hideo Matsufuru	High Energy Accelerator Research Organization, Japan
Robert Makin	Western Michigan University, USA

Workshop Program Committee Members

Hiroshi Daisaka	Hitotsubashi University, Japan
James Rhodes	Western Michigan University, USA
Naohito Nakasato	University of Aizu, Japan
Issaku Kanamori	RIKEN, Japan
Takahiro Ueda	Juntendo University, Japan
Keigo Nitadori	RIKEN, Japan
Fukuko Yuasa	KEK, Japan
Shigeyoshi Aoyama	KEK, Japan
Tatsumi Aoyama	University of Tokyo, Japan
Tasnim Gharaibeh	Kalamazoo College, USA

Cities, Technologies and Planning (CTP 2026)

Workshop Organizers

Giuseppe Borruso	University of Trieste, Italy
Beniamino Murgante	University of Basilicata, Italy
Malgorzata Hanzl	Lodz University of Technology, Poland
Anastasia Stratigea	National Technical University of Athens, Greece
Ljiljana Zivkovic	Republic Geodetic Authority, Serbia
Ginevra Balletto	University of Cagliari, Italy

Workshop Program Committee Members

Silvia Battino	University of Sassari, Italy
Mara Ladu	University of Cagliari, Italy
Maria del Mar Munoz	University of Cádiz, Spain
Ainhoa Amaro	EVOCONS, Spain
Enrico Dagostini	University of Malta, Malta

Francesca Krasna	Università degli Studi di Trieste, Italy
Brisol García García	Universidad Politécnica de Quintana Roo, Mexico
Giovanni Mauro	University of Campania Luigi Vanvitelli, Italy
Maria Ronza	University of Naples Federico II, Italy
Clara Di Fazio	University of Trieste, Italy
Andrea Gallo	Università degli Studi di Trieste, Italy
Francesca Sinatra	University of Trieste, Italy

Evaluating Inner Areas Potentials (EIAP 2026)

Workshop Organizers

Diana Rolando	Politecnico di Torino, Italy
Manuela Rebaudengo	Politecnico di Torino, Italy
Alice Barreca	Politecnico di Torino, Italy
Giorgia Malavasi	Politecnico di Torino, Italy

Workshop Program Committee Members

John Accordino	Virginia Commonwealth University, USA
Sara Favargiotti	University of Trento, Italy
Barbara Lino	Università degli Studi di Palermo, Italy
Maddalena Ferretti	Università Politecnica delle Marche, Italy
Eugenio Muccio	University of Naples Federico II, Italy
Giuliano Poli	University of Naples Federico II, Italy
Giulia Datola	Politecnico di Milano, Italy
Elena Todella	University of Parma, Italy
Asja Aulisio	Politecnico di Torino, Italy
Marco Rossitti	Politecnico di Milano, Italy
Martina Spinelli	Politecnico Torino, Italy
Amina Pereno	Politecnico di Torino, Italy
Francesco Bruzzone	Università Iuav di Venezia, Italy
Alexandra Stankulova	Politecnico di Torino, Italy
Benedetta Di Leo	Università Politecnica delle Marche, Italy
Edda Donati	University of Rome Tor Vergata, Italy
Vanessa Assumma	University of Bologna, Italy
Beatrice Mecca	Politecnico di Torino, Italy

sustainablE mobiLity and Last mIle lOgisTic in urban context (ELLIOT 2026)

Workshop Organizers

Tiziana Campisi	University of Enna Kore, Italy
Socrates Basbas	Aristotle University of Thessaloniki, Greece
Florin Rusca	National University of Science and Technology Politehnica Bucharest, Romania
Grigorios Fountas	Aristotle University of Thessaloniki, Greece
Aleksandra Deluka-Tibljas	University of Rijeka, Croatia

Workshop Program Committee Members

Francesco Russo	University of Reggio Calabria, Italy
Oana Dinu	National University of Science and Technology Politehnica Bucharest, Romania
Antonio Polimeni	Università di Messina, Italy
Antonio Russo	University of Enna Kore, Italy
Elena Cocuzza	Univesity of Catania, Italy
Corrado Rindone	Università degli studi Mediterranea di Reggio Calabria, Italy

Econometrics and Multidimensional Evaluation in Urban Environment (EMEUE 2026)

Workshop Organizers

Maria Cerreta	University of Naples Federico II, Italy
Giuseppe Ciciriello	University of Naples Federico II, Italy
Simona Panaro	University of Sussex Business School, UK
Piero Zizzania	Università of Naples Federico II, Italy
Pierluigi Morano	Polytechnic University of Bari, Italy
Francesco Tajani	Sapienza University of Rome, Italy

Workshop Program Committee Members

Giuliano Poli	Università Federico II di Napoli, Italy
Ludovica La Rocca	University of Naples Federico II, Italy
Sabrina Sacco	University of Naples Federico II, Italy

Ecosystem Services in Spacial Planning for climate neutral urban and rural areas (ESSP 2026)

Workshop Organizers

Sabrina Lai	University of Cagliari, Italy
Carmela Gargiulo	Università Federico II di Napoli, Italy
Beniamino Murgante	University of Basilicata, Italy
Francesco Scorza	University of Basilicata, Italy
Corrado Zoppi	University of Cagliari, Italy
Floriana Zucaro	University of Naples Federico II, Italy

Workshop Program Committee Members

Alfonso Annunziata	University of Basilicata, Italy
Ginevra Balletto	University of Cagliari, Italy
Ivan Blečić	University of Cagliari, Italy
Giuseppe Borruso	University of Trieste, Italy
Barbara Caselli	Università degli Studi di Parma, Italy
Maria Cerreta	University of Naples Federico II, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Federica Isola	University of Cagliari, Italy
Francesca Leccis	University of Cagliari, Italy
Federica Leone	University of Cagliari, Italy
Silvia Rossetti	Università degli Studi di Parma, Italy
Luigi Santopietro	Institute of Methodologies for Environmental Analysis CNR IMAA, Italy
Simone Corrado	University of Basilicata, Italy
Rachele Vanessa Gatto	Università della Basilicata, Italy

The 16th International Workshop on Future Information System Technologies and Applications (FiSTA 2026)

Workshop Organizers

Bernady O. Apduhan	Kyushu Sangyo University, Japan
Rafael Santos	National Institute for Space Research (INPE), Brazil

Workshop Program Committee Members

Agustinus Waluyo	La Trobe University, Australia
Earl Ryan Aleluya	Mindanao State University Iligan Institute of Technology, Philippines
Eric Pardede	La Trobe University, Australia
Kai Cheng	Kyushu Sangyo University, Japan
Fenghui Yao	Tennessee State University, USA
Yusuke Gotoh	Okayama University, Japan
Kazuaki Tanaka	Kyushu Institute of Technology, Japan
Toshihiro Yamauchi	Okayama University, Japan
Toshihiro Uchibayashi	Kyushu University, Japan
Krishnamoorthy Ranganathan	Chennai Institute of Technology, India
Orven Llantos	MSU-Iligan Institute of Technology, Philippines
Anton Louise De Ocampo	Batangas State University, Philippines

Geographical Analysis, Urban Modeling, Spatial Statistics (Geog-And-Mod 2026)

Workshop Organizers

Giuseppe Borruso	University of Trieste, Italy
Beniamino Murgante	University of Basilicata, Italy
Hartmut Asche	Hasso Plattner Institute for Digital Engineering, Germany
Rodrigo Tapia-McClung	Centro de Investigación en Ciencias de Información Geoespacial, Mexico
Andreas Fricke	Hasso Plattner Institute for Digital Engineering, University of Potsdam, Germany

Workshop Program Committee Members

Ginevra Balletto	University of Cagliari, Italy
Mara Ladu	University of Cagliari, Italy
Gabriella Schoier	Università di Trieste, Italy
Marco Mazzarino	Iuav University of Venice, Italy
Malgorzata Hanzl	Lodz University of Technology, Poland
Anastasia Stratigea	National Technical University of Athens, Greece
Maria del Mar Munoz	University of Cádiz, Spain
Veronica Camerada	Università degli Studi di Sassari, Italy

Generative and Neuro-Symbolic Artificial Intelligence for Real-World Systems (GNS-AI 2026)

Workshop Organizers

Hugo Peixoto	University of Minho, Portugal
Manuel Rodrigues	University of Minho, Portugal
Regina Sousa	University of Minho, Portugal

Workshop Program Committee Members

Dalila Duraes	Universidade do Minho, Portugal
Paulo Novais	Universidade do Minho, Portugal
Laura Cavalcante	Instituto Superior de Engenharia do Porto, Portugal
Diogo Martinho	Instituto Superior de Engenharia de Porto, Portugal
Hélia Guerra	University of the Azores, Portugal
Rohit Raja	Guru Ghasidas Vishwavidyalaya, India
Hien Nguyen	University of Information Technology, Vietnam
Antonio Carlos Da Silva Abelha	Universidade do Minho, Portugal
Cristiana Neto	University of Minho, Portugal
José Machado	University of Minho, Portugal
Luís Conceicao	Instituto Superior de Engenharia do Porto, Portugal
Omar Gomez	Universidad Nacional de Colombia, Colombia
Luis Mendes Gomes	University of the Azores, Portugal
Nicolas Lori	University of the Azores, Portugal
Sara Balderas Díaz	University of Cádiz, Spain
Gabriel Guerrero Contreras	University of Cádiz, Spain
Francisco Zayas Gato	Universidade da Coruña, Spain
Diana Ferreira	University of Minho, Portugal
Jolanta Miliauskaite	Vilnius University, Lithuania

Geomatics for Resource Monitoring and Management (GRMM 2026)

Workshop Organizers

Alberico Sonnessa	Politecnico di Bari, Italy
Eufemia Tarantino	Politecnico di Bari, Italy
Marco Locurcio	Polytechnic University of Bari, Italy

Workshop Program Committee Members

Valeria Belloni	Sapienza University of Rome, Italy
Roberta Ravanelli	University of Liège, Belgium
Michela Ravanelli	Sapienza Università di Roma, Italy
Alessandra Mascitelli	University of Chieti Pescara, Italy
Francesco Bimbo	University of Foggia, Italy
Francesco Di Capua	University of Basilicata, Italy
Gaetano Falcone	Università degli Studi di Napoli Federico II, Italy
Francesco Tajani	Sapienza University of Rome, Italy
Paola Amoruso	LUM Giuseppe Degennaro University, Italy
Pierfrancesco De Paola	University of Naples Federico II, Italy
Isidora Đurić	University of Novi Sad, Serbia
Rossella Petti	Politecnico di Bari, Italy
Valeria Cera	University of Naples Federico II, Italy
Antonino Iannuzzo	Universitá degli studi del Sannio, Italy
Pier Francesco Recchi	Università di Napoli Federico II, Italy
Giuseppe Francesco Cesare Lama	University of Naples Federico II, Italy
Ester Scotto di Perta	University of Naples Federico II, Italy
Antonino Maltese	University of Palermo, Italy
Andrea Montanino	University of Naples, Italy
Carla Cavallo	University of Salerno, Italy
Dario Esposito	Polytechnic University of Bari, Italy
Rosaria Parente	Universitas Mercatorum, Italy
Claudia Panciera	Water Research Institute (IRSA) CNR, Italy
Pasquale Perrini	Consorzio Interuniversitario per l'Idrologia, Italy
Beniamino Onorati	Università degli Studi della Basilicata, Italy
Francesco Sica	Sapienza University of Rome, Italy

International Workshop on Information and Knowledge in the Internet of Things (IKIT 2026)

Workshop Organizer

Teresa Guarda	Universidad Estatal Península de Santa Elena, Ecuador

Workshop Program Committee Members

Isabel Lopes	Instituto Politécnico de Bragança, Portugal
Luis Mazon	Bitrum Research Group, Spain

Arnulfo Alanis	Tecnológico Nacional de México, Campus Tijuana, Mexico
Filipe Mota Pinto	Instituto Politécnico de Leiria, Portugal
Luis Enrique Chuquimarca Jiménez	Universidad Estatal Penìnsula de Santa Elena, Ecuador
Gustavo Gatica	Universidad Andres Bello, Chile
Jairo Coronado Hernandez	Universidad de la Costa, Colombia
Maria Fernanda Augusto	BITrum Research Center, Spain
José María Díaz Nafría	Universidad a Distancia de Madrid, Spain
Modestos Stavrakis	University of the Aegean, Greece
Isabel Ribeiro	Instituto Politécnico Bragança, Portugal
Bruno Sousa	University of Coimbra, Portugal
Datzania Villao	Universidad Estatal Península de Santa Elena, Ecuador

Intelligent Sustainable Development: from Models to Real-World Impact (ISD 2026)

Workshop Organizers

Diego Brandão	Federal Center for Technological Education of Rio de Janeiro, Brazil
Angela Ferreira	Polytechnic University of Bragança, Portugal
Eduardo Bezerra	Federal Center for Technological Education of Rio de Janeiro, Brazil
Raphael Guerra	Universidade Federal Fluminense, Brazil
Douglas Cardoso	University of Coimbra, Portugal
Rui Pedro Lopes	Instituto Politecnico de Bragança, Portugal

Workshop Program Committee Members

Wagner Telles	Universidade Federal Fluminense, Brazil
Jader Lugon Junior	Instituto Federal Fluminense, Brazil
Edcarllos Santos	Universidade Federal do Sul da Bahia, Brazil
Kennedy Fernandes	Universidade Federal de Juiz de Fora, Brazil
Israel Mendonca	Kumamoto University, Japan
Pedro Henrique Gonzalez	Federal University of Rio de Janeiro, Brazil
Pablo Luiz Araujo Munhoz	Federal University of Ouro Preto, Brazil
Elloa Guedes	Amazonas State University, Brazil
Danielli Araújo Lima	Instituto Federal do Triângulo Mineiro, Brazil
Ary Henrique Oliveira	Universidade Federal de Tocantins, Brazil

Paula Odete Fernandes	Instituto Politecnico de Bragança, Portugal
Hiago Rocha	National Laboratory for Scientific Computing, Brazil
Andre Luiz Coutinho Merlo	Centro Federal de Educacao Tecnologica Celso Suckow da Fonseca CEFET RJ, Brazil
Kennedy Fernandes	Universidade Federal de Juiz de Fora, Brazil
Ary Henrique Oliveira	Universidade Federal de Tocantins, Brazil
Mayara Omai	Centro Federal de Educacao Tecnologica Celso Suckow da Fonseca CEFET RJ, Brazil

Multidimensional Evolutionary Evaluations for Transformative Approaches (MEETA 2026)

Workshop Organizers

Maria Cerreta	University of Naples Federico II, Italy
Giuliano Poli	Università Federico II di Napoli, Italy
Gaia Daldanise	CNR Institute for Research on Innovation and Services for Development, Italy
Ludovica La Rocca	University of Federico II of Naples, Italy
Maria Somma	University of Naples Federico II, Italy
Sabrina Sacco	University of Naples Federico II, Italy

Workshop Program Committee Members

Chiara Mazzarella	TUDelft, The Netherlands
Simona Panaro	University of Sussex Business School, UK
Eugenio Muccio	University of Naples Federico II, Italy

Building multi-dimensional Models for assessing complex Environmental Systems (MES 2026)

Workshop Organizers

Vanessa Assumma	University of Bologna, Italy
Caterina Caprioli	Politecnico di Torino, Italy
Giulia Datola	Politecnico di Milano, Italy
Federico DellAnna	Politecnico di Torino, Italy
Marta Dell'Ovo	Politecnico di Milano, Italy
Marco Rossitti	Politecnico di Milano, Italy

Workshop Program Committee Members

Mariarosaria Angrisano	Pegaso University, Italy
Giuliano Poli	University of Naples Federico II, Italy
Francesco Sica	Sapienza University of Rome, Italy
Eugenio Muccio	University of Naples Federico II, Italy
Sebastiano Barbieri	Politecnico di Torino, Italy
Giulio Cavana	Politecnico di Torino, Italy
Marta Bottero	Politecnico di Torino, Italy
Francesca Torrieri	Politecnico di Milano, Italy
Maksims Feofilovs	Riga Technical University, Latvia
Daniela Maiullari	Delft University of Technology, The Netherlands
Elena Todella	University of Parma, Italy
Yvonne Russo	University of Campania L. Vanvitelli, Italy
Tommaso Lazzarin	University of Padua, Italy
Francesco Tajani	Sapienza University of Rome, Italy
Claudia Amadei	University of Padua, Italy
Isabella Maria Lami	Politecnico di Torino, Italy
Federico Vagnon	Politecnico di Torino, Italy
Francesca Torrieri	Politecnico di Milano, Italy
Elena De Angelis	Politecnico di Torino, Italy
Piergiorgio Farfariello	Politecnico di Torino, Italy
Francesca Biondi	Politecnico di Torino, Italy
Giulio Cavana	Politecnico di Torino, Italy
Fabiana Forte	Università della Campania Luigi Vanvitelli, Italy
Siyu Qiu	University of Bologna, Italy
Eugenio Muccio	University of Naples Federico II, Italy
Iuliia Kozlova	University of Bologna, Italy
Danny Casprini	Politecnico di Milano, Italy
Franco Corti	University of Padua, Italy
Elisa Sbardella	Politecnico di Torino, Italy
Andrea Sarcina	Università di Padova, Italy
Daniela Menna	Università degli studi della Campania Luigi Vanvitelli, Italy
Sebastiano Barbieri	Politecnico di Torino, Italy
Ana Soares	University of Minho, Portugal
Francesco Sica	Sapienza University of Rome, Italy
Chiara D'Alpaos	University of Padua, Italy
Alessandra Oppio	Politecnico di Milano, Italy
Maria Rosaria Guarini	Sapienza University of Rome, Italy
Alessia Crisopulli	Politecnico di Milano, Italy
Elisabetta Colucci	Politecnico di Torino, Italy

Giorgia Sugoni	Links Foundation, Italy
Lucia Petronio	University of Bologna, Italy
Marco Alvise Bragadin	University of Bologna, Italy

Computational Mechanics, Intelligence, and Data (MIND 2026)

Workshop Organizers

Lelio Campanile	Università degli Studi della Campania Luigi Vanvitelli, Italy
Renato Zona	Università della Campania Vanvitelli, Italy

Workshop Program Committee Members

Vincenzo Minutolo	Università degli Studi della Campania Luigi Vanvitelli, Italy
Fiammetta Marulli	University of Campania Luigi Vanvitelli, Italy
Simone Palladino	Università degli Studi della Campania Luigi Vanvitelli, Italy
Ali Forooghi	University of Campania Luigi Vanvitelli, Italy
Mario Argenziano	University of Palermo, Italy
Donato Perfetto	University of Campania Luigi Vanvitelli, Italy

Modelling Liveable Cities: Techniques, Methods, Challenges, and Perspectives Behind the X-Minute City (MLC 2026)

Workshop Organizers

Federico Mara	University of Pisa, Italy
Valerio Cutini	University of Pisa, Italy

Workshop Program Committee Members

Beniamino Murgante	University of Basilicata, Italy
Chiara Garau	University of Cagliari, Italy
Francesco Zullo	University of LAquila, Italy
Alessandro Marucci	University of L Aquila, Italy
Chiara Di Dato	University of L Aquila, Italy
Giampiero Lombardini	University of Genoa, Italy

Alfonso Annunziata	University of Basilicata, Italy
Simone Corrado	University of Basilicata, Italy
Müslüm Hacar	Yildiz Technical University, Turkey
Federica Deri	University of Pisa, Italy
Chiara Anselmi	University of Pisa, Italy

Mathematical Methods for Image Processing and Understanding 2026 (MMIPU 2026)

Workshop Organizers

Ivan Gerace	Università degli Studi di Perugia, Italy
Gianluca Vinti	Università degli Studi di Perugia, Italy
Arianna Travaglini	Lucian Blaga University of Sibiu, Romania

Workshop Program Committee Members

Marco Baioletti	University of Perugia, Italy
Marco Donatelli	University of Insubria, Italy
Muhammad Hanif	Ghulam Ishaq Khan Institute of Engineering Sciences and Technology, Pakistan
Francesco Marchetti	University of Padua, Italy
Wolfgang Erb	Università degli Studi di Padua, Italy
Ana Maria Acu	Lucian Blaga University of Sibiu, Romania
Simone Rebegoldi	University of Modena and Reggio Emilia, Italy
Valentina Giorgetti	Università degli Studi di Perugia, Italy
Vita Leonessa	University of Basilicata, Italy

MOdels and indicators for assessing and measuring the urban settlement deVElopment in the view of NET ZERO by 2050 (MOVEto0 2026)

Workshop Organizers

Lorena Fiorini	University of L'Aquila, Italy
Lucia Saganeiti	Institute of Methodologies for Environmental Analysis (IMAA) CNR, Italy
Angela Pilogallo	Institute of Methodologies for Environmental Analysis (IMAA) CNR, Italy

Francesco Zullo	University of L'Aquila, Italy
Alessandro Marucci	University of L'Aquila, Italy

Workshop Program Committee Members

Ginevra Balletto	University of Cagliari, Italy
Giuseppe Borruso	University of Trieste, Italy
Chiara Garau	University of Cagliari, Italy
Beniamino Murgante	University of Basilicata, Italy
Ljiljana Zivkovic	Republic Geodetic Authority, Serbia
Luigi Santopietro	Institute of Methodologies for Environmental Analysis National Research Council of Italy CNR IMAA, Italy
Ilaria Delponte	University of Genoa, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Chiara Di Dato	University of L'Aquila, Italy

Numerical Aspects in Machine Learning For Dynamical Systems (NAMLYS 2026)

Workshop Organizers

Pasquale De Luca	Parthenope University of Naples, Italy
Livia Marcellino	University of Naples Parthenope, Italy
Valeria Mele	University of Naples Federico II, Italy

Workshop Program Committee Members

Dario Caramiello	University of Naples Parthenope, Italy
Giuseppe Fiorillo	Parthenope University of Naples, Italy
Marco Lettiero	University of Naples Parthenope, Italy
Francesco Peluso	University of Naples Parthenope, Italy

Numerical Methods and Machine Learning in Science & Engineering (NMLSE 2026)

Workshop Organizers

Cecília Coelho	Helmut Schmidt University, Germany
M. Fernanda P. Costa	University of Minho, Portugal

Luís Ferrás	University of Minho, Portugal
Eliana Costa e Silva	Polytechnic of Porto, Portugal

Workshop Program Committee Members

Osvaldo Gramaxo de Freitas	Universitat de València, Spain
Estela Bicho	University of Minho, Portugal
Flora Ferreira	University of Porto, Portugal
Bruno Veloso	Universidade do Porto, Portugal
Luís Louro	University of Minho, Portugal
Weronika Wojtak	CCG/ZGDV Institute, Portugal
Wolfram Erlhagen	University of Minho, Portugal

6th Workshop on Privacy in the Cloud/Edge/IoT World 2026 (PCEIoT 2026)

Workshop Organizers

Lelio Campanile	Università degli Studi della Campania Luigi Vanvitelli, Italy
Mauro Iacono	Università degli Studi della Campania Luigi Vanvitelli, Italy
Michele Mastroianni	University of Foggia, Italy

Workshop Program Committee Members

Fiammetta Marulli	University of Campania Luigi Vanvitelli, Italy
Christian Riccio	Università degli Studi della Campania Luigi Vanvitelli, Italy
Arcangelo Castiglione	University of Salerno, Italy
Gianni D'Angelo	University of Salerno, Italy
Biagio Boi	University of Salerno, Italy
Luigi Piero Di Bonito	Università degli Studi della Campania, Italy

Preserving Our Past: Spatial and Remote Sensing Technologies for Cultural Heritage in a Changing Climate (POP 2026)

Workshop Organizers

Maria Danese	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy
Rosa Lasaponara	Institute of Methodologies for Environmental Analysis (CNR IMAA), Italy
Nicola Masini	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy

Workshop Program Committee Members

Dario Gioia	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy
Marilisa Biscione	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy
Valentina Florio	CNR, Italy
Antonio Minervino Amodio	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy

Scientific Computing Infrastructure (SCI 2026)

Workshop Organizers

Vladimir Korkhov	Saint Petersburg State University, Russia
Elena Stankova	Saint Petersburg State University, Russia

Workshop Program Committee Members

Adam Belloum	University of Amsterdam, The Netherlands
Dmitrii Vasiunin	Deutsche Telekom Cloud Services E.P.E., Greece
Serob Balyan	Institute for Informatics and Automation Problems, National Academy of Sciences of Armenia, Armenia
Suren Abrahamyan	Institute for Informatics and Automation Problems of NAS RA, Armenia
Ashot Gevorkyan	Institute for Informatics and Automation Problems NAS of Armenia, Armenia

Michal Hnatič	P.J. Šafárik University in Košice, Slovak Republic
Martin Vala	Pavol Jozef Šafárik University, Slovak Republic
Nodir Zaynalov	Samarkand Branch of Tashkent University of Information Technologies Named After Muhammad al Khwarizmi, Uzbekistan
Michail Panteleyev	Saint Petersburg Electrotechnical University, Russia
Alexander Degtyarev	Saint Petersburg State University, Russia
Alexander Bogdanov	St. Petersburg State University, Russia
Nataliia Kulabukhova	AutoTech, Russia
Alexander Degtyarev	St.Petersburg State University, Russia
Alexander Bogdanov	St. Petersburg State University, Russia

Socio-Economic and Environmental Models for Land Use Management (SEMLUM 2026)

Workshop Organizers

Debora Anelli	Sapienza University of Rome, Italy
Pierluigi Morano	Polytechnic University of Bari, Italy
Francesco Tajani	Sapienza University of Rome, Italy
Felicia Di Liddo	Polytechnic University of Bari, Italy

Workshop Program Committee Members

Benedetto Manganelli	Università degli Studi della Basilicata, Italy
Emma Sabatelli	Sapienza Università di Roma, Italy
María Paz Sáez-Pérez	University of Granada, Spain
Yasmine Selim	Ain Shams University, Egypt
Daniela Tavano	University of Calabria, Italy
Francesco Calabrò	Mediterranea University of Reggio Calabria, Italy
Pierfrancesco De Paola	University of Naples Federico II, Italy
Francesco Paolo Del Giudice	University of Rome La Sapienza, Italy
Paola Amoruso	LUM Giuseppe Degennaro University, Italy
Rubina Canesi	Università degli Studi di Padova, Italy
Daniele La Rosa	University of Catania, Italy
Giuseppe Cerullo	Sapienza University of Rome, Italy
Antonella Roma	Sapienza University of Rome, Italy
Francesca Fariello	Sapienza University of Rome, Italy

Ports and logistics of the future - Smartness and Sustainability (SmartPorts 2026)

Workshop Organizers

Andrea Gallo	Università degli Studi di Trieste, Italy
Giuseppe Borruso	University of Trieste, Italy
Gianfranco Fancello	University of Cagliari, Italy
Veronica Camerada	Università degli studi di Sassari, Italy
Clara Di Fazio	University of Trieste, Italy

Workshop Program Committee Members

Patrizia Serra	University of Cagliari, Italy
Francesco Bruzzone	Università Iuav di Venezia, Italy
Marco Mazzarino	University Iuav of Venice, Italy
Maria del Mar Munoz	University of Cádiz, Spain
Ainhoa Amaro	EVOCONS, Spain
Brunella Brundu	University of Sassari, Italy
Giovanni Mauro	University of Campania Luigi Vanvitelli, Italy
Maria Ronza	University of Naples Federico II, Italy
Marcello Tadini	Università del Piemonte Orientale, Italy
Luca Toneatti	University of Trieste, Italy
Mara Ladu	University of Cagliari, Italy
Ginevra Balletto	University of Cagliari, Italy
Francesca Sinatra	University of Trieste, Italy

Sustainable evolution of long-Distance frEight and paSsenger Transport (SOLIDEST 2026)

Workshop Organizers

Antonio Russo	University of Enna Kore, Italy
Giuseppe Musolino	Università Mediterranea di Reggio Calabria, Italy
Socrates Basbas	Aristotle University of Thessaloniki, Greece
Oana Dinu	National University of Science and Technology Politehnica Bucharest, Romania
Sanja Šurdonja	University of Rijeka, Croatia
Elen Twrdy	University of Ljubljana, Slovenia

Workshop Program Committee Members

Tiziana Campisi	University of Enna Kore, Italy
Massimo Di Gangi	Università di Messina, Italy
Antonio Polimeni	Università di Messina, Italy
Efstathios Bouhouras	Aristotle University of Thessaloniki, Greece
Florin Rusca	National University of Science and Technology Politehnica Bucharest, Romania
Vittorio Astarita	University of Calabria, Italy
Giuseppe Guido	Università della Calabria, Italy
Stefanos Tsigdinos	National Technical University of Athens, Greece
Chiara Garau	University of Cagliari, Italy

Sustainability Performance Assessment: Models, Approaches, and Applications Toward Interdisciplinary and Integrated Solutions (SPA 2026)

Workshop Organizers

Francesco Scorza	University of Basilicata, Italy
Sabrina Lai	University of Cagliari, Italy
Jolanta Dvarioniene	Kaunas University of Technology, Lithuania
Michele Campagna	University of Cagliari, Italy

Workshop Program Committee Members

Corrado Zoppi	University of Cagliari, Italy
Maria Cerreta	University of Naples Federico II, Italy
Giuseppe Borruso	University of Trieste, Italy
Chiara Garau	University of Cagliari, Italy
Ginevra Balletto	University of Cagliari, Italy
Alenka Poplin	Iowa State University, USA
Francesco Zullo	University of L'Aquila, Italy
Alessandro Marucci	University of L'Aquila, Italy
Federica Leone	University of Cagliari, Italy
Simone Corrado	University of Basilicata, Italy
Alfonso Annunziata	University of Basilicata, Italy
Rachele Vanessa Gatto	Università della Basilicata, Italy
Francesco Paolo Del Giudice	University of Basilicata, Italy

Specifics of smart cities development in Europe (SPEED 2026)

Workshop Organizers

Chiara Garau	University of Cagliari, Italy
Katarína Vitálišová	Matej Bel University, Slovak Republic
Marco Fanfani	University of Florence, Italy
Anna Vaňová	Matej Bel University, Slovak Republic
Kamila Borsekova	Matej Bel University, Slovak Republic
Paola Zamperlin	University of Pisa, Italy

Workshop Program Committee Members

Francesca Maltinti	University of Cagliari, Italy
Silvia Rossetti	Università degli Studi di Parma, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Floriana Zucaro	University of Naples Federico II, Italy
Giovanna Concu	University of Cagliari, Italy
Paolo Nesi	University of Florence, Italy
Pierfrancesco Bellini	University of Florence, Italy
Gerardo Carpentieri	University of Naples Federico II, Italy
Radoslav Kožiak	Matej Bel University, Slovak Republic
Lucia Dobrotová	Matej Bel University in Banská Bystrica, Slovak Republic
Veronika Piscová	Institute of Landscape Ecology Slovak Academy of Sciences, Slovak Republic
Darina Rojíková	Matej Bel University, Slovak Republic
Katarína Sýkorová	Matej Bel University, Slovak Republic
Silvia Ondrejičková	Comenius University Bratislava, Slovak Republic
Vladimír Ondrejička	Comenius University Bratislava, Slovak Republic
Maros Finka	Comenius University Bratislava, Slovak Republic
Jan Nevima	Masaryk University, Czech Republic
Christian Corral Burau	Leibniz University Hannover, Germany
Kamila Turečková	Silesian University in Opava, Czech Republic
Luciano Alessandro Ipsaro Palesi	University of Florence, Italy
Enrico Collini	University of Florence, Italy
Gianni Pantaleo	University of Florence, Italy
Krishnamoorthy Ranganathan	Chennai Institute of Technology, India

Smart, Safe, and Health Cities (SSHC 2026)

Workshop Organizers

Chiara Garau	University of Cagliari, Italy
Gerardo Carpentieri	University of Naples Federico II, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Martina Carra	University of Brescia, Italy
Vincenza Torrisi	University of Catania, Italy
Silvia Rossetti	Università degli Studi di Parma, Italy

Workshop Program Committee Members

Romano Fistola	University of Naples Federico II, Italy
Francesca Maltinti	University of Cagliari, Italy
Barbara Caselli	Università degli Studi di Parma, Italy
Tiziana Campisi	University of Enna Kore, Italy
Katarína Sýkorová	Matej Bel University, Slovak Republic
Tazyeen Alam	Università degli Studi di Cagliari, Italy
Francesco Zullo	University of L'Aquila, Italy
Giovanna Concu	University of Cagliari, Italy
Anastasia Stratigea	National Technical University of Athens, Greece
Ezio Riggi	National Research Council, Italy
Miriam Distefano	CNR, Italy

Smart and Sustainable Island Communities (SSIC 2026)

Workshop Organizers

Chiara Garau	University of Cagliari, Italy
Anastasia Stratigea	National Technical University of Athens, Greece
Yiota Theodora	National Technical University of Athens, Greece
Giovanna Concu	University of Cagliari, Italy
Tiziana Campisi	University of Enna Kore, Italy
Tanja Congiu	University of Sassari, Italy

Workshop Program Committee Members

Francesca Maltinti	University of Cagliari, Italy
Mauro Coni	University of Cagliari, Italy
Paola Zamperlin	University of Florence, Italy

Tarek Teba	University of Portsmouth, UK
Gerardo Carpentieri	University of Naples Federico II, Italy
Floriana Zucaro	University of Naples Federico II, Italy
Roberto Minunno	Curtin University, Australia
Apostolos Lagarias	University of Thessaly, Greece
Giuseppina Vacca	University of Cagliari, Italy
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Tazyeen Alam	Università degli Studi di Cagliari, Italy
Mauro D'Apuzzo	University of Cassino and Southern Lazio, Italy
Mana Dastoum	Polytechnic University of Madrid, Spain
Alfonso Annunziata	University of Basilicata, Italy
Barbara Caselli	Università degli Studi di Parma, Italy
Silvia Rossetti	Università degli Studi di Parma, Italy
Gian Felice Giaccu	University of Sassari, Italy
Anastasia Stratigea	National Technical University of Athens, Greece
Vincenza Torrisi	University of Catania, Italy

Sustainable Tourism Economy, Silk Road, Reuse and Enhancement of Cultural Heritage (STESilkRoad 2026)

Workshop Organizers

Celestina Fazia	Università di Enna Kore, Italy
Alexandru Ilies	University of Oradea, Romania
Dorina Camelia Ilies	University of Oradea, Romania
Cristina Natoli	Ministero dell'Ambiente e della Sicurezza Energetica, Italy
Bahodirhon Safarov	Samarkand Branch of Tashkent State University of Economics, Uzbekistan
Clara Stella Vicari Aversa	Università degli Studi Mediterranea di Reggio Calabria, Italy

Workshop Program Committee Members

Giulia Fernanda Grazia Catania	Università di Enna Kore, Italy
Federica Sortino	Università degli Studi di Enna Kore, Italy
Thowayeb Hassan	King Faisal University, Saudi Arabia
Tatar Corina	University of Oradea, Romania
Maria Francesca Faro	Ministry of Education and Merit (MIM), Italy

Theoretical and Computational Chemistry and its Applications (TCCM 2026)

Workshop Organizers

Noelia Faginas Lago	Università di Perugia, Italy
Andrea Lombardi	Università di Perugia, Italy
Marcos Mandado Alonso	University of Vigo, Spain

Workshop Program Committee Members

Jose Manuel Hermida Ramon	Universidade de Vigo, Spain
Luca Mancini	University of Perugia, Italy
Stefano Falcinelli	University of Perugia, Italy
Ángel Martín Pendás	Universidad de Oviedo, Spain
Enrique M. Cabaleiro Lago	Universidade de Santiago de Compostela, Spain
Giacomo Giorgi	Università degli Studi di Perugia, Italy
Stefano Evangelisti	Université de Toulouse, France
Nicolás Ramos Berdullas	University of Vigo, Spain
Juan José Nogueira	Universidad Autónoma de Madrid, Spain
Giacomo Pannacci	Università degli Studi di Perugia, Italy

Transport Infrastructures for Smart Cities (TISC 2026)

Workshop Organizers

Francesca Maltinti	University of Cagliari, Italy
Mauro Coni	University of Cagliari, Italy
Chiara Garau	University of Cagliari, Italy
Benedetto Barabino	University of Brescia, Italy
Nicoletta Rassu	Università di Cagliari, Italy
James Rombi	University of Cagliari, Italy

Workshop Program Committee Members

Tiziana Campisi	University of Enna Kore, Italy
Mauro D'Apuzzo	University of Cassino and Southern Lazio, Italy
Roberto Ventura	Università degli Studi di Brescia, Italy
Roberto Minunno	Curtin University, Australia

Giovanna Concu	University of Cagliari, Italy
Claudia Loggia	University of KwaZulu-Natal, South Africa

Transforming Urban Analytics: The Impact of Crowdsourced Mapping and Advanced AI Techniques on Future Cities (Tr-UrbAna 2026)

Workshop Organizers

Ayse Giz Gulnerman Gengec	Ankara Hacı Bayram Veli University, Turkey
Müslüm Hacar	Yildiz Technical University, Turkey

Workshop Program Committee Members

Beniamino Murgante	University of Basilicata, Italy
Abdulkadir Memduhoğlu	Harran University, Turkey
Zeynel Abidin Polat	Izmir Katip Celebi University, Turkey
Güzide Miray Perihanoglu	Van Yuzuncu Yil University, Turkey
Tugba Memisoglu Baykal	Hacı Bayram Veli University, Turkey
Federico Mara	University of Pisa, Italy
Himmet Karaman	Istanbul Technical University, Turkey
Özge Öztürk Hacar	Yildiz Technical University, Turkey

From structural to TRAnsformative-change of City Environment: challenges & solutions & perspectives (TRACE 2026)

Workshop Organizers

Maria Rosaria Guarini	Sapienza University of Rome, Italy
Pierluigi Morano	Polytechnic University of Bari, Italy
Francesco Sica	Sapienza University of Rome, Italy
Alejandro Segura	Universidad Politécnica de Madrid, Spain
Yilsy María Nuñez Guerrero	Universidad Politécnica de Madrid, Spain

Workshop Program Committee Members

Maria Cerreta	University of Naples Federico II, Italy
Francesco Scorza	University of Basilicata, Italy

Francesco Calabrò	Mediterranea University of Reggio Calabria, Italy
Felicia Di Liddo	Polytechnic University of Bari, Italy
Marco Locurcio	Polytechnic University of Bari, Italy
Valeria Saiu	University of Cagliari, Italy
Lucia Della Spina	University Mediterranea of Reggio Calabria, Italy
Fabrizio Battisti	University of Florence, Italy
Giuseppe Cerullo	Sapienza University of Rome, Italy
Fabiana Forte	Università della Campania Luigi Vanvitelli, Italy
Marta Dell'Ovo	Politecnico di Milano, Italy
Cristina Coscia	Politecnico di Torino, Italy
Debora Anelli	Sapienza University of Rome, Italy

Urban Artificial Intelligence for Modelling, Analysis & Planning (Urban plAnnIng 2026)

Workshop Organizers

Chiara Garau	University of Cagliari, Italy
Romano Fistola	University of Naples Federico II, Italy
Beniamino Murgante	University of Basilicata, Italy
Dessislava Petrova Antonova	Sofia University, GATE Institute, Bulgaria
Gianni Pantaleo	University of Florence, Italy
Rosa Anna La Rocca	University of Naples Federico II

Workshop Program Committee Members

Paola Zamperlin	University of Florence, Italy
Francesco Zullo	University of L'Aquila, Italy
Gerardo Carpentieri	University of Naples Federico II, Italy
Luciano Alessandro Ipsaro Palesi	University of Florence, Italy
Tazyeen Alam	Università degli Studi di Cagliari, Italy
Reza Askarizad	University of Cagliari, Italy
Mana Dastoum	Polytechnic University of Madrid, Spain
Alfonso Annunziata	University of Basilicata, Italy

Urban Space Accessibility and Mobilities (USAM 2026)

Workshop Organizers

Chiara Garau	University of Cagliari, Italy
Alessandro Plaisant	University of Sassari, Italy
Barbara Caselli	Università degli Studi di Parma, Italy
Mauro D'Apuzzo	University of Cassino and Southern Lazio, Italy
Francesco Pinna	University of Cagliari, Italy
Silvio Nocera	IUAV University of Venice, Italy

Workshop Program Committee Members

Mauro Coni	University of Cagliari, Italy
Martina Carra	University of Brescia, Italy
Tiziana Campisi	University of Enna Kore, Italy
Reza Askarizad	University of Cagliari, Italy
Giuseppe Cappelli	University of Cassino and Southern Lazio, Italy
Floriana Zucaro	University of Naples Federico II, Italy
Gerardo Carpentieri	University of Naples Federico II, Italy
Youssef El Ganadi	International University of Rabat, Morocco
Yiota Theodora	National Technical Univeristy of Athens, Greece
Silvia Rossetti	Università degli Studi di Parma, Italy
Aime Tsinda	University of Rwanda, Rwanda
Paola Zamperlin	University of Florence, Italy
Luciano Alessandro Ipsaro Palesi	University of Florence, Italy
Tazyeen Alam	Università degli Studi di Cagliari, Italy
Mariano Pernetti	Università della Campania Luigi Vanvitelli, Italy
Salvatore Leonardi	University of Catania, Italy
Francesco Galante	University of Naples Federico II, Italy

Virtual Reality, Augmented Reality and Applications (VRA 2026)

Workshop Organizers

Damiano Perri	University of Perugia, Italy
Osvaldo Gervasi	University of Perugia, Italy
Paolo Nesi	University of Florence, Italy

Workshop Program Committee Members

Taihoon Kim	Chonnam National University, South Korea
JungYoon Kim	Gachon University, South Korea
David Berti	ART SpA, Italy

Workshop on Advanced and Computational Methods for Earth Science applications (WACM4ES 2026)

Workshop Organizers

Luca Piroddi	University of Cagliari, Italy
Patrizia Caprizzi	University of Palermo, Italy
Marilena Cozzolino	Università del Molise, Italy
Giuseppina Vacca	University of Cagliari, Italy
Luigi Capozzoli	Hydrogeosite Laboratory of CNR-IMAA, Italy
Nikos Papadopoulos	Foundation for Research and Technology Hellas, Greece

Workshop Program Committee Members

Sebastiano D'Amico	University of Malta, Malta
Enrica Vecchi	University of Cagliari, Italy
Donatella Rita Fiorino	Università degli Studi di Cagliari, Italy
Chiara Garau	University of Cagliari, Italy
Emanuele Colica	University of Malta, Malta
Raffaele Martorana	Università di Palermo, Italy
Raffaele Persico	Università della Calabria, Italy
Andrea Angelini	Institute of Heritage Science of the National Research Council of Italy (CNR ISPC), Italy
Paolo Mauriello	Università del Molise, Italy
Dedalo Marchetti	Italian National Institute of Geophysics and Volcanology, Italy
Veronica Pazzi	University of Florence, Italy
Agnese Innocenti	University of Florence, Italy
Maurizio Ercoli	University of Perugia, Italy
Luciano Galone	University of Malta, Malta
Amr Abd El-Raouf	Zagazig University, Egypt
Peter Iregbeyen	University of Malta, Malta
Pietro Milillo	University of Houston, USA
Tommaso Beni	University of Florence, Italy

Cristiano Fidani	Central Italy Electromagnetic Network, Italy
Deodato Tapete	Italian Space Agency (ASI), Italy
Enzo Rizzo	University of Ferrara, Italy
Guido Luzi	Centre Tecnològic de Telecomunicacions de Catalunya, Spain
Sina Saneiyan	Binghamton University, USA
Alex Sendrós	Universitat de Barcelona, Spain
Mercedes Solla	Universidade de Vigo, Spain
Alessandra Carollo	University of Palermo, Italy
Andrea Vergnano	Università di Torino, Italy
Giorgio Alaia	Università di Perugia, Italy
Giorgio De Donno	Sapienza University of Rome, Italy
Samuel Bignardi	University of Chieti and Pescara, Italy
Athos Agapiou	Cyprus University of Technology, Cyprus
Gregory Tsokas	Aristotle University of Thessaloniki, Greece
Alexandru Hegyi	Oslo University, Norway
Ilaria Barone	Università degli Studi di Padova, Italy
Yuan Qi	University of Hong Kong, China
Carmen Cuenca García	Universitat de València, Spain
Ilaria Catapano	Institute for Electromagnetic Sensing of the Environment (IREA) CNR, Italy
Mariano Lisi	University of Basilicata, Italy
Adrien Dimech	École de Technologie Supérieure, Canada
Davide Campo	Independent Researcher, UK
Chiara Torre	University of Malta, Malta
Anna Forte	University of Bologna, Italy
Adam Gauci	University of Malta, Malta
Vittoria Borghese	TNO, The Netherlands
Andrea Scianna	Institute for High Performance Computing and Networking (ICAR) CNR, Italy
Maria Antonia Brovelli	Politecnico di Milano, Italy
Pietro Milillo	University of Houston, USA
Gregory Tsokas	Aristotle University of Thessaloniki, Greece

Sponsoring Organizations

ICCSA 2026 would not have been possible without tremendous support from many organizations and institutions, for which all organizers and participants of ICCSA 2026 express their sincere gratitude:



Universidade do Minho
Escola de Engenharia

University of Minho, Portugal (https://www.uminho.pt/)



University of Perugia, Italy (https://www.unipg.it)



University of Basilicata, Italy (http://www.unibas.it)

University of Cagliari, Italy (https://unica.it/)

Monash University, Australia (https://www.monash.edu/)

Kyushu Sangyo University, Japan (https://www.kyusan-u.ac.jp/)

Venue

ICCSA 2026 took place at: **the University of Minho campus:**
Campus de Gualtar
4710-057 Braga
Portugal





Additional Reviewers

Amr Abd Elraouf	Zagazig University, Egypt
Lidia Aceto	Università del Piemonte Orientale, Italy
Ana Maria Acu	Lucian Blaga University of Sibiu, Romania
Vera Afreixo	University of Aveiro, Portugal
Akif Akgul	Hitit University, Turkey
Sania Akhtar	Ghulam Ishaq Khan Institute of Engineering Sciences and Technology, Pakistan
Giorgio Alaia	Università di Perugia, Italy
Tazyeen Alam	Università degli Studi di Cagliari, Italy
Arnulfo Alanis	Tecnológico Nacional de México, Mexico
Sandra Aleixo	University of Lisbon, Portugal
Earl Ryan Aleluya	Mindanao State University Iligan Institute of Technology, Philippines
Muhammad Junaid Ali	Université de Haute Alsace, France
Filipe Alvelos	Universidade do Minho, Portugal
Claudia Amadei	University of Padua, Italy
Paulo Amorim	Escola de Matemática Aplicada da Fundação Getulio Vargas, Brazil
Ana Paula Amorim	Universidade do Minho, Portugal
Paola Amoruso	LUM Giuseppe Degennaro University, Italy
Debora Anelli	Sapienza University of Rome, Italy
Mariarosaria Angrisano	Pegaso University, Italy
Alfonso Annunziata	University of Basilicata, Italy
Chiara Anselmi	Università di Pisa, Italy
Shigeyoshi Aoyama	KEK, Japan
Tatsumi Aoyama	University of Tokyo, Japan
Bernady O. Apduhan	Kyushu Sangyo University, Japan
Pablo Luiz Araujo Munhoz	Federal University of Ouro Preto, Brazil
Muhammad Waqas Arshad	University of Bologna, Italy
Daniela Ascenzi	Università degli Studi di Trento, Italy
Laura Ascione	Università degli Studi di Napoli Federico II, Italy
Reza Askarizad	University of Cagliari, Italy
Vanessa Assumma	University of Bologna, Italy
Vittorio Astarita	University of Calabria, Italy
Maria Fernanda Augusto	BITrum Research Center, Spain
Asja Aulisio	Politecnico di Torino, Italy
Samat Baigereyev	D. Serikbayev East Kazakhstan Technical University, Kazakhstan
Marco Baioletti	University of Perugia, Italy
Sara Balderas Díaz	University of Cádiz, Spain

Ginevra Balletto	University of Cagliari, Italy
Serob Balyan	Institute for Informatics and Automation Problems, National Academy of Sciences of Armenia, Armenia
Benedetto Barabino	University of Brescia, Italy
Sebastiano Barbieri	Politecnico di Torino, Italy
Ilaria Barone	Università degli Studi di Padova, Italy
Alice Barreca	Politecnico di Torino, Italy
Socrates Basbas	Aristotle University of Thessaloniki, Greece
Rosaria Battarra	Consiglio Nazionale delle Ricerche, Italy
Silvia Battino	University of Sassari, Italy
Fabrizio Battisti	University of Florence, Italy
Patricia Beites	University of Beira Interior, Portugal
David Berti	Art Spa, Italy
Samuel Bignardi	University of Chieti and Pescara, Italy
Francesco Bimbo	University of Foggia, Italy
Francesca Biondi	Politecnico di Torino, Italy
Marilisa Biscione	CNR, Italy
Alexander Bogdanov	St. Petersburg State University, Russia
Biagio Boi	University of Salerno, Italy
Alessandro Bombini	Istituto Nazionale di Fisica Nucleare, Italy
Ana Borges	Instituto Politécnico do Porto, Portugal
Vittoria Borghese	TNO, The Netherlands
Giuseppe Borruso	University of Trieste, Italy
Kamila Borsekova	Matej Bel University, Slovak Republic
Efstathios Bouhouras	Aristotle University of Thessaloniki, Greece
Ana Cristina Braga	University of Minho, Portugal
Diego Brandão	Federal Center for Technological Education of Rio de Janeiro, Brazil
Maria Antonia Brovelli	Politecnico di Milano, Italy
Francesco Bruzzone	Università Iuav di Venezia, Italy
Francesca Buglione	CNR ISPC, Italy
Enrique M. Cabaleiro Lago	Universidade de Santiago de Compostela, Spain
Francesco Calabrò	Mediterranea University of Reggio Calabria, Italy
Rogerio Calazan	Instituto de Estudos do Mar Almirante Paulo Moreira - Brazilian Navy, Brazil
Lelio Campanile	Università degli Studi della Campania Luigi Vanvitelli, Italy
Tiziana Campisi	University of Enna Kore, Italy
Dario Campisi	University of Perugia, Italy
Rubina Canesi	Università degli Studi di Padova, Italy
Alessandra Capolupo	Polytechnic University of Bari, Italy

Caterina Caprioli	Politecnico di Torino, Italy
Dario Caramiello	University of Naples Parthenope, Italy
Gerardo Carpentieri	University of Naples Federico II, Italy
Martina Carra	University of Brescia, Italy
Barbara Caselli	Università degli Studi di Parma, Italy
Danny Casprini	Politecnico di Milano, Italy
Arcangelo Castiglione	University of Salerno, Italy
Giulia Fernanda Grazia Catania	Università di Enna Kore, Italy
Ilaria Catapano	CNR IREA, Italy
Laura Cavalcante	Instituto Superior de Engenharia do Porto, Portugal
Giulio Cavana	Politecnico di Torino, Italy
Valeria Cera	University of Naples Federico II, Italy
Maria Cerreta	University of Naples Federico II, Italy
Giuseppe Cerullo	Sapienza University of Rome, Italy
Daniele Cesini	CNAF, INFN, Italy
Kai Cheng	Kyushu Sangyo University, Japan
Giuseppe Ciciriello	University Of Naples Federico II, Italy
Birol Ciloglugil	Ege University, Turkey
Cecília Coelho	Helmut Schmidt University, Germany
Enrico Collini	University of Florence, Italy
Luís Conceicao	Instituto Superior de Engenharia do Porto, Portugal
Giovanna Concu	University of Cagliari, Italy
Tatar Corina	University of Oradea, Romania
Alejandro Cornejo	Universitat Politècnica de Catalunya, Spain
Simone Corrado	University of Basilicata, Italy
Florbela Correia	Polytechnic Institute of Viana do Castelo, Portugal
Franco Corti	University of Padua, Italy
Cristina Coscia	Politecnico di Torino, Italy
Lino Costa	University of Minho, Portugal
Eliana da Costa e Silva	Polytechnic Institute of Porto, Portugal
Alessandro Costantini	INFN, Italy
Alessia Crisopulli	Politecnico di Milano, Italy
Valerio Cutini	University of Pisa, Italy
Mauro D'Apuzzo	University of Cassino and Southern Lazio, Italy
Hiroshi Daisaka	Hitotsubashi University, Japan
Sebastiano Damico	University of Malta, Malta
Maria Danese	CNR ISPC, Italy
Gianni Dangelo	University of Salerno, Italy
Mana Dastoum	Polytechnic University of Madrid, Spain

Giulia Datola	Politecnico di Milano, Italy
Regina De Almeida	University of Trás-os-Montes and Alto Douro, Portugal
Elena De Angelis	Politecnico di Torino, Italy
Elise de Doncker	Western Michigan University, USA
Pasquale De Luca	Parthenope University of Naples, Italy
Luiza De Macedo Mourelle	State University of Rio de Janeiro, Brazil
Pierfrancesco De Paola	University of Naples Federico II, Italy
Alexander Degtyarev	St.Petersburg State University, Russia
Francesco Paolo Del Giudice	University of Basilicata, Italy
Piera Della Morte	Pegaso Telematic University, Italy
Federico Dellanna	Politecnico di Torino, Italy
Marta Dellovo	Politecnico di Milano, Italy
Federica Deri	Università di Pisa, Italy
Francesco Di Capua	University of Basilicata, Italy
Clara Di Fazio	University of Trieste, Italy
Gabriella Di Genova	Università degli Studi di Perugia, Italy
Benedetta Di Leo	Università Politecnica delle Marche, Italy
Felicia Di Liddo	Polytechnic University of Bari, Italy
Joana Dias	University of Coimbra, Portugal
Oana Dinu	National University of Science and Technology Politehnica Bucharest, Romania
Lucia Dobrotová	Matej Bel University in Banská Bystrica, Slovak Republic
Edda Donati	University of Rome Tor Vergata, Italy
Aziz Dursun	Virginia Tech University, USA
Zahra Ebadi Ansaroudi	Fondazione Bruno Kessler, Italy
Maurizio Ercoli	University of Perugia, Italy
Dario Esposito	Polytechnic University of Bari, Italy
Noelia Faginas Lago	Università di Perugia, Italy
Maria Irene Falcao	University of Minho, Portugal
Stefano Falcinelli	University of Perugia, Italy
Marco Fanfani	University of Florence, Italy
Marcos Farias	Instituto de Engenharia Nuclear, Brazil
Francesca Fariello	Sapienza University of Rome, Italy
Maria Francesca Faro	Ministry of Education and Merit (MIM), Italy
Eslam Farsimadan	University of Salerno, Italy
Sara Favargiotti	University of Trento, Italy
Celestina Fazia	Università di Enna Kore, Italy
Maksims Feofilovs	Riga Technical University, Latvia
Florbela P. Fernandes Cedri	Instituto Politecnico de Bragança, Portugal
Paula Odete Fernandes	Instituto Politecnico de Bragança, Portugal

Flora Ferreira	University of Porto, Portugal
Maddalena Ferretti	Università Politecnica delle Marche, Italy
Luís Ferrás	University of Porto, Portugal
Massimo Ficco	University of Salerno, Italy
Maros Finka	Masaryk University, Slovak Republic
Ugo Fiore	University of Salerno, Italy
Lorena Fiorini	University of L'Aquila, Italy
Romano Fistola	University of Naples Federico II, Italy
Beatriz Flamia Azevedo	Instituto Politecnico de Bragança, Portugal
Valentina Florio	CNR, Italy
Ali Forooghi	University of Campania Luigi Vanvitelli, Italy
Anna Forte	University of Bologna, Italy
Grigorios Fountas	Aristotle University of Thessaloniki, Greece
Andreas Fricke	Hasso Plattner Institute for Digital Engineering, University of Potsdam, Germany
Luigi Fusco Girard	Pegaso Telematic University, Italy
Francesco Galante	University of Naples Federico II, Italy
Stefano Galassi	University of Florence, Italy
Andrea Gallo	Università degli Studi di Trieste, Italy
Chiara Garau	University of Cagliari, Italy
Gustavo Gatica	Universidad Andrés Bello, Chile
Rachele Vanessa Gatto	Università della Basilicata, Italy
Ivan Gerace	Università degli Studi di Perugia, Italy
Amal Gerges	University of Cagliari, Italy
Osvaldo Gervasi	University of Perugia, Italy
Tasnim Gharaibeh	Kalamazoo College, USA
Gian Felice Giaccu	University of Sassari, Italy
Dario Gioia	CNR ISPC, Italy
Luca Giommi	INFN CNAF, Italy
Valentina Giorgetti	Università degli Studi di Perugial, Italy
Giacomo Giorgi	Università degli Studi di Perugia, Italy
Hiago Mayk Gomes De Araújo Rocha	National Laboratory for Scientific Computing, Brazil
Omar Gomez	Universidad Nacional de Colombia, Colombia
A. Manuela Goncalves	University of Minho, Portugal
Joaquim Goncalves	Polytechnic Institute of Cávado and Ave, Portugal
Pedro Henrique Gonzalez	Federal University of Rio de Janeiro, Brazil
Yusuke Gotoh	Okayama University, Japan
Osvaldo Gramaxo De Freitas	Universitat de València, Spain
Antonia Gravagnuolo	National Research Council Italy (CNR) Institute of Heritage Science, Italy

Teresa Guarda	Universidad Estatal Península De Santa Elena, Ecuador
Maria Rosaria Guarini	Sapienza University of Rome, Italy
Raphael Guerra	Universidade Federal Fluminense, Brazil
Hélia Guerra	University of the Azores, Portugal
Gabriel Guerrero Contreras	University of Cádiz, Spain
Carmen Guida	Università degli Studi di Napoli Federico II, Italy
Eligius Hendrix	Universidad de Málaga, Spain
Jose Manuel Hermida Ramon	Universidade de Vigo, Spain
Antonino Iannuzzo	Universitá degli Studi del Sannio, Italy
Luciano Alessandro Ipsaro Palesi	University of Florence, Italy
Issaku Kanamori	Riken, Japan
Yeliz Karaca	UMass Chan Medical School and Massachusetts Institute of Technology, USA
Himmet Karaman	Istanbul Technical University, Turkey
Tatyana Khachkova	IPGG SB RAS, Russia
Jungyoon Kim	Gachon University, South Korea
Vladimir Korkhov	Saint Petersburg State University, Russia
Iuliia Kozlova	University of Bologna, Italy
Radoslav KožIak	Matej Bel University, Slovak Republic
Ludovica La Rocca	University of Naples Federico II, Italy
Rosa Anna La Rocca	University of Naples Federico II, Italy
Adnane Labbaci	Ibn Tofail University, Morocco
Sabrina Lai	University of Cagliari, Italy
Giuseppe Francesco Cesare Lama	University of Naples Federico Ii, Italy
Isabella Maria Lami	Politecnico di Torino, Italy
Salvatore Leonardi	University of Catania, Italy
Federica Leone	University of Cagliari, Italy
Vita Leonessa	University of Basilicata, Italy
Barbara Lino	Università degli Studi di Palermo, Italy
Vadim Lisitsa	Institute of Petroleum Geology and Geophysics SB RAS, Russia
Marco Locurcio	Polytechnic University of Bari, Italy
Andrea Lombardi	Università di Perugia, Italy
Cristina Lopes	Polytechnic Institute of Porto, Portugal
Isabel Lopes	Instituto Politécnico de Bragança, Portugal, Italy
Rui Pedro Lopes	Instituto Politécnico de Bragança, Portugal
Nicolas Lori	University of the Azores, Portugal
Luís Louro	University of Minho, Portugal
Jader Lugon Junior	Instituto Federal Fluminense, Brazil
Guido Luzi	Centre Tecnològic de Telecomunicacions de Catalunya, Spain

Antonio S. Macias	University of A Coruña, Spain
Daniela Maiullari	Delft University of Technology, The Netherlands
Robert Makin	Western Michigan University, USA
Giorgia Malavasi	Politecnico di Torino, Italy
Francesca Maltinti	University of Cagliari, Italy
Luca Mancini	University of Perugia, Italy
Marcos Mandado Alonso	University of Vigo, Spain
Benedetto Manganelli	Università degli Studi della Basilicata, Italy
Federico Mara	University of Pisa, Italy
Livia Marcellino	University of Naples Parthenope, Italy
Dedalo Marchetti	INGV, Italy
Francesco Marchetti	University of Padua, Italy
Alessandra Marra	University of Salerno, Italy
Barbara Martelli	INFN, Italy
Valerio Martinelli	Università degli Studi di Napoli Federico II, Italy
Ángel Martín Pendás	Universidad de Oviedo, Spain
Alessandra Mascitelli	University of Chieti Pescara, Italy
Michele Mastroianni	University of Foggia, Italy
Luís Miguel Matos	Universidade do Minho, Portugal
Hideo Matsufuru	High Energy Accelerator Research Organization, Japan
Luis Mazon	BITrum Research Group, Spain
Beatrice Mecca	Politecnico di Torino, Italy
Valeria Mele	University of Naples Federico II, Italy
Abdulkadir Memduhoğlu	Harran University, Turkey
Joao Mendes	Instituto Politecnico de Braganca, Portugal
Luis Mendes Gomes	University of the Azores, Portugal
Paolo Mengoni	Hong Kong Baptist University, China
Daniela Menna	Università degli Studi della Campania Luigi Vanvitelli, Italy
Antonio Minervino Amodio	CNR ISPC, Italy
Roberto Minunno	Curtin University, Australia
Fernando Miranda	Universidade do Minho, Portugal
Andrea Montanino	International Centre for Numerical Methods in Engineering, Spain
Andrea Montanino	University of Naples Parthenope, Italy
Ana Monteiro	University of Coimbra, Portugal
Leila Moradi	University of Salerno, Italy
Joao Morais	ITAM, Mexico
Filipe Mota Pinto	Instituto Politécnico de Leiria, Portugal
Ricardo Moura	NOVA University of Lisbon, Portugal

Maria Mourao	Polytechnic Institute of Viana do Castelo, Portugal
Eugenio Muccio	University of Naples Federico II, Italy
Beniamino Murgante	University of Basilicata, Italy
Giuseppe Musolino	Università Mediterranea di Reggio Calabria, Italy
Naohito Nakasato	University of Aizu, Japan
Isabel Natário	Nova University of Lisbon, Portugal
Nadia Nedjah	State University of Rio de Janeiro, Brazil
Grazia Neglia	Pegaso Telematic University, Italy
Paolo Nesi	University of Florence, Italy
Keigo Nitadori	Riken, Japan
Juan José Nogueira	Universidad Autónoma de Madrid, Spain
Mikhail Novikov	Sobolev Institute of Mathematics SB RAS, Russia
Yilsy María Nuñez Guerrero	Universidad Politécnica de Madrid, Spain
Ana Oliveira	Academia Militar, Portugal
Carlo Olivieri	Pegaso University, Italy
Vladimír Ondrejička	Comenius University in Bratislava, Slovak Republic
Beniamino Onorati	Università degli Studi della Basilicata, Italy
Fabian Bernal Orozco Barrera	UNAM, Mexico
M. Fernanda P. Costa	University of Minho, Portugal
António Pacheco	Instituto Superior Técnico, Portugal
Emerson Paiva	University of Minho, Portugal
Simone Palladino	Università degli Studi della Campania L. Vanvitelli, Italy
Francesco Palmieri	University of Salerno, Italy
Simona Panaro	University of Sussex Business School, UK
Claudia Panciera	IRSA CNR, Italy
Giacomo Pannacci	Università degli Studi di Perugia, Italy
Gianni Pantaleo	University of Florence, Italy
Eric Pardede	La Trobe University, Australia
Rosaria Parente	Universitas Mercatorum, Italy
Vittorio Paris	University of Bergamo, Italy
Luca Parodi	University of Genoa, Italy
Rui Pascoal	University of Coimbra, Portugal
Hugo Peixoto	University of Minho, Portugal
Francesco Peluso	University of Naples Parthenope, Italy
Ana Pereira	Polytechnic Institute of Bragança, Portugal
Amina Pereno	Politecnico di Torino, Italy
Martin Perez Perez	University of Vigo, Spain
Donato Perfetto	University of Campania Luigi Vanvitelli, Italy
Güzide Miray Perihanoglu	Van Yuzuncu Yil University, Turkey

Mariano Pernetti	Università della Campania Luigi Vanvitelli, Italy
Damiano Perri	University of Perugia, Italy
Pasquale Perrini	Consorzio Interuniversitario per l'Idrologia, Italy
Raffaele Persico	Università della Calabria, Italy
Lucia Petronio	University of Bologna, Italy
Rossella Petti	Politecnico di Bari, Italy
Francesco Picariello	Universitas Mercatorum, Italy
Angela Pilogallo	IMAA CNR, Italy
Natalia Pingaro	Politecnico di Milano, Italy
Telmo Miguel Pires Pinto	University of Minho, Portugal
Veronika Piscová	Institute of Landscape Ecology, Slovak Academy of Sciences, Slovak Republic
Zeynel Abidin Polat	Izmir Katip Celebi University, Turkey
Giuliano Poli	Università di Napoli Federico II, Italy
Antonio Polimeni	Università di Messina, Italy
Maurizio Pollino	ENEA, Italy
Mohammad Pourfouladi	Politecnico di Milano, Italy
Dmitry Prokhorov	IPGG SB RAS, Russia
Faiza Qamar	National University of Sciences and Technology, Pakistan
Yuan Qi	University of Hong Kong, China
Siyu Qiu	University of Bologna, Italy
Arunasalam Rahunanthan	Central State University, USA
Joelmir Ramos	Universidade do Estado do Rio de Janeiro, Brazil
Bruna Ramos	University of Minho, Portugal
Nicolás Ramos Berdullas	University of Vigo, Spain
Krishnamoorthy Ranganathan	Chennai Institute of Technology, India
Nicoletta Rassu	Università di Cagliari, Italy
Roberta Ravanelli	University of Liège, Belgium
Michela Ravanelli	Sapienza Università di Roma, Italy
Manuela Rebaudengo	Politecnico di Torino, Italy
Simone Rebegoldi	University of Modena and Reggio Emilia, Italy
Pier Francesco Recchi	Università di Napoli Federico II, Italy
Rommel Regis	Saint Joseph's University, USA
Marco Reis	University of Coimbra, Portugal
James Rhodes	Western Michigan University, USA
Luigi Maciel Ribeiro	State University of Rio de Janeiro, Brazil
Corrado Rindone	Università degli Studi Mediterranea di Reggio Calabria, Italy
Ana Maria A. C. Rocha	Universidade do Minho, Portugal
Humberto Rocha	University of Coimbra, Portugal
Carlos Rodrigues	Instituto Politécnico de Bragança, Portugal

Ivan Rodriguez Conde	University of Arkansas at Little Rock, USA
Darina Rojíková	Matej Bel University, Slovak Republic
Diana Rolando	Politecnico di Torino, Italy
Antonella Roma	Sapienza University of Rome, Italy
James Rombi	University of Cagliari, Italy
Elisabetta Ronchieri	INFN CNAF, Italy
Marzio Rosi	University of Perugia, Italy
Silvia Rossetti	Università degli Studi di Parma, Italy
Marco Rossitti	Politecnico di Milano, Italy
Florin Rusca	National University of Science and Technology Politehnica Bucharest, Romania
Yvonne Russo	University of Campania L. Vanvitelli, Italy
Antonio Russo	University of Enna Kore, Italy
Ana Sofia Rézio	Polytechnic Institute of Lisbon, Portugal
Emma Sabatelli	Sapienza Università di Roma, Italy
Matteo Sacchet	University of Turin, Italy
Sabrina Sacco	University of Naples Federico II, Italy
Lucia Saganeiti	Institute of Methodologies for Environmental Analysis, National Research Council, Italy
Valeria Saiu	University of Cagliari, Italy
Francesco Santini	University of Perugia, Italy
Rafael Santos	National Institute for Space Research (INPE), Brazil
Edcarllos Santos	Universidade Federal do Sul da Bahia, Brazil
Valentino Santucci	Università per Stranieri di Perugia, Italy
Andrea Sarcina	Università di Padova, Italy
Ilaria Scarcelli	Politecnico di Bari, Italy
Andrea Scianna	National Research Council of Italy, Italy
Francesco Scorza	University of Basilicata, Italy
Ester Scotto Di Perta	University of Naples Federico II, Italy
Alejandro Segura	Universidad Politécnica de Madrid, Spain
Yasmine Selim	Ain Shams University, Egypt
Ines Sena Cedri	Instituto Politécnico de Bragança and University of Minho, Portugal
Alex Sendrós	Universitat de Barcelona, Spain
Ricardo Severino	University of Minho, Portugal
Mojtaba Sharifzadeh	Università di Trento, Italy
Luigi Sibille	Norwegian University of Science and Technology, Norway
Francesco Sica	Sapienza University of Rome, Italy
Adriano Silva	Instituto Politecnico de Bragança, Portugal
Angela Silva	Universidade do Minho, Portugal

Paula Simões	Military Academy, Portuguese Army, Portugal
Francesca Sinatra	University of Trieste, Italy
Mercedes Solla	Universidade de Vigo, Spain
Sergey Solovyev	Sobolev Institute of Mathematics SB RAS, Russia
Maria Somma	University of Naples Federico II, Italy
Federica Sortino	Università degli Studi di Enna Kore, Italy
Eleonora Sottile	University of Cagliari, Italy
Martina Spinelli	Politecnico di Torino, Italy
Elena Stankova	Saint Petersburg State University, Russia
Alexandra Stankulova	Politecnico di Torino, Italy
Giorgia Sugoni	Links Foundation, Italy
Katarína Sýkorová	Matej Bel University, Slovak Republic
Francesco Tajani	Sapienza University of Rome, Italy
Deodato Tapete	Italian Space Agency (ASI), Italy
Rodrigo Tapia-McClung	Centro de Investigación en Ciencias de Información Geoespacial, Mexico
Maria Tartari	CNR, Italy
Daniela Tavano	University of Calabria, Italy
Tarek Teba	University of Portsmouth, UK
Ana Paula Teixeira	Universidade do Minho, Portugal
Wagner Telles	Universidade Federal Fluminense, Brazil
Luis Teodoro	University of Oslo, Norway
Maria Filomena Teodoro	Instituto Superior Técnico, Portugal
Yiota Theodora	National Technical University of Athens, Greece
Elena Todella	University of Parma, Italy
Graca Tomaz	Instituto Politécnico da Guarda, Portugal
Chiara Torre	University of Malta, Malta
Francesca Torrieri	Politecnico di Milano, Italy
Vincenza Torrisi	University of Catania, Italy
Arianna Travaglini	Lucian Blaga University of Sibiu, Romania
Stefanos Tsigdinos	National Technical University of Athens, Greece
Kamila Turečková	Silesian University in Opava, Czech Republic
Elen Twrdy	University of Ljubljana, Slovenia
Toshihiro Uchibayashi	Kyushu University, Japan
Takahiro Ueda	Juntendo University, Japan
Generoso Vaiano	Pegaso University, Italy
Aleksandar Valjarević	University of Belgrade, Serbia
Gianmarco Vanuzzo	University of Perugia, Italy
Dmitrii Vasiunin	Deutsche Telekom Cloud Services E.P.E., Greece
Mohammad Vazir Panah	University of Salerno, Italy
Anna Vaňová	Matej Bel University, Slovak Republic
Enrica Vecchi	University of Cagliari, Italy

Roberto Ventura	Università degli Studi di Brescia, Italy
Ferdinando Verardi	Pegaso Telematic University, Italy
Clara Stella Vicari Aversa	Università degli Studi Mediterranea di Reggio Calabria, Italy
Gianluca Vinti	Università degli Studi di Perugia, Italy
Katarína Vitálišová	Matej Bel University, Slovak Republic
Weronika Wojtak	CCG/ZGDV Institute, Portugal
Fukuko Yuasa	KEK, Japan
Madina Yussubaliyeva	D. Serikbayev East Kazakhstan Technical University, Kazakhstan
Paola Zamperlin	University of Florence, Italy
Francisco Zayas Gato	Universidade da Coruña, Spain
Nodir Zaynalov	Samarkand Branch of Tashkent University of Information Technologies Named After Muhammad Al Khwarizmi, Uzbekistan
Ljiljana Zivkovic	Republic Geodetic Authority, Serbia
Piero Zizzania	University of Naples Federico II, Italy
Renato Zona	Università della Campania L. Vanvitelli, Italy
Floriana Zucaro	University of Naples Federico II, Italy
Francesco Zullo	University of L'Aquila, Italy
Giovanni Zurlo	INFN CNAF, Italy
Özge Öztürk Hacar	Yildiz Technical University, Turkey
Isidora Đurić	University of Novi Sad, Serbia

Plenary Lectures

Digital Twins for Cities – The Rise of the Socio-Technical System

Claudia van der Laag

Department of Built Environment, Oslo Metropolitan University, Norway

Abstract. Cities are entering a new era in which digital twins are no longer simply technical replicas of urban systems, but evolving socio-technical infrastructures that shape how cities are planned, governed, and experienced. Across the world, urban digital twins are being adopted to address some of the most pressing challenges of our time — climate adaptation, urban resilience, mobility transitions, decarbonisation, infrastructure management, and inclusive governance. Yet as these technologies become increasingly sophisticated, a fundamental question emerges: can digital twins truly support more democratic and human-centred cities, or will they reinforce existing structures of power and expertise? In this keynote, Claudia van der Laag explores the transformation of urban digital twins from data-driven modelling tools into collaborative platforms for urban decision-making and societal change. It argues that the future of digital twins lies not only in advances in artificial intelligence, real-time data integration, simulation, and predictive analytics, but equally in the ability to connect technological innovation with governance, participation, and human agency. Drawing on emerging international research and practical case studies, the keynote examines the complex ecosystem surrounding urban digital twins, bringing together planners, policymakers, researchers, technologists, and communities. It highlights how digital twins increasingly function as meeting points between multiple forms of knowledge — technical, political, social, and experiential — and why interdisciplinary collaboration is essential for their success. The talk critically reflects on the tensions between technological ambition and civic inclusion, exploring themes such as transparency, interoperability, participation, and trust. It also discusses how digital twins can become powerful environments for scenario exploration, urban experimentation, and collective learning, enabling cities to better navigate uncertainty and complexity in an age of rapid urban transformation. Ultimately, this keynote presents a vision of digital twins not as neutral technological systems, but as living socio-technical infrastructures capable of reshaping urban governance itself. The challenge ahead is therefore not only how accurately cities can be modelled, but how inclusively they can be imagined, negotiated, and governed through these emerging digital environments.

Prescriptive AI in Practice: Bridging Optimisation, Generative AI and Agentic Decision-Making

Bernardo Almada-Lobo

Universidade do Porto, Faculdade de Engenharia, INESC TEC, Portugal

Abstract. Prescriptive AI, viewed through the lens of mathematical optimisation and operations research, addresses complex business decision-making problems involving limited resources, multiple decision options and competing objectives. This talk considers the full decision pipeline: problem framing, model formulation, computational solution, implementation, monitoring and organisational adoption. The practical success of prescriptive AI depends less on algorithmic sophistication alone than on the quality of the translation between real-world decision contexts and formal optimisation models. In practice, this translation remains a major bottleneck. Objectives are often ambiguous, constraints are incomplete or tacit, data may be fragmented, and the operational consequences of a recommendation may be poorly understood. As a result, optimisation projects can fail even when the underlying mathematical model is technically sound. The talk will address the persistent gap between optimisation theory and deployment practice. While optimisation has generated substantial value in domains such as supply chain planning, logistics, production, revenue management, workforce planning and inventory control, many initiatives still face a significant risk of under-delivery. Common failure modes include incorrect problem framing, weak data readiness, brittle assumptions, infeasible or non-actionable recommendations, limited explainability, insufficient ownership by business users and inadequate post-deployment monitoring. Closing this gap requires moving from one-off model building to an operational lifecycle in which models are stress-tested, embedded in workflows, monitored for drift and revised as the business environment changes. Recent developments in generative and agentic AI create new opportunities to reduce these barriers. Generative AI can support the elicitation of decision-maker preferences, the structuring of unstructured operational knowledge, the drafting of candidate formulations and the explanation of model outputs in business language. Agentic AI can extend this role by orchestrating multi-step workflows that involve model formulation, solver-code generation, execution, debugging, infeasibility diagnosis, sensitivity analysis and human review. However, these capabilities should not be interpreted as replacing optimisation expertise. Rather, they create a new interface between human judgement, analytical modelling and computational decision support. This talk will discuss a cross-fertilisation agenda between optimisation, generative AI and agentic AI, with a focus on how these technologies should be calibrated to the decision at hand.

Net-Zero Communities through Life Cycle Thinking: Canadian Lessons and Global Implications

Kasun Hewage

School of Engineering (Okanagan Campus), Clean Energy Research Centre (CERC), University of British Columbia (UBC), Canada

Abstract. This keynote presentation discusses how life cycle thinking supports the transformation of communities toward sustainable, resilient, and net-zero futures. Drawing on Canadian research, practice, and policy experience, the presentation highlights how integrated life cycle approaches can move beyond conventional carbon accounting to inform practical decision-making across buildings, infrastructure, energy systems, waste management, water systems, and industrial processes [1–8]. The keynote features research conducted at the Life Cycle Management Laboratory (LCML) at the University of British Columbia, Canada, an internationally recognized research hub led by Kasun Hewage. LCML brings together engineers, scientists, policy researchers, postdoctoral fellows, visiting scholars, and graduate researchers to translate data-intensive sustainability research into actionable climate solutions. The laboratory's work is organized around six interconnected research themes: energy systems, urban development, water systems, waste management, construction and asset management, and industrial processes and products. Together, these themes provide an integrated platform for advancing life cycle–based pathways to net-zero communities. Supported by strong partnerships with government, industry, and community stakeholders, LCML has contributed to a substantial body of applied sustainability research, including more than 400 journal articles, 100 conference papers, and numerous technical reports that have informed real-world planning and implementation. This keynote presentation highlights selected success stories that demonstrate how life cycle thinking can be translated into practical strategies for reducing environmental impacts while strengthening community resilience. These insights are especially relevant to the interdisciplinary focus of ICCSA 2026.

References

1. Weerasinghe, O., Hewage, K., Razi, F., Fatmi, M., Sadiq, R.: Applicability of land use models for sustainable urban planning: a review. J. Urban Plann. Dev. ASCE, **151**(3) (2025). https://doi.org/10.1061/JUPDDM.UPENG-5705

2. Sultana, S. R., et al.: Looking beyond energy efficiency targets: life cycle sustainability of mechanical and water heating equipment in Canadian homes. J. Clean. Prod. Elsevier, **473**, 143469 (2024). https://doi.org/10.1016/j.jclepro.2024.143469
3. Liyanage, D. R., Hewage, K., Hussain, S. H., Razi, F., Sadiq, R.: Climate adaptation of existing buildings: a critical review on planning energy retrofit strategies for future climate. Renew. Sustain. Energy Rev. Elsevier, **199**, 114476 (2024). https://doi.org/10.1016/j.rser.2024.114476
4. Liyanage, D. R., Hewage, K., Karunathilake, H., Sadiq, R.: Feasibility of carbon-capturing in building heating systems: a life cycle thinking-based approach. International J. Greenhouse Gas Control. Elsevier, **132** (2023). https://doi.org/10.1016/j.ijggc.2023.104056
5. Moudgil, V., Sadiq, R., Bakhtavar, E., Paudel, A., Hewage, K.: A cloud-oriented data-analysis framework to analyze peak demand dynamics in institutional building clusters. Sustain. Cities Soci. Elsevier, **111**, 105553 (2024). https://doi.org/10.1016/j.scs.2024.105553
6. Liyanage, D. R., Hewage, K., Ghobadi, M., Sadiq, R.: Thermal resiliency of single-family housing stock under extreme hot and cold conditions. Energy Build. Elsevier, **323**, 114809 (2024). https://doi.org/10.1016/j.enbuild.2024.114809
7. Wijayasekera, S. C., et al.: Sustainability of waste-to-hydrogen conversion pathways: a life cycle thinking-based assessment. Energy Convers. Manage. Elsevier, **270** (2022). https://doi.org/10.1016/j.enconman.2022.116218
8. Jayasena, A., Hewage, K., Siddiqui, O., Sadiq, R.: Socio-economic and environmental cost-benefit analysis of passive houses: a life cycle perspective. J. Clean. Prod. Elsevier, **373** (2022). https://doi.org/10.1016/j.jclepro.2022.133718

Machine Learning Based Methodologies for Grape and Wine Production

Pedro Melo-Pinto

University of Trás-os-Montes and Alto Douro, Engineering Department, Centre for the Research and Technology of Agro-Environmental and Biological Sciences, Portugal

Abstract. Wine production is a major area in the Portuguese economy. The sector accounts for approximately 2.6% of Portugal's GDP, and is crucial to the development of inland regions. Wine quality depends on grape quality, which is a result of different factors such as harvesting at the optimal point of maturity, or selecting the grapes according to quality features. The growing importance of precision agriculture demands novel methods to make the most out of the data generated. Machine learning approaches are particularly suitable for modelling complex processes, where too many arguments are involved. Many of the agricultural processes such as yield estimation or ripening fall into this category. However, several challenges may arise: even if models are available for particular cases, general models are difficult to come by, largely due to the high variability involved; with such models, there is a need for huge amounts of data, in a way which will not always be easy to obtain. This is particularly true in the case of grape production and the wine industry, where the number of grape varieties, climate, soil, etc., introduce a degree of variability that is uncommon in other sectors. One way of obtaining more information in such cases has been the use of images, whether RGB, multi-spectral, or hyperspectral. The use of images introduces additional challenges, such as data size and lighting conditions, which underscores the need for data-driven modelling methodologies. The use of machine learning (ML) techniques combined with images has been establishing itself as a viable alternative for predicting yields or quality parameters in grapes, and assisting with critical harvesting decisions. Several ML approaches have been proposed to handle such data characteristics, but selecting a suitable methodology that best addresses the problem under study and making sure it generalizes well is, most of the time, a cumbersome task. We will address three wine and grape production problems, yield estimation, grape quality assessment, and illness detection, presenting different difficulties. Yield estimation is a major aspect in the wine production industry. The number of berries is one of the most relevant yield aspects, and can be estimated using computer vision, namely image segmentation algorithms. Lighting is one of the key factors in computer vision approaches, and occlusions—which are common in this problem—can be minimized, but require estimation models to achieve reliable results. Early illness detection is, of course, of the utmost importance, with an impact on production and quality. Detecting early symptoms and distinguishing them from other conditions is a difficult task, which makes the use of machine learning algorithms advisable, using automatic visual inspection of the leaves, addressed as a pixel-based classification problem. However, the available datasets are limited and

complex to compile, as different diseases typically require very specific time-frames and locations. The use of generative models to overcome this obstacle, in conjunction with robust classification models, leads to models that are intended to be as general as possible. The evaluation of the grape's maturation profile based on the evolution of their oenological parameters over time is one of the most important aspects for ensuring wine quality, playing a decisive role in the definition of the optimal time for harvesting. The use of hyperspectral imaging combined with learning algorithms represents a faster and non-destructive grape ripeness assessment method. We focus in two fundamental aspects to address the resulting data variability: a possible wavelength bands selection (with the purpose of reducing the dimensionality of data without losing predictive power) and the generalization ability of the ML model under such demanding conditions.

Contents

Blockchain and Distributed Ledgers: Technologies and Applications (BDLTA 2026)

Blockchain as Infrastructure for CBDC: Technological Foundations and Challenges 3
Hu Bitai, Valery Khvatov, Egor Savkov, and Aleksandr V. Bogdanov

From Supply Chain Finance to Paradigm Reconstruction: The Synergistic Evolution of CBDC and Blockchain 22
Hu Bitai, Valery Khvatov, Egor Savkov, and Aleksandr V. Bogdanov

Orchestra-Paxos: A Revision of the Paxos with Dynamic Number of Quorum Participants 39
Mikhail Diagilev and Vladimir Korkhov

Computational and Applied Mathematics (CAM 2026)

A Mach-Number-Weighted Nonlinear Multiscale Viscosity Method for the Compressible Navier–Stokes Equations 55
Elaine Bernine, Isaac P. Santos, and Lucia Catabriga

Dynamics Between Order and Chaos for a Family of Two-Dimensional Cellular Automata 73
Ana Maria Leitão, Ricardo Severino, and Maria João Alves

Appell-Vietoris Polynomials: Structural Properties and Zero Bounds 84
Isabel Cação, M. Irene Falcão, Helmuth R. Malonek, and Graça Tomaz

Interpolation Methods and Cryptographic Primitives in Secret Sharing 99
Thays Rocha Neri Ferreira and Fábio Borges

Computational and Applied Statistics for Data & AI (CAS 2026)

Predicting Periodontal Disease Using Interpretable Machine Learning Models 119
Joana M. A. Lopes, Ana C. Braga, and Marta M. M. Relvas

Modeling Error Propagation Quantization In Neural Networks 133
Jasur Kiyamov, Aleksandr Dik, and Egor Savkov

Email Scanner – A Machine Learning Approach to Staffing/Non-Staffing Automatic Classification 144
Sandra M. Aleixo, Cátia Dias, Artur Ferreira, Mariana Lopes, Inês Prazeres, and Ana Silveira

Forecast Difficulty Detection via Latent Geometry in N-BEATS 163
Lauryane Santos Siqueira, Leonardo Nogueira Matos, Paulo Novais, and Pedro José Costa de Oliveira

Virtuality in Teams: A Multilevel Analysis of Its Effects on Innovation and Satisfaction 181
Isabel Dórdio Dimas, Teresa Rebelo, Paulo Renato Lourenço, and Marta Alves

A Daily-Grid Hazards-to-Dynamics Framework for Inpatient Flow 192
Cecilia Castro, Ana Jacinta Soares, Ana Paula Amorim, and Alexandre Magalhães

Convolutional Autoencoder Compression for LoRa Visual Transmission on Embedded Systems 210
João Mendes, Inês Sena, José Lima, and Ana I. Pereira

Classification of Postal Distribution Centres 224
Sandra M. Aleixo, Célia Fernandes, Paulo Ramos, Margarida Ramires, and Rute Milho

Lightweight Transformer-Based Event Embeddings with Log-Frequency Session Weighting for Purchase Prediction in E-Commerce Clickstream Data 242
Deniz Guzelcay, Sefa Çayır, Konuralp Bol, Halil İbrahim Simsek, Onur Cilgin, and Mehmet S. Aktas

A Modified Bias-Corrected Estimator for the Log-Logistic Shape Parameter ... 260
Ayana Mateus and Frederico Caeiro

A Comparative NLP Study of Semantic and Thematic Changes in AI Research Pre - and Post-generative AI 272
Inês Sena, João Mendes, and Ana I. Pereira

Customer Churn Prediction in the Telecom Industry Via a Clustering-Random Forest-XGBoost Approach: Discovering the Role of Delta in PCA Loadings 289
Chien-Sing Lee and Sharini Sivadas Nanda

A Multiclass Model for Financial Distress and Bankruptcy Prediction 308
Ana Sousa, Ana C. Braga, and Jorge Cunha

Simulation-Based Data Generation for On-Demand Bus-as-a-Service Analytics 327
Vânia Sousa, Martim Santos, Bruno Silva, Yanelys Llerena, Tiago Sá, António C. Vieira, and Maribel Y. Santos

Survival Analysis of HIV Seroconversion Using Interval-Censored Data: Nonparametric and Regression Approaches 344
Rui Alves, Filipa Pinto, Aurélio Sidumo, Carla Moreira, Luís Meira-Machado, Paula Meireles, Miguel Rocha, and Maria João Novais

Experimental Limits of Machine Learning on Near-Random Data: A Case Study of the Romanian Lottery System 358
Alexandru Dinu, Sara-Gabriela Modan, and Mădălin Frunzete

The Generalized Jackknife and the Extremal Index Estimation 372
Dora Prata Gomes and M. Manuela Neves

Circular Economy and Territorial Competitiveness (CETC 2026)

Measuring Urban Competitiveness: An Application at the Urban Districts Level in the Galicia Cities of Vigo and Coruna (Spain) 387
Gerardo Carpentieri, Carmen Guida, Manuel Fernandez-Grela, Yolanda Pena Boquete, and Jackiline Salazar-Torres

Mapping Territorial Competitiveness and Circular Economy Through Inclusive Management: A Bibliometric Analysis and Case Study of the CirculEC Project 402
Madina Yussubaliyeva, Zhadyra Konurbayeva, Madina Rakhimberdinova, Oxana Denissova, Allfiya Zakimova, and Darya Surova

Region-Focused Architectural Design in the Context of the Circular Economy 419
Oxana Gurskaya, Yelena Feoktistova, Tatiana Inozemtseva, and Madina Yussubaliyeva

Creative Destruction and Regional Resilience: The Case of the Volvo Investment in the Košice Region 437
Silvia Ručinská, Lenka Jakubócová, and Ondrej Mitaľ

Circular Ecological Heritage Communities: Regeneration and Evaluation Strategies for Historic Villages, Landscapes and Rural Territories (CHEC 2026)

Regenerating Historical, Landscape, and Environmental Heritage Through the Municipal Territorial Plan in Lombardy 455
Dina Jovanovic, Martina Mulinacci, Davide Cerati, and Daniela Oreni

Co-creating Circular Tourism Action Plans in Remote Heritage Sites: Evidence from Three Scottish Castles 471
Antonia Gravagnuolo and Patrizia Riganti

Circular Regeneration Pathways for Rural Heritage: The Role of Adaptive Reuse Across 18 European Villages 488
Piera Della Morte

Heritage-Led Regeneration and Multidimensional Impact Assessment in a Historic Coastal City: The Case of the Community Cooperative *"Identity and Beauty"* in Sciacca, Sicily (Italy) 506
Antonia Gravagnuolo, Francesca Buglione, and Mariarosaria Angrisano

Strategic Design for Energy Self-sufficiency in Historic Villages: Renewable Energy Communities and Circular Regeneration in Grecìa Salentina 523
Grazia Neglia, Ferdinando Verardi, and Ippolita Mecca

Co-designing Alternatives of Adaptive Reuse: The Case of the Church of the Holy Name of Jesus in Moncalieri 541
Maria Tartari, Giulio Cavana, Mariarosaria Angrisano, Gravagnuolo Antonia, Luigi Fusco Girard, and Marta Bottero

A Multiscale Methodology for the Circular Regeneration of Historic Villages, Cultural Landscapes and Rural Territories 560
Albina Scioti and Mariarosaria Angrisano

Cyber Intelligence and Applications (CIA 2026)

Enabling AI-Driven Cybersecurity Through Reproducible Cloud-Native Security Architectures 575
Gianni D' Angelo, Carmine Ambrosino, Simona Ettari, Alfonso Petrosino, and Francesco Palmieri

LADDM: A Flexible Tool for Log Anomaly Detection 592
Gianni D'Angelo, Raffaele Esposito, Mariano Caiazzo, and Francesco Palmieri

Analyzing Ensemble Stacking Technologies for Supervised Anomaly Detection 607
Arcangelo Castiglione, Gianni D' Angelo, Francesco Palmieri, and Antonio Scarfó

Toward Secure and Intelligent Subject Authentication Using Brainprints and Enhanced AMEEGNet 624
Adnan Elahi Khan Khalil, Jesus Arturo Perez-Díaz, José Antonio Cantoral-Ceballos, and Javier M. Antelis

Optimizing Intrusion Detection for SDN-Enabled IoT Using of Lightweight Deep Learning Architectures 641
Mustafa T. Saleh and Ali H. Hamad

Author Index 657

Blockchain and Distributed Ledgers: Technologies and Applications (BDLTA 2026)

Blockchain as Infrastructure for CBDC: Technological Foundations and Challenges

Hu Bitai[1], Valery Khvatov[1], Egor Savkov[2], and Aleksandr V. Bogdanov[2(✉)]

[1] St. Petersburg State University, 7/9 Universitetskaya Emb., St. Petersburg 199034, Russia
[2] DGT Technologies AG, Toronto, Canada
a.v.bogdanov@spbu.ru

Abstract. Central bank digital currencies (CBDC) are now in pilot or production in more than 130 jurisdictions [1], and the choice of underlying ledger architecture has become a first-order policy question. This paper examines blockchain as a candidate infrastructure for CBDC. We compare centralised, permissioned-distributed, and hybrid ledger designs along the dimensions of throughput, consensus mechanism, privacy, and governance. Three strands of recent technical work are surveyed and located in this design space: zero-knowledge proofs together with the Rayls-style controlled-anonymity protocol; programmable smart contracts for monetary-policy logic; and post-quantum cryptography for long-horizon security. We then quantify four operational scenarios using the AB2023 tokenomic model of Khvatov and Bogdanov [19], reimplemented in Python and calibrated to published F-BFT measurements [20]. Across a 100-month horizon, our simulation yields consortium throughput between 28 and 68 CTPS depending on regime, traces the Gini coefficient of token holdings from 0.89 to 0.20 (−77%) under the gradual-decentralisation scenario, and shows that doubling the slashing penalty approximately halves attack-success probability below the BFT 1/3 threshold but is ineffective above it. The throughput model agrees with [20] to within 12% mean relative error for cluster sizes $n \leq 18$. Two implications follow. Pure node-count expansion does not buy decentralisation cheaply: throughput dips by roughly 8% during the multi-cluster transition before sharding and H-NET segmentation [21] recover and then exceed the baseline. And economic deterrence has a hard ceiling at the BFT threshold, so technical Sybil resistance remains non-substitutable. The contribution of this paper is therefore a parameterised, empirically anchored vocabulary for techno-economic discussion of CBDC infrastructure, rather than a new architecture in its own right.

Keywords: CBDC · Blockchain · Distributed ledger technology · Smart contracts · F-BFT consensus · AB2023 model · Tokenomics simulation

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 3–21, 2027.
https://doi.org/10.1007/978-3-032-30518-3_1

1 Introduction

1.1 Motivation and Global Context

As of 2025, more than 130 countries and regions have been exploring or piloting central bank digital currencies (CBDC) worldwide [1]. Notable cases include China's e-CNY (more than 300 million users), The European Central Bank's digital euro project, Nigeria's eNaira, The Bahamas' sand dollar and Eastern Caribbean DCASH. The current global momentum indicates that the spread of money's digitalisation needs to rely on strong, expandable, and reliable infrastructure.

Among the candidate technologies, blockchain and distributed ledger technology (DLT) have garnered considerable attention due to their characteristics of decentralisation, immutability, traceability and programmability [2].

1.2 Research Questions and Contributions

This paper addresses four research questions:

1. How can blockchain serve as infrastructure for CBDC, and what architectural options does it open up?
2. Which technical capabilities – privacy protection, programmability, and post-quantum security – have been demonstrated in current pilots?
3. Which technical bottlenecks (performance, post-quantum threats, consensus attacks) limit blockchain-based CBDC, and how do recent proposals address them?
4. How can a tokenomic model help quantify the techno-economic trade-offs that arise once design choices are made?

Our contributions are:

- A side-by-side analysis of three CBDC ledger architectures – centralised, permissioned-distributed, and hybrid – along throughput, governance, privacy, and interoperability axes (Sect. 2; Table 1).
- A survey of privacy-preserving techniques for CBDC, organised around the Rayls controlled-anonymity protocol [10,11] and the post-quantum migration path identified by NIST [9,17] (Sect. 3).
- A four-scenario quantitative study of the AB2023 tokenomic model of Khvatov and Bogdanov [19]. We re-implement and parameterise the framework in Python under four operational regimes, validate its F-BFT throughput function against published cluster measurements [20], and report Gini-coefficient trajectories, throughput envelopes, and Monte Carlo collusion-attack outcomes across 100 simulated months (Sect. 5; Tables 4–8; Figs. 2–4).
- Three design implications drawn from the simulation: an ∼8% throughput dip in the multi-cluster transition window, a hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold, and the observation that horizontal scaling via sharding and H-NET segmentation [21] is the binding constraint when retail uptake grows.

– A position statement on AI–CBDC integration that locates it at the operational-monitoring layer rather than the consensus layer (Sect. 6).

2 Blockchain Architecture Choices for CBDC

CBDC architectures span a continuum from fully centralised to fully decentralised designs [3,4]; the choice trades off settlement performance, governance control, and cross-border interoperability.

2.1 Centralized Ledger Architecture

Centralized architectures meet the requirements of high-throughput systems with quick settlement finality. Centralised ledgers of China's e-CNY can handle more than 100,000 transactions per second and ensure sub-second finality [5]. The system uses a two-layer structure; the central bank manages the main ledger, and commercial banks' operating institutions handle users' wallets and provide front-end services. It has full jurisdiction over monetary policy and is simple, showing strong performance.

2.2 Permissioned Distributed Ledger

On the other hand, permissioned DLTs allow multiple independent parties to agree upon each transaction independently. The BIS-led mBridge project integrates central banks and commercial banks of various countries via a permissioned DLT, which is based on Corda. It achieves cross-border settlement in under 10 s, compared to 2–5 days for correspondent banking, while enhancing transparency and reducing reconciliation costs [7]. The consensus mechanism uses a variation of Practical Byzantine Fault Tolerance (PBFT) with a limited set of trusted nodes.

2.3 Hybrid Architectures

The combination of a centralized retail system and a dispersed wholesale network via DLT technology. Many central banks have taken this approach; for example, e-CNY (using a centralised ledger system in retail while exploring DLT for cross-border features), and the ECB has also considered its own digital euro using both methods – retaining control through traditional systems but adopting DLT across clearinghouses – [6].

The IMF proposes the following three routes: [8]

– Compatible ledger: The central bank's RTGS system remains unchanged; bridges connect to external DLT platforms for specific applications.
– Single ledger: Tokenised reserves and other tokenised assets are stored in the same DLT system to ensure atomic settlement.
– Separate Ledgers: Delegating the task of ledgering to another entity and remaining responsible for issuing the currency.

2.4 Comparison and Trade-Off Analysis

Table 1 shows the main compromises. Based on the purpose of selection policies, retail inclusion and high-throughput prefer centralised or mixed types; cross-border multi-jurisdictional settlement favours distributed forms.

Table 1. Technological trade-off matrix of CBDC architectures.

Design Parameter	Centralized Ledger	Permissioned DLT	Hybrid Architecture
Throughput	High	Medium	High
Transparency	Low	High	Medium
Governance	Strong Central Bank Control	Shared	Mixed
Cyber Risk	Single-point failure	Byzantine fault tolerance	Mixed risk
Settlement Finality	Instant	Program-dependent	Tiered
Interoperability	Limited	High	Medium

2.5 Consensus Mechanisms for Permissioned DLT

Permissioned DLTs commonly rely on Byzantine fault-tolerant (BFT) consensus. Hyperledger Fabric uses PBFT and reaches finality in under two seconds for networks of up to 100 nodes. Tendermint (Cosmos) and HotStuff (Diem) extend the design with chained leader rotation, improving scalability beyond the classical PBFT $O(n^2)$ message-complexity ceiling. CBDC consortia typically operate with high mutual trust among validators, which means a simpler BFT family—F-BFT, IBFT, or hardened PBFT—is usually sufficient for the throughput envelope that retail and wholesale settlement require.

2.6 Performance and Scalability of Hybrid Designs

Recent pilot data illustrate the scalability gap. Hybrid retail-wholesale designs have handled peaks above 10,000 transactions per second in the e-CNY rollout [5]; cross-border settlement on the mBridge platform completes in roughly 10 s per transaction [7]. The gap is not coincidental: per-transaction latency on a fully distributed ledger grows with the geographical and organisational dispersion of validators, and the trade-off worsens superlinearly as transaction volume scales. Sharding and layer-2 constructions are the standard mitigations, and we return to their quantitative effect on consortium throughput in Sect. 5.

3 Key Technological Capabilities

3.1 Privacy Computing and Controlled Anonymity

"Controlled anonymity" balances users' privacy with AML/CFT requirements. Privacy-preserving technologies such as zero-knowledge proofs (ZKP), multi-party computation (MPC) and oblivious transfer can protect transactional privacy while ensuring regulatory enforcement [9].

Zero-Knowledge Proofs: Principles and Variants. ZKPs enable a prover to convince a verifier that a statement is true without revealing any other details. ZKPs can also show whether a transaction meets certain conditions (e.g., the user has sufficient funds) while keeping private keys hidden.

Two main families are relevant:

- **zk-SNARKs** (succinct non-interactive argument of knowledge): They generate very small proofs (around hundreds of bytes) and can be verified quickly; however, they require a trusted setup. Examples include Groth16 and Plonk.
- **zk-STARKs** (scalable transparent arguments of knowledge): They eliminate the trusted setup and are post-quantum secure, but produce larger proofs (tens of kilobytes) and are slower to verify.

Hybrid constructions combining the two families have appeared in the literature, attempting to retain SNARK proof size with STARK transparency; the engineering trade-offs remain under active study.

The Rayls System. The Rayls system, deployed at a major clearinghouse, has achieved high throughput and private CBDC transactions. It combines zk-SNARKs to verify transaction authenticity and operates in a permissioned network achieving more than 1,000 transactions per second with sub-two-second finality [10,11].

Tiered Anonymity Models. A typical design pattern is hierarchical anonymity: small transactions enjoy strong privacy, while large ones trigger identity verification. Sweden's e-krona pilot and the ECB's digital euro adopt graded variants of this pattern. Cryptographic constructions such as range proofs and selective disclosure ensure that ordinary user activity remains private from regulators while suspicious patterns above declared thresholds remain inspectable [12] (Fig. 1).

3.2 Smart Contracts and Programmable Money

Conditional automatic processing of money-related functions through smart contracts. Welfare subsidies may specify certain goods or periods of use to prevent fraud and achieve targeted assistance [13]. The tokenised reserve, together with smart contracts, achieves atomic delivery-vs-payment (DvP) for settlement of securities transactions.

Project Pine and Programmable Monetary Policy. Project Pine, co-operated by the BIS Innovation Hub and the Federal Reserve Bank of New York, implemented programmed monetary-policy instruments through smart contract technology. It automatically performed reverse-repo operations in a simulation environment, responded promptly to liquidity shocks by activating emergency loan facilities, and handled asset purchases immediately [14].

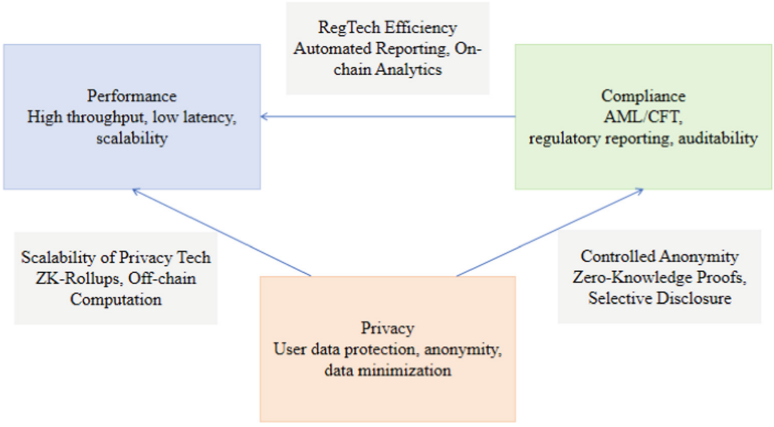

Fig. 1. CBDC Design Triangle (Privacy—Compliance—Performance).

Formal Verification of Smart Contracts. Smart-contract vulnerabilities have caused substantial losses in decentralised finance—the DAO incident (2016) and the Wormhole bridge exploit (2022) are paradigmatic examples [17]. Common defect classes are re-entrancy, integer overflow, timestamp dependency, and access-control flaws. The CBDC setting raises the stakes considerably: a single contract bug at retail scale can propagate into systemic settlement risk, so verification cannot be optional.

Several approaches exist:

- Model checking: Exhaustively explores states of a contract to verify properties. Tools like Manticore and Oyen are automated.
- Proof-checking via interactive theorem provers; for example, Coq and Isabelle. KEVM provides an executable semantics of the EVM for formal verification.
- Symbolic execution: Analyses all possible execution paths. Mythril, Slither.io and other tools are widely available.

For CBDC contracts, none of the three approaches alone is sufficient. A defensible deployment process combines model checking on the high-stakes execution paths, machine-checked proofs for the core monetary-policy logic, and symbolic-execution sweeps for property-violation detection, supplemented by an external bug-bounty programme over an extended testing window.

3.3 Technology Development Support

Global blockchain patents mainly focus on privacy protection, data sharing and high-performance transaction systems [12,14]. Based on a 2023 patent map study, about 60% of blockchain applications involving finance have been filed recently; there has been a significant rise in privacy protection and performance improvement technologies since 2021 [15].

Standardisation efforts are coordinated across central banks and standard-setting bodies: the ISO 20022 messaging standard is being extended to cover digital currency, and the BIS Innovation Hub runs cross-border CBDC interoperability programmes [16].

4 Technical Challenges and Research Frontiers

4.1 Performance Bottlenecks

Public blockchains such as Bitcoin and Ethereum each achieve at most a few tens of transactions per second. Permissioned blockchains improve their consensus algorithms and can process more than ten thousand transactions per second in some situations; nevertheless, they cannot be used nationwide for a massive number of retail purchases from this point on [5].

Sharding and Layer 2. Sharding partitions the network into parallel chains, each processing a subset of transactions independently. Ethereum's sharding roadmap targets 100,000 tps. Layer-2 constructions (rollups in particular) batch transactions off-chain and post only succinct proofs to L1. ZK-rollups have been demonstrated at over 10,000 tps with sub-second user-perceived finality. Two costs limit the gain. Trusted setup is one. The other is proof-generation overhead, which currently requires GPU acceleration or dedicated hardware to meet retail-payment latency budgets at scale.

4.2 Security Risks

Post-Quantum Threats. Commonly used asymmetric protocols (e.g., RSA, ECC) are vulnerable to quantum attacks. The NIST standardisation process has selected lattice-based (ML-KEM, ML-DSA) and multivariate cryptographic systems for replacement. However, their larger key sizes (up to 40× larger) pose integration challenges for constrained CBDC wallets [17]. Post-quantum signature lengths can reach 2–3 KB; thus extra storage space needs to be configured in the system. It may also impair offline payment functions or mobile-wallet capabilities.

Consensus Attacks. Permissioned blockchains do not eliminate collusion risk. Coordinated validator behaviour above the BFT threshold ($p \geq 1/3$) compromises consensus deterministically; below the threshold, attack viability depends on the cost-benefit ratio between economic stakes and slashing penalties. Permissioned DLTs operate with a small set of trusted nodes, which reduces the Sybil surface but raises the impact per compromised node. Hybrid CBDC architectures with central-bank veto authority lower this risk further at the cost of partial decentralisation. Section 5 reports a quantitative attack-economics analysis (Scenario D) that locates the deterrence boundary numerically.

Smart Contract Vulnerabilities. The vulnerability classes catalogued in Sect. 3.2.2 (re-entrancy, arithmetic overflow, timestamp dependence, access-control failures) appear repeatedly in CBDC pilot post-mortems. Formal verification and security audits are necessary but not sufficient: machine-checked proofs cover only the properties that have been specified, and symbolic-execution tools miss bugs that lie outside their state-space coverage. For a CBDC deployment the practical approach is layered: formal verification of core monetary logic, automated static analysis on every commit, an external audit for each major release, and a long-running bug-bounty programme attached to a testnet that mirrors production traffic.

4.3 Interoperability

In the future, there will be many types of CBDCs, private stablecoins and legacy payment systems. Solving the problem of intercommunication among different types of ledgers is therefore critical.

Cross-Chain Protocols. Standardisation based on ISO 20022 data format is the topic area of this project [16,18]. Cross-chain protocols can be divided into:

- Notary Schemes: An intermediary confirms transactions on one chain and issues a signature on another. Simple yet centralised.
- Side chains/Relays: Chain header verification enables trustless cross-chain transfer. Examples include Polkadot and Cosmos.
- Atomic swaps based on hash-time locked contracts; peer-to-peer exchange.

CBDCs may use either notary or side-chain mechanisms for their own decentralisation, with the central bank serving as a trusted notary for international cross-border payments.

The mBridge Approach. The mBridge platform takes a different approach. It replaces inter-ledger bridges with a single permissioned ledger shared by participating central banks, each operating a validator node within a Corda-based consortium [7]. Cross-border settlement on this design avoids the trust assumptions and reconciliation overhead of bridge-based interoperability altogether: a transaction that would traditionally cross multiple ledgers is committed once, on a ledger jointly governed by the participating monetary authorities. The architectural cost is the requirement that all participants agree on the validator set and the upgrade path; the architectural benefit is that governance integration becomes a precondition rather than an afterthought.

Without effective interoperability across these designs, seamless cross-border payment will not be achievable, and CBDC deployments risk fragmenting into isolated payment islands rather than coalescing into a global fabric.

4.4 Summary of Challenges

The five challenge categories surveyed in 4.1–4.3 are summarised in Table 2, together with severity, current solutions, and active research frontiers. Two patterns are worth noting. Performance bottlenecks and consensus attacks have well-established mitigations (sharding, layer-2, economic penalties) but trade off against each other in practice—increasing throughput typically expands the attack surface. Quantum threats and cross-chain interoperability remain open problems whose timelines depend on external factors (NIST standardisation cadence, ISO 20022 adoption). The performance and consensus-attack dimensions, in particular, are the two challenges most amenable to quantitative analysis from a tokenomic-modelling perspective. Section 5 takes up exactly those dimensions, using the AB2023 framework of Khvatov and Bogdanov [19] to quantify throughput trajectories under three deployment regimes and the boundary at which economic deterrence ceases to operate.

Table 2. Assessment of CBDC Technical Challenges and Countermeasures.

Challenge	Specifics	Severity	Solutions	Frontiers
Performance	Throughput, latency, energy	High	Sharding, layer-2, hybrid	ZK-rollups, hardware acceleration
Quantum threats	Asymmetric crypto vulnerability	High	Post-quantum crypto	Lattice-based, multivariate crypto
Consensus attacks	Collusion, Sybil	Moderate	Economic penalties, node authentication	Verifiable random functions
Smart contract	Code defects	High	Formal verification, auditing	ZK-proof verification
Cross-chain	Inconsistent standards	High	ISO 20022, cross-chain protocols	Universal cross-chain communication

5 AB2023 Tokenomic Model: Implementation and Quantitative Results

To examine how design choices in a hybrid blockchain-based CBDC propagate into measurable outcomes, this section reports a controlled simulation study based on the AB2023 tokenomic model formalised by Khvatov and Bogdanov [19]. We first describe the model components we re-implemented and the calibration procedure (5.1). We then specify four parameter regimes (5.2) and

report time-series and comparative results (5.3–5.4). Finally, we present a validation against published F-BFT measurements [20] and a sensitivity analysis (5.5–5.6).

Tokenomic modelling for blockchain-based monetary systems has accumulated a small but heterogeneous literature, with each entry foregrounding a different facet of the design problem (Table 3). Pricing dynamics, algebraic semantics, governance taxonomies, and design guidelines have all been treated separately. AB2023 is the only entry calibrated to a permissioned consortium consensus protocol – F-BFT – and parameterised for the joint trajectory of nodes, users, and token holdings over a multi-year horizon, which is what makes it operative for the questions a CBDC operator actually faces.

Table 3. Comparison of tokenomic models for blockchain-based monetary systems.

Model	Focus	Key features
Cong, Li & Wang (2021) [22]	Dynamic pricing	Equilibrium from transactional demand; network externalities
Letychevsky et al. (2019) [23]	Algebraic verification	Formalisation via insertion modelling; algebraic verification of token-economy properties
Jurjens et al. (2022) [24]	Decentralised incentivisation	Token frameworks for Data Spaces; technical, behavioural, value, coordination, and archetype dimensions
Lesavre, Varin & Yaga (2021) [25]	Token design	NIST IR 8301: design and management guidelines for blockchain-network tokens
AB2023 [19]	Hybrid permissioned-blockchain CBDC	100-month time-series simulation; F-BFT throughput model calibrated to [20]; four decentralisation indicators (ANI, NDI, CII, TDI) with composite D-metric; Monte Carlo collusion-attack module

5.1 Model Implementation

The AB2023 framework as specified by Khvatov and Bogdanov [19] couples seven modules: (i) a Richards-type node-growth process; (ii) a logistic-Gompertz-node-driven (LGND) user-growth process; (iii) a token-supply rule; (iv) a token-price stochastic process; (v) a token-distribution dynamic targeting an evolving wealth concentration; (vi) decentralisation indicators (ANI, NDI, CII, TDI) aggregated into a composite metric D; and (vii) a throughput model rooted in the F-BFT consensus protocol [20,21]. We re-implemented modules (i)–(iii) and (v)–(vii)

in Python (NumPy/Pandas/Matplotlib only, ≈550 lines), following the equations of [19] faithfully where they are specified. Module (iv) is present in our code but is not reported in this paper: market-dynamics analysis lies outside the computational-science focus we adopt here, which centres on consensus throughput and on decentralisation as functions of network structure. Cross-cluster network latency is omitted entirely. Where the throughput model intersects the experimental regime reported in [20] – Docker-based clusters of 6 to 21 nodes – our predictions agree with the published measurements to within 12% relative error for $n \leq 18$ (Table 7). The implementation is reproducible from a fixed random seed (20240508).

The throughput module is anchored to the asynchronous-stream measurements of Bogdanov et al. [20], which report CTPS in the range 0.6–1.7 for cluster sizes $n \in \{6, 12, 15, 18, 21\}$. Outside this range we extrapolate with a federated multi-cluster model: N nodes are partitioned into $\lceil N/15 \rceil$ clusters with arbiter-ring coordination cost $1/(1+(k/k^*)^2)$ at $k^* = 12$. A hardware multiplier of 5.5× projects from the Docker test bed to production infrastructure, consistent with the 5–6× H-NET latency reduction in [21]; sharding adds $0.75 \times (n_{\text{shards}} - 1)$.

5.2 Scenarios and Parameter Settings

We retain the four illustrative scenarios from [19] and instantiate each with parameter values listed in Table 4. Each scenario runs for 100 monthly steps.

Table 4. Scenario parameter settings. N_0 initial nodes; $N_{\max}$ carrying capacity; r_0 baseline node-growth rate; U_0, U_∞ initial and saturating users; T_{sh} sharding activation month; T_{seg} H-NET segment activation month; p malicious-node share (D only).

Parameter	A: Centralised	B: Decentralising	C: High adoption	D: Attack
N_0, $N_{\max}$	100, 115	100, 600	100, 320	inherits B at month 60
r_0	0.025	0.07	0.05	—
U_0, U_∞	10 K, 80 K	10 K, 200 K	10 K, 1.2 M	—
Consortium share (initial)	0.65	0.65	0.45	—
Wealth concentration (target)	0.78	0.55	0.62	—
T_{sh}, T_{seg}	none	42, 66	50, none	—
p sweep	—	—	—	{0.05, 0.10, . . . , 0.40}
Penalty regimes	—	—	—	1× standard, 2× slashing

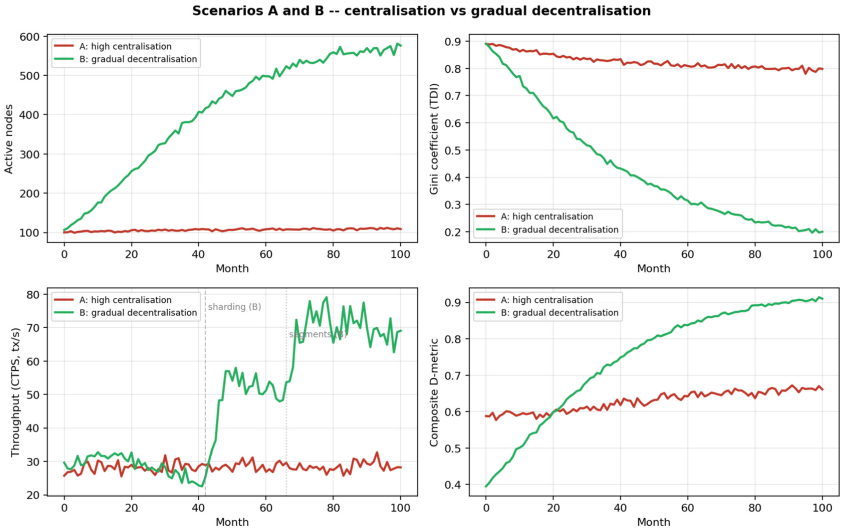

Fig. 2. Four-panel comparison of Scenarios A and B: (a) active nodes over time; (b) Gini coefficient; (c) CTPS, with sharding and segment activation marked; (d) composite D-metric. (Color figure online)

5.3 Time-Series Results: Scenarios A, B, C

Scenario A maintains the network at 108 ± 4 active nodes throughout the simulation (Fig. 2a). The Gini coefficient for the token distribution decays slowly from 0.89 to 0.80 (-10%); the Consortium Influence Index (CII) holds near 0.60. CTPS remains in the 26–32 band. The composite D-metric rises modestly from 0.59 to 0.66. This regime models a retail CBDC such as e-CNY in its current operational form: stable performance, persistent governance concentration, and limited dispersion of holdings.

Scenario B drives the network from 100 to 575 active nodes over 100 months (Fig. 2a, green curve). The Gini coefficient falls from 0.89 to 0.20 (-77%), and CII falls from 0.65 to 0.11 (-83%). CTPS traces the predicted trade-off (Fig. 2c, lower-left): an initial plateau of $\sim$30 CTPS during the high-throughput cluster regime (months 0–24), followed by a slow decline to $\sim$27 CTPS at month 36 as arbiter-ring overhead grows. Sharding activation at month 42 lifts CTPS to $\sim$54 within six months, and H-NET segment activation at month 66 brings it to $\sim$73 by month 72, settling at 68 by month 100. The composite D-metric rises from 0.40 to 0.91 ($+128\%$), confirming that node-count expansion translates into substantive decentralisation rather than nominal participation. Selected milestones are summarised in Table 5.

Scenario C tests the effect of demand-side pressure on network scalability. The user base follows an S-curve from 50 K to its 1.32 M ceiling around month 50 (Fig. 3). Despite the network expanding to 301 nodes by month 100, throughput tracks Scenario B at the same node count ($\sim$52 CTPS), confirming that

Table 5. Scenario B selected milestones.

Month	Active Nodes	Gini (TDI)	CII	CTPS	D-metric
0	106	0.89	0.65	29.6	0.40
24	282	0.58	0.42	28.8	0.63
36	381	0.45	0.34	27.2	0.73
48	461	0.37	0.28	54.0	0.80
60	498	0.31	0.23	49.4	0.84
72	538	0.27	0.19	73.4	0.87
100	575	0.20	0.11	67.7	0.91

the throughput function is governed by node count rather than by user load. The result indicates that high user uptake places pressure on horizontal scaling (sharding, segmentation) rather than on consensus per se, motivating the architectural patterns discussed in Sect. 4.1.

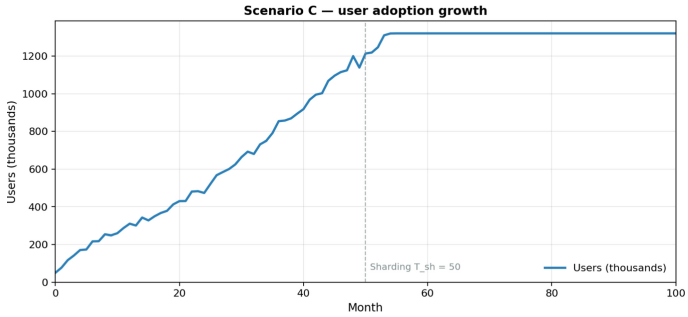

Fig. 3. Scenario C user adoption: monthly active users following the S-curve over 100 months.

5.4 Coordinated Collusion Attack: Scenario D

We freeze the Scenario B network state at month 60 (498 active nodes, 425 K users, mature token supply) and sweep the malicious-node share p from 0.05 to 0.40 under two penalty regimes. For each (p, penalty) pair we run 30 000 Monte Carlo rounds. Each round draws 50 validators by hypergeometric sampling, declares per-round consensus compromised when malicious validators reach the BFT threshold $\lceil n/3 \rceil$, and tests anomaly detection with probability $0.20 + 0.6\,p \times$ penalty. The compounded attack-success probability accounts for repeated attempts before detection closes the window. Attacker cost equals stake locked in malicious nodes plus an opportunity-cost adjustment proportional to the penalty multiplier.

Results are reported in Table 6 and Fig. 4. Three observations follow.

First, below the BFT threshold ($p < 1/3$), attack success scales steeply but predictably with p. At $p = 0.20$ the standard-penalty regime yields a 2.1% attack-success probability at a cost of 1.78 M tokens ($\approx 0.18\%$ of total supply T_0), while the doubled-penalty regime reduces this to 1.3% and raises cost to 2.57 M tokens. Doubling the penalty approximately halves attack probability across the whole sub-threshold range without changing the threshold itself.

Second, an absorbing transition occurs at $p \approx 0.35$: consensus failure becomes deterministic, attack probability saturates near 0.85 in both regimes, and economic deterrence loses traction. This separates two qualitatively distinct attacker postures: rational actors aiming to maximise expected return are deterred by penalties below the threshold, while well-funded adversaries who tolerate the cost can disable consensus altogether. The implication for CBDC governance is that economic incentives are necessary but not sufficient; technical safeguards against Sybil entry – node identity verification, geographically distributed onboarding, and stake caps – remain non-substitutable above $p = 0.30$.

Third, time-to-detection ranges from 4.4 months at $p = 0.05$ down to 2.3 months at $p = 0.40$, with the doubled-penalty regime detecting 30–40% faster. This horizon should calibrate the latency requirements of supplemental anomaly-detection layers; CBDC monitoring infrastructure based on AI methods (Sect. 6) needs to operate within roughly a quarter-year detection budget under a high-effort attack.

Table 6. Attack outcomes by malicious share p and penalty regime.

p	P(success) 1×	Cost 1× (M tok)	P(success) 2×	Cost 2× (M tok)	TTD 1× (mo)	TTD 2× (mo)
0.05	0.000	0.49	0.000	0.71	4.35	3.85
0.20	0.021	1.78	0.013	2.57	3.12	2.27
0.25	0.155	2.22	0.085	3.21	2.85	1.99
0.30	0.422	2.67	0.230	3.85	2.63	1.79
0.33	0.586	2.93	0.305	4.24	2.51	1.69
0.35	0.847	3.11	0.847	4.50	—	—
0.40	0.860	3.55	0.857	5.15	—	—

5.5 Validation Against Published F-BFT Measurements

To establish the throughput model's empirical credibility we compare its predictions in the small-cluster Docker regime against the asynchronous-stream test reported in Bogdanov et al. [20]. The model is run with the production hardware multiplier disabled (1.0×) so the comparison is on equivalent footing. Results are in Table 7.

The model agrees with measurement to within 21% for $n \leq 18$, and to within 12% on average across that range. The 59% error at $n = 21$ reflects a regime our

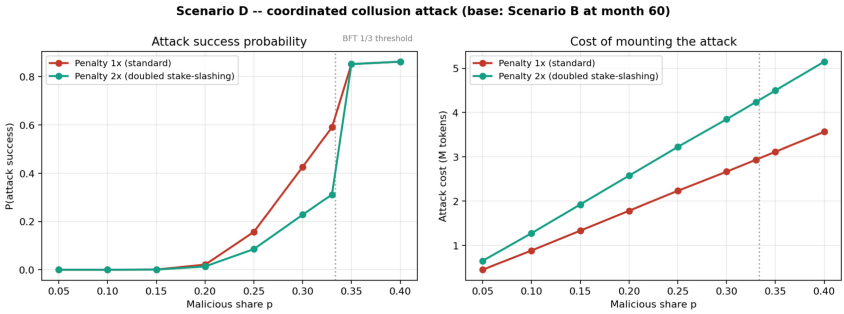

Fig. 4. Attack-success probability and attacker cost versus malicious share p, two penalty regimes, with the BFT 1/3 threshold marked.

Table 7. Model versus published F-BFT measurements [20].

n	Measured CTPS [20]	Model CTPS	Relative error
6	1.23	1.14	7.3%
12	1.49	1.46	2.0%
15	1.72	1.40	18.6%
18	1.09	1.31	20.2%
21	0.81	1.29	59.3%

model does not simulate: at this cluster size the test bed in [20] reports a commit ratio of 0.28, indicating workload saturation. We note this limitation explicitly. For the cluster sizes used inside our federated regime (clusters of ∼15 nodes), the small-cluster validation is the relevant benchmark, and the agreement supports use of the throughput model for the comparative analyses in 5.3–5.4.

5.6 Sensitivity Analysis

To check whether Scenario B results depend sensitively on parameter choices, we perturb the four most consequential parameters by ±20% (or +50% for the attrition rate a_0) and re-run the simulation. End-of-horizon CTPS and node count are reported in Table 8.

The end-state CTPS varies by less than 25% across the perturbation grid, and the qualitative ordering of regimes is preserved. The model is most sensitive to $N_{\max}$: networks pushed beyond the per-cluster optimum suffer disproportionately, reinforcing the principal design lesson – namely, that node-count expansion in a CBDC context must be paired with active scalability investment (sharding, segmentation), not pursued as an end in itself.

Table 8. Sensitivity of Scenario B end-state to parameter perturbations.

Perturbation	Nodes at $t = 100$	CTPS at $t = 100$
Baseline B	577	67.7
r_0 +20%	603	71.6
r_0 −20%	532	71.8
$N_{\max}$ +20%	688	61.6
$N_{\max}$ −20%	458	83.7
a_0 +50%	559	79.0

5.7 Implications for System Design

The four scenarios collectively support three claims that go beyond what the unparameterised AB2023 description in [19] could establish.

First, the centralisation–performance trade-off in Scenario B is real but bounded in time. The ∼8% drop in CTPS between months 0 and 36 is recoverable: after sharding and H-NET segmentation are activated, throughput exceeds the baseline-B starting point by month 72. Central banks contemplating a transition from centralised to permissioned-distributed CBDC must therefore plan for a transient performance valley spanning roughly three years under the parameters used here.

Second, the attack-economics analysis in Scenario D quantifies the boundary at which economic deterrence ceases to operate. Doubling the penalty approximately halves attack probability across the entire $p < 1/3$ region but does nothing once p exceeds the BFT threshold. Stake-slashing alone is therefore inadequate as a security architecture above moderate Sybil-resistance thresholds.

Third, Scenario C demonstrates that user-side adoption alone does not stress consensus performance directly; the bottleneck remains the node-count regime. CBDC operators expecting rapid retail uptake should plan capacity for horizontal scaling (sharding, segmentation) rather than for consensus tuning.

Together, these results provide a quantitative vocabulary for techno-economic discussions about CBDC infrastructure: an empirically anchored throughput function, a parameterised attack-economics model, and a composite decentralisation index whose movement can be traced to specific design interventions. The full simulation code, parameter files, and CSV outputs are available as supplementary material.

6 Future Directions: AI and CBDC

The intersection of CBDC and AI is starting to receive concrete operational interest, mostly at the monitoring and decision-support layer rather than at the consensus layer. Four directions look most useful:

- Anomaly detection over transaction streams can identify behavioural patterns associated with fraud, money laundering, and account-takeover attacks—categories that have well-established AML/CFT regulatory frameworks but where rule-based detectors miss novel attack patterns.
- Autonomous policy execution: AI agents can adjust smart-contract parameters dynamically in response to macroeconomic indicators, augmenting (not replacing) the central-bank policy committee. Pilot work in this direction has appeared at BIS Innovation Hub events [14].
- Adaptive consensus tuning: dynamic adjustment of block-size, validator-set composition, and shard allocation based on real-time load—the optimisation problem here is well-suited to reinforcement learning, with throughput and security as the competing objectives.
- User-facing assistants: natural-language interfaces and personalised financial agents that lower the barrier to CBDC adoption among less-technical users.

The integration is not free of complications. Explainability requirements for regulatory decisions, privacy guarantees for user data feeding AI models, and fairness constraints on algorithmic decisions all need to be specified before deployment, not retrofitted. Federated learning and differentially private training are two of the technical handles being studied; governance frameworks for AI in monetary infrastructure remain an open research area.

7 Conclusion

An integrated analysis is provided for the impact of blockchain technology on CBDCs below.

Major findings:

- CBDC ledger architectures fall on a spectrum from centralised to permissioned-distributed, with hybrid designs increasingly preferred when both regulatory control and cross-border interoperability are required.
- Privacy is achievable through multi-tier anonymity combined with zero-knowledge proofs; programmable smart contracts enable automated payment and policy execution; and post-quantum migration paths have been mapped by NIST [9,17] for long-horizon security.
- The principal technical bottlenecks—performance, quantum threats, consensus attacks, and cross-chain interoperability—admit partial solutions today (sharding, post-quantum cryptography, economic penalties, ISO 20022 standardisation), but no single solution is dominant; continued research on each axis remains essential.
- The AB2023 simulation reported in Sect. 5 yields specific operational findings: consortium throughput between 28 and 68 CTPS depending on regime, a Gini coefficient trajectory from 0.89 to 0.20 under gradual decentralisation, and a hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold.

– AI integration at the CBDC monitoring layer offers gains in fraud detection and adaptive operations, but raises new concerns about explainability, privacy, and algorithmic fairness that require their own technical and governance treatment.

No single architecture fits all use cases; appropriate choices emerge from joint technical and institutional considerations. Pilot-data validation of the AB2023 model and the development of international technical standards remain open agenda items.

Acknowledgments. The authors acknowledge the Lyapunov Laboratory at SPBU for the computational environment used to run the simulation experiments. Any errors are our own.

Disclosure of Interests. The authors declare no competing interests.

References

1. World Bank: Global Payment Systems Survey (2024)
2. Frizzo-Barker, J., et al.: Blockchain as a disruptive technology for business. Int. J. Inf. Manage. **51**, 102029 (2020)
3. Bank for International Settlements: 2023 BIS annual economic report (2023)
4. Kumhof, M., Noone, C.: Payment efficiency gains from central bank digital currency. BIS Q. Rev., 22–35 (2024)
5. People's Bank of China: E-CNY white paper (2024 update) (2024)
6. European Central Bank: Report on a digital euro (2023)
7. mBridge Project Consortium: Phase III technical implementation report (2024)
8. International Monetary Fund.: Conflict of laws in cross-border CBDC transactions. (IMF Policy Paper 2025-003) (2025)
9. National Institute of Standards and Technology: Post-quantum cryptography for financial networks (NIST SP 800-208) (2024)
10. Yaksetig, M., et al.: Rayls II: Fast, Private, and Compliant CBDCs. Cryptology ePrint Archive, Paper 2025/1638 (2025)
11. Yaksetig, M., Xu, J.: Rayls: A Novel Design for CBDCs. Cryptology ePrint Archive, Paper 2025/1639 (2025)
12. Yao, Q.: Controlled anonymity in CBDC system architecture. J. Digit. Currency Res. **7**(2), 513–529 (2024)
13. Central Bank of The Bahamas: Sand Dollar design paper: Programmable fiscal rules (2023)
14. Bank for International Settlements Innovation Hub: Annual report 2023 (2023)
15. Liu, X., Shan, X.G., Jiang, N.: Analysis of Sino-US blockchain technology competition situation based on patent information. Sci. Technol. Progress Policy **37**(18) (2020)
16. Committee on Payments and Market Infrastructures: Interlinking payment systems and the role of application programming interfaces. BIS (2023)
17. National Institute of Standards and Technology: Status report on the third round of the NIST post-quantum cryptography standardization process (NIST IR 8413) (2022)

18. Financial Stability Board: FSB report on the international regulation of crypto-asset activities (2023)
19. Khvatov, V., Bogdanov, A.: DGT Network Tokenomics — AB2023 Model Specification. Technical Documentation Rev. 21 (2023). https://dgtnetwork.gitbook.io/dgt-docs (Section 5)
20. Bogdanov, A., et al.: Testing and comparative analysis of the F-BFT-based DLT solution. In: Gervasi, O., et al. (eds.) ICCSA 2021. LNCS, vol. 12952, pp. 31–41. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86973-1_3
21. Bogdanov, A., Khvatov, V., Uteshev, A., Shchegoleva, N.: Virtual blockchain network for secure financial and industrial applications. In: ICCSA 2024 Workshops, LNCS, vol. 14815, pp. 142–161. Springer (2024)
22. Cong, L.W., Li, Y., Wang, N.: Tokenomics: dynamic adoption and valuation. Rev. Financ. Stud. **34**(3), 1105–1155 (2021). https://doi.org/10.1093/rfs/hhaa089
23. Letychevsky, O., Peschanenko, V., Radchenko, V., Poltoratzkyi, M., Kovalenko, P., Mogylko, S.: Formal verification of token economy models. In: 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), pp. 201–204. IEEE (2019). https://doi.org/10.1109/BLOC.2019.8751318
24. Jurjens, J., Scheider, S., Yildirim, F., Henke, M.: Tokenomics: decentralized incentivization in the context of data spaces. In: Otto, B., ten Hompel, M., Wrobel, S. (eds.) Designing Data Spaces, pp. 91–110. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-93975-5_6
25. Lesavre, L., Varin, P., Yaga, D.: Blockchain Networks: Token Design and Management Overview. NIST Interagency Report 8301 (2021). https://doi.org/10.6028/NIST.IR.8301

From Supply Chain Finance to Paradigm Reconstruction: The Synergistic Evolution of CBDC and Blockchain

Hu Bitai[1], Valery Khvatov[1], Egor Savkov[2], and Aleksandr V. Bogdanov[2(✉)]

[1] St. Petersburg State University, 7/9 Universitetskaya Emb., St. Petersburg 199034, Russia
[2] DGT Technologies AG, Toronto, Canada
a.v.bogdanov@spbu.ru

Abstract. More than 130 jurisdictions now run CBDC pilots [1]; supply chain finance (SCF) is one of the densest test beds for joint deployment with distributed ledger technology [13]. We organise the joint-deployment problem along three axes: blockchain plus CBDC against the three SCF frictions (information asymmetry, settlement delay, SME credit constraint); three institutionally distinct national patterns (Singapore PBM [8], EU privacy-centred digital euro [6], U.S. market-driven framework [9]); and a four-feature paradigm—intelligence, trustworthiness, inclusion, cross-border—tested on a tokenomic simulation. We re-implement the AB2023 model of Khvatov and Bogdanov [14] in Python and run four scenarios over 100 months: e-CNY-style retail (115 nodes, ∼28 CTPS), wholesale permissioned-distributed (575 nodes, Gini 0.89 → 0.20, CTPS 30 → 27 → 68 with sharding at month 42 and H-NET segmentation [16] at month 66), high-adoption SCF (301 nodes, 1.32 M users, ∼52 CTPS), and a 30 000-round Monte Carlo collusion sweep. Doubling the slashing penalty halves attack-success below the BFT 1/3 threshold; above it both regimes saturate near $P = 0.85$. The throughput model agrees with F-BFT measurements [15] to within 12% mean error for $n \leq 18$. Two implications follow. The retail/wholesale CBDC choice is a quantifiable trade-off; the throughput valley between months 24 and 42 in Scenario B costs roughly three years if scalability investment lags behind node growth. Trustworthiness needs three layers: token-economic deterrence below BFT, technical Sybil resistance above it, and AI anomaly detection on a 2–4 month latency budget. The contribution is a parameterised vocabulary linking paradigm goals to operational outcomes, not a new architecture.

Keywords: CBDC · Blockchain · Supply chain finance · Cross-border payments · Programmable money · AB2023 model · Tokenomics simulation

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 22–38, 2027.
https://doi.org/10.1007/978-3-032-30518-3_2

1 Introduction

1.1 Motivation and Global Context

CBDC have crossed from research literature to operational infrastructure in over 130 jurisdictions, and 94% of central banks report some form of CBDC engagement [1,4]. China's e-CNY has surpassed 300 million users [5]; the Eurosystem's digital-euro project sits in its preparation phase [6]; mBridge has reached Phase III multi-jurisdiction settlement [7]. The open question is no longer adoption but design: how the architectural and policy choices behind a CBDC interact with the application domains that absorb it first.

Supply chain finance (SCF) is the densest of those domains. The sector's three legacy frictions—multi-tier information asymmetry, five-to-ten-day letter-of-credit settlement, persistent SME credit constraint—map onto what blockchain and CBDC jointly address [12,13]: blockchain provides shared, immutable record-keeping and programmable execution; CBDC provides final settlement in central-bank money without intermediary chains. We take that joint deployment as our empirical anchor and use a tokenomic simulation to attach quantitative price tags to the policy choices it forces.

1.2 Literature Review

Four overlapping research streams cover the CBDC–blockchain–SCF intersection.

Stream 1—Blockchain in Supply Chain Finance. Frizzo-Barker et al. [2] surveyed 158 articles on blockchain as a disruptive technology and identified trade finance among the densest application clusters. Tang and Chen [19], on 88 listed Chinese logistics firms 2010–2023, report a 0.759 difference-in-differences coefficient of blockchain adoption on supply-chain resilience, mediated by innovation capacity and total factor productivity. Li and Du [18], surveying 1,760 papers, find digital currency, market efficiency, and SCF applications dominant. The stream treats the settlement asset as exogenous.

Stream 2—CBDC Design and Macro-effects. Kumhof and Noone [10] formalise CBDC design principles for financial stability; the BIS [3] provides annual cross-jurisdictional surveys. The stream is theoretically strong on monetary effects but rarely engages the technical substrate.

Stream 3—Cross-Border and Multi-CBDC. mBridge [7] has produced operational evidence that wholesale CBDC bridges compress settlement from days to minutes; the BIS Innovation Hub's Project Agorá [11] extends the architecture to a tokenised correspondent-banking layer. Empirical analysis remains thin because the implementations are recent.

Stream 4—Comparative National Approaches. MAS has published Project Orchid materials including the Purpose Bound Money protocol [8]; the ECB [6] documents digital-euro design with privacy as the governing constraint; the U.S. Federal Reserve [9] has issued discussion papers without committing to retail CBDC.

The streams rarely meet. SCF–blockchain studies fix the settlement asset; CBDC studies black-box the substrate; comparative studies stay qualitative. What is missing is a parameterised framework in which CBDC design choices, blockchain architecture choices, and SCF deployment patterns generate measurable joint outcomes. This paper supplies one.

1.3 Research Design and Contributions

We combine three methodological elements: comparative case analysis of national approaches, a paradigm-reconstruction synthesis built on the comparative material, and a parametric simulation of the AB2023 tokenomic model linking paradigm goals to operational outcomes. The research questions are: (i) which SCF pain points yield to joint blockchain–CBDC solutions and which resist them? (ii) what differentiates the Singapore, EU, and U.S. patterns in terms of deployment viability? (iii) at what threshold do paradigm goals become parameter-attainable rather than aspirational?

Contributions: a structured analysis of CBDC–blockchain synergy in SCF organised by three pain points and three solution mechanisms (Sect. 2; Table 1); a comparative framework over the three national patterns with explicit institutional and SCF implications (Sect. 3; Table 2); a four-dimensional paradigm derived from comparison rather than asserted (Sect. 4); a 100-month parametric simulation under four CBDC archetypes, validated against published F-BFT measurements [15] to within 12% mean error and reporting quantitative trade-offs that link paradigm goals to operational outcomes (Sect. 5; Tables 3–7; Figs. 1–2). Every paradigm claim in Sect. 4 maps to a measurable AB2023 indicator in Sect. 5; every Sect. 5 number reproduces from Table 3 with random seed 20240508.

The paper proceeds: Sect. 2 covers SCF pain points and joint solutions. Section 3 presents the comparative cases. Section 4 develops the four-feature paradigm. Section 5 reports the simulation. Section 6 closes with policy implications and an agenda.

2 Synergistic Solutions in Supply Chain Finance

SCF moves working capital and liquidity along multi-party trade networks [13]. Three structural inefficiencies persist that neither blockchain nor CBDC fully solves alone: information asymmetry across multi-tier supplier networks [18,19], in which financial institutions cannot verify at low cost that an invoice corresponds to a genuine trade, leading to elevated SME financing premia; operational delay, with letter-of-credit transactions taking five to ten business days from application to fund release [12], the bottleneck being manual reconciliation across banks, freight forwarders, customs, and treasuries rather than any single step; and SME-tier liquidity constraints [13,18], where bank programmes favour anchor enterprises over the long tail of suppliers, with the resulting credit shortfall being one of the most-cited frictions in the empirical SCF literature [19].

2.1 Blockchain Mechanisms for Trust and Transparency

Three mechanisms account for most reported benefits. *Immutable record-keeping* [2]: order placement, shipment events, inspection reports, and invoice issuance write to a shared ledger that authorised participants verify independently, which suppresses invoice double-pledging—one of the larger historical loss categories for trade-finance lenders. *Smart-contract automation* [12]: business rules such as payment on inspection or against bill of lading encode as contracts that trigger on condition match; IoT sensor evidence (RFID inventory, GPS shipping data) feeds them as oracles, eliminating most manual reconciliation. *Tokenisation of real-world assets* [18]: accounts receivable, warehouse receipts, and other trade assets become fungible or non-fungible tokens with fractional ownership, which expands the funding pool for suppliers who lacked a liquid receivables market. The empirical case is no longer purely conceptual: Tang and Chen [19] report a 0.759 DiD coefficient on supply-chain resilience over 88 firms (2010–2023), with innovation capacity and total factor productivity as significant mediators. The settlement-asset choice—in particular whether that asset is a CBDC—is the next question.

2.2 The Enabling Role of CBDC: Value-Flow Integration

CBDC, as a direct liability of the central bank, supports final settlement in central-bank money without intermediary chains [10]; integrated with a blockchain trade platform, payment occurs atomically with title transfer (DvP at machine speed), collapsing multi-day letter-of-credit settlement risk. CBDC issuance can carry conditional restrictions on use [14]: in an SCF context, a buyer releases payment spendable only by designated sub-suppliers, preventing leakage to unrelated obligations—the mechanism Singapore's PBM implements directly [8]. For cross-border settlement, the largest cost line in international trade finance, multi-CBDC platforms such as mBridge let participating commercial banks transact directly on a shared ledger, compressing settlement from days to seconds and reducing aggregate cost by roughly 40–60% [7].

3 Global Exploration Paths: A Comparative Analysis

National CBDC–blockchain integration diverges along institutional axes that the technology alone does not determine. Three cases bracket most of the variation observed: Singapore's design-driven, ecosystem-building approach; the EU's privacy-first, compliance-bound trajectory; and the U.S. market-driven, regulation-first stance. The differences condition which deployment patterns are SCF-viable in each jurisdiction.

Table 1. Supply chain finance use cases under joint blockchain–CBDC deployment.

Use case	Blockchain role	CBDC role	Reported gain	Example
Receivables financing	Authenticity, double-pledge prevention	Real-time, programmable release	Cycle days → hours	e-CNY pilot [5]
Cross-border payment	Multi-lateral shared ledger	Synchronised cross-currency settlement	2–5 days → seconds; cost ↓ 40–60%	mBridge Phase III [7]
Trade letters of credit	Full documentation on-chain	Conditional automatic payment	5–10 days → within 24 h	Singapore PBM [8]
Advance-payment financing	Logistics tracking	Targeted fund release	Risk-controlled	Bahamas Sand Dollar (illustr.)

3.1 Singapore: Purpose Bound Money and a Layered Stack

MAS has anchored CBDC innovation in two flagship programmes, Project Ubin (wholesale) and Project Orchid (retail and programmable money) [8]. The architecture has three features. First, a layered ecosystem: wholesale CBDC for inter-bank settlement, tokenised bank liabilities for commercial-bank money on chain, and regulated stablecoins for specific commercial use cases, with PBM as the universal programmable layer across all three [8]. Second, PBM proper: an issuer attaches conditions—payee, purpose, expiry—at the protocol layer rather than the application layer; the Amazon–HSBC pilot used PBM to tokenise accounts payable, letting merchants obtain unsecured working capital against the resulting digital receivables [8]. Third, wholesale CBDC live testing for inter-bank settlement, with planned extension to cross-border securities settlement.

The pattern is design-driven: the regulator participates in protocol design rather than waiting for market consolidation, shaping standards while the ecosystem is still small and tractable.

3.2 European Union: Privacy at the Core

The ECB has pursued the digital euro since 2021 with privacy and data protection as governing constraints [6]. Public consultations identified privacy as the top user concern, and the design embeds GDPR-aligned protections including offline transaction modes for low-value retail payments. Programmability is deliberately limited compared with Singapore's PBM, emphasising payment sovereignty and inclusion over conditional spending [6]. While the retail digital euro receives most public attention, the ECB and several national central banks (notably Banque de France) have run wholesale CBDC experiments for cross-border and securities settlement [11]. The pattern is compliance-first: innovation must fit existing legal frameworks, not reshape them.

3.3 United States: Market-Driven, Regulation-First

The Federal Reserve has moved cautiously on retail CBDC while encouraging private-sector innovation [9]. The U.S. has prioritised establishing a regulatory framework for digital assets, including stablecoins, through executive orders and Treasury policy with consumer protection, financial stability, and illicit-finance prevention as stated priorities [9]. Innovation in blockchain payments and SCF is driven by technology companies and financial institutions; the Canton Network targets wholesale settlement and trade-finance use cases. The Fed has not committed to a digital dollar, citing the need for Congressional authorisation [9]. The market-driven approach exploits private-sector dynamism but carries fragmentation risk: without mandated interoperability standards, parallel platforms can proliferate without converging.

3.4 Comparative Insights

Table 2 summarises how each pattern translates into SCF-relevant outcomes. The three paths represent stable points in the design space rather than transient choices. A fourth pattern—China's e-CNY rollout, with its centralised governance and large user base [5]—supplies the reference for Sect. 5's retail-CBDC scenario.

Table 2. Comparison of global CBDC exploration paths and their SCF implications.

Region	Core characteristic	Focus in SCF	Regulatory orientation
Singapore	Layered ecosystem; PBM	Enterprise programmable financing; wholesale settlement	Proactive design; protocol-level regulator involvement
European Union	Privacy first; limited programmability	Retail SME payment; wholesale pilots (Banque de France)	Compliance-first; strong legal framework
United States	Market-driven; private-sector-led	Blockchain SCF platforms; future regulated stablecoins	Sandbox guidance; risk-control emphasis

4 Paradigm Reconstruction: Towards a New Financial System

CBDC–blockchain integration is not an incremental upgrade to payment rails but a reshaping of how the financial system is governed, audited, and connected

across borders. Four features define the emerging paradigm: intelligence, trustworthiness, inclusion, and cross-border functionality. Each is illustrated below with an SCF hook—since SCF remains our empirical anchor—and each maps to a measurable AB2023 indicator in Sect. 5. The institutional preconditions for the paradigm (monetary classification as M0 [10]; legal-tender status under amended central-bank acts as in eNaira and the e-CNY [5]; data-protection regimes that constrain programmability under GDPR [6] or accept it under PBM [8]; and financial-stability holding caps such as the ECB's ∼3,000 retail cap [6] or the e-CNY's tiered identity-keyed limits [5]) shape Sect. 5 Scenario A's adoption ceiling directly.

4.1 Intelligent Finance

CBDC carrying spending conditions in the protocol layer makes conditional payment routine [14]: a trade-finance loan releases automatically when a bill of lading hashes onto the ledger; a green-bond disbursement flows only to certified low-carbon suppliers [10]. Both patterns hit Table 1's accounts-receivable-financing and milestone-triggered-LC use cases. Anonymised, aggregated CBDC transaction data, processed through statistical and ML pipelines, can give central banks early warning of liquidity stress at the sectoral or regional level [3]—shifting monetary aggregates from a coarse unit-of-account variable to a finer addressable grain, which makes targeted SCF relief programmes feasible in days rather than quarters.

4.2 Trustworthy Finance

Legacy finance trust rests on the reputation and regulatory backing of centralised institutions: banks, clearinghouses, rating agencies. Blockchain layers cryptographic trust on top, reducing what each institution carries single-handedly [2]. Key trade and settlement records write to a shared ledger that no single participant rewrites unilaterally [18]; this addresses the asymmetries embedded in asset-backed securities, where buyers historically had to trust originators on receivable quality. With receivables on chain, the verification cost shifts from the originator's reputation to cryptographic attestation. Through zero-knowledge proofs and verifiable credentials, users prove specific attributes—creditworthiness, regulatory status—without exposing underlying data [14], so SMEs demonstrate credit history without surrendering full transaction logs and corporate buyers verify supplier compliance without aggregating supplier data centrally.

4.3 Inclusive Finance

Financial inclusion remains binding in much of the global South where physical banking is sparse [1]. A CBDC wallet runs on a smartphone or, in low-bandwidth deployments, a feature phone with a SIM-based hard wallet. The Bahamas Sand Dollar was designed precisely to serve residents across more

than 700 islands where conventional banking is uneconomic [1]; Nigeria's eNaira targets the unbanked. CBDC transaction data, combined with other digital footprints—agricultural IoT readings, e-commerce histories, mobile-money records—supports credit assessment for smallholder farmers and informal businesses that legacy bank systems cannot rate [19]. For SCF this opens working-capital credit down to the long tail of small suppliers, conditional on the data-protection rules covered above.

4.4 Cross-Border Finance

Cross-border payment is the single most inefficient global-finance segment: settlement still takes days, retail remittance costs run 6–11% of principal, reconciliation across correspondent-banking chains is the dominant error source [3]. mBridge lets participating banks exchange CBDC directly across the platform's shared ledger, bypassing correspondent intermediaries [7]: settlement compresses from days to seconds and aggregate cost falls by 40–60%. Project Agorá extends the architecture to a tokenised correspondent-banking layer [11]. For SCF this collapses the asynchronous-LC class of transactions to synchronous settlement, removing principal risk from a class that has carried it for centuries. Currency competition follows: a CBDC participating in mBridge or Agorá inherits the platform's reach, while a non-participating one carries a structural disadvantage in trade settlement.

5 Quantitative Calibration via the AB2023 Tokenomic Model

The four paradigm goals of Sect. 4 are aspirational until tied to measurable parameter regimes. Each goal is mapped here to a CBDC policy archetype, each archetype to AB2023 parameter settings [14], each parameter setting to a 100-month trajectory. The four scenarios should be read as calibrations of policy intent, not predictions of any specific deployment.

5.1 Shared Simulation Framework and Scope

The framework, parameter values, and CSV outputs reported below are shared with the companion paper Hu and Bogdanov (2026) [17], which discusses the same AB2023 implementation from a technical-validation standpoint. The present section re-interprets the four scenarios as policy parameterisations of Sect. 4's paradigm goals rather than as evidence about consensus or scalability per se. Raw CSV data and source code are supplementary material common to both papers.

The AB2023 framework couples seven modules: (i) a Richards-type node-growth process; (ii) a logistic-Gompertz node-driven (LGND) user-growth process; (iii) a token-supply rule with mint and burn; (iv) a stake-weighted token-distribution dynamic; (v) decentralisation indicators—Active Node Index (ANI),

Network Diversity Index (NDI), Consortium Influence Index (CII), Token Distribution Inequality (TDI, Gini)—aggregated into composite metric D; (vi) an F-BFT throughput model anchored to Bogdanov et al. [15]; (vii) a Monte Carlo collusion-attack module. We re-implemented these modules in Python (NumPy and Pandas only, ≈550 lines) following the equations of [14]. Two simplifications are explicit: a geometric Brownian motion replaces the full equilibrium token-price formulation, and cross-cluster network latency is omitted. Reproducibility is anchored at random seed 20240508.

The throughput module deserves a separate note since it carries the empirical weight of the section. We anchor its small-cluster behaviour to the asynchronous-stream measurements of [15] (CTPS in 0.6–1.7 for $n \in \{6, 12, 15, 18, 21\}$). Outside that range, N nodes partition into $\lceil N/15 \rceil$ clusters with arbiter-ring coordination cost $1/(1 + (k/k^*)^2)$ at $k^* = 12$. A 5.5× hardware multiplier projects from the Docker test bed to production infrastructure, consistent with the 5–6× H-NET latency reduction reported in [16]. Sharding adds $0.75 \times (n_{\text{shards}} - 1)$.

5.2 Mapping Policy Archetypes to Scenarios

The four scenarios (Table 3) tie back to Sects. 3 and 4. Scenario A (e-CNY archetype, $N_{\max} = 115$, $U_\infty = 80\,\text{K}$) models a tightly governed retail rollout where the central bank retains direct node-admission control [5]. Scenario B (mBridge / Agorá archetype, $N_{\max} = 600$, sharding at month 42, H-NET segmentation at month 66) models active expansion with explicit interoperability and gradual decentralisation goals [7]. Scenario C (Singapore PBM / tokenised-

Table 3. Scenario parameter settings. N_0 initial nodes; $N_{\max}$ carrying capacity; r_0 baseline node-growth rate; U_0, U_∞ initial and saturating users; T_{sh}, T_{seg} sharding and H-NET segmentation activation months; p malicious-node share (D only).

Parameter	A: Retail	B: Wholesale	C: High-adoption SCF	D: Attack
N_0, $N_{\max}$	100, 115	100, 600	100, 320	inherits B at month 60
r_0	0.025	0.07	0.05	—
U_0, U_∞	10 K, 80 K	10 K, 200 K	10 K, 1.2 M	—
Consortium share (initial)	0.65	0.65	0.45	—
Wealth concentration (target)	0.78	0.55	0.62	—
T_{sh}, T_{seg}	none	42, 66	50, none	—
p sweep	—	—	—	$\{0.05, 0.10, \ldots, 0.40\}$
Penalty regimes	—	—	—	1× standard, 2× slashing

receivables, $U_\infty = 1.2\,\mathrm{M}$, sharding at month 50) tests whether consensus remains the binding constraint under aggressive demand-side pressure [8]. Scenario D freezes Scenario B at month 60 (498 nodes, 425 K users) and sweeps malicious-share $p \in [0.05, 0.40]$ under two penalty regimes, quantifying the economic-incentive boundary at which CBDC governance deters collusion without off-chain enforcement.

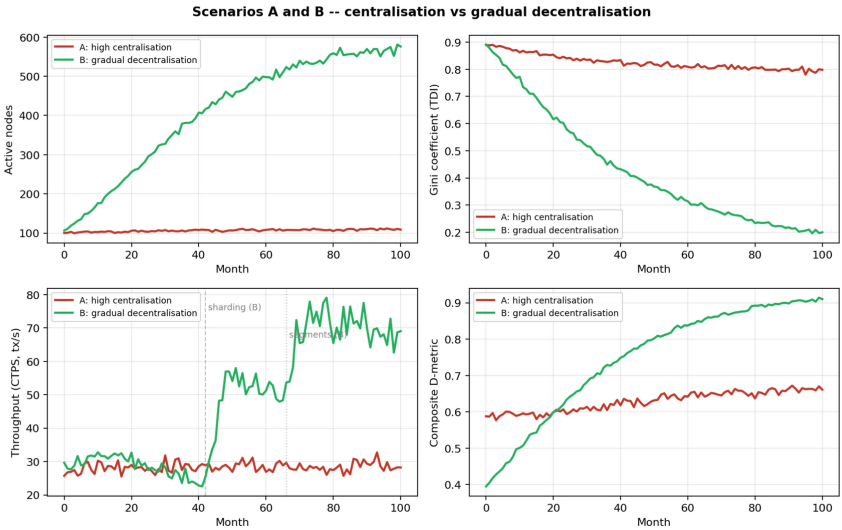

Fig. 1. Four-panel comparison of Scenarios A and B: (a) active nodes; (b) Gini coefficient; (c) CTPS, with sharding (month 42) and H-NET segment (month 66) marked; (d) composite D-metric. (Color figure online)

5.3 Time-Series Results

Scenario A. The network stabilises at 108 ± 4 active nodes (Fig. 1a). Gini drifts $0.89 \rightarrow 0.80$ (-10%); CII holds near 0.60; CTPS settles in 26–32; composite D rises $0.59 \rightarrow 0.66$. The trajectory matches the e-CNY profile in [5], where the two-tier issuance arrangement preserves the central bank as sole issuer and authorised intermediaries handle distribution.

Scenario B. Active nodes rise $100 \rightarrow 575$ (Fig. 1a, green). Gini falls $0.89 \rightarrow 0.20$ (-77%); CII falls $0.65 \rightarrow 0.11$ (-83%). CTPS traces a non-monotone path: an initial $\approx$30 plateau (months 0–24), a slow decline to $\approx$27 at month 36 as arbiter-ring overhead grows, sharding lifting it to $\approx$54 by month 48, H-NET segmentation carrying it to $\approx$73 by month 72, end-of-horizon settling at 68. Composite D rises $0.40 \rightarrow 0.91$ ($+128\%$): node-count expansion translates into substantive decentralisation rather than nominal participation. Milestones are in Table 4. The CTPS minimum at month 36 is not a model artefact—it is the

predictable consequence of arbiter-ring overhead growing faster than per-cluster throughput in the federated regime [16].

Table 4. Scenario B selected milestones.

Month	Active Nodes	Gini (TDI)	CII	CTPS	D-metric
0	106	0.89	0.65	29.6	0.40
24	282	0.58	0.42	28.8	0.63
36	381	0.45	0.34	27.2	0.73
48	461	0.37	0.28	54.0	0.80
60	498	0.31	0.23	49.4	0.84
72	538	0.27	0.19	73.4	0.87
100	575	0.20	0.11	67.7	0.91

Scenario C. The user base scales 50 K $\rightarrow$ 1.32 M around month 50 (S-curve); node count expands to 301 by month 100 as financial institutions, logistics providers, and corporate treasurers enter the platform supporting PBM-style conditional payments and accounts-receivable tokenisation. CTPS at month 100 is $\approx$52, tracking Scenario B at the same node count. Throughput here is governed by node count and consensus topology, not user load: even at >1 M active wallets the binding constraint is federated cluster topology, not per-transaction consensus load. SCF deployments scaling toward Singapore-PBM volumes need horizontal scaling (sharding, segmentation) more than BFT optimisation.

5.4 Governance Under Coordinated Attack: Scenario D

The trustworthiness paradigm of Sect. 4.2 rests on the claim that economic incentives deter collusion within a CBDC ledger. Scenario D quantifies the boundary at which that claim breaks down. We freeze the Scenario B network at month 60 (498 nodes, 425 K users, mature token supply) and sweep malicious-share p from 0.05 to 0.40 under two penalty regimes. For each $(p, \text{penalty})$ pair we run 30 000 Monte Carlo rounds. Each round draws 50 validators by hypergeometric sampling, declares per-round consensus compromised when malicious validators reach the BFT threshold $\lceil n/3 \rceil$, and tests anomaly detection with probability $0.20{+}0.6\,p{\times}\text{penalty}$. Compounded attack-success accounts for repeated attempts before detection closes the window. Attacker cost equals stake locked in malicious nodes plus an opportunity-cost adjustment scaled by the penalty multiplier.

Three observations follow from Table 5 and Fig. 2. (i) Below the BFT threshold ($p < 1/3$), attack success scales steeply but predictably with p. At $p = 0.20$ the standard-penalty regime yields a 2.1% success probability at a cost of 1.78 M tokens ($\approx$0.18% of total supply T_0); doubling slashing drops success to 1.3% and raises cost to 2.57 M tokens. The penalty multiplier roughly halves attack probability across the entire sub-threshold range without shifting the threshold

itself. (ii) An absorbing transition occurs near $p = 0.35$: consensus failure is deterministic, attack probability saturates near 0.85 in both regimes, economic deterrence loses traction. Rational adversaries who maximise expected return are deterred below the threshold; well-funded adversaries who tolerate cost can disable consensus altogether. Above $p = 0.30$, technical Sybil safeguards—node identity verification, geographically distributed onboarding, stake caps—are non-substitutable. (iii) Time-to-detection runs 4.4 months at $p = 0.05$ down to 2.3 months at $p = 0.40$, with the doubled-penalty regime detecting 30–40% faster. AI monitoring infrastructure must operate within roughly a quarter-year detection budget under high-effort attack, giving the AI–CBDC research agenda a quantitative target rather than a qualitative wish.

5.5 Validation Against Published F-BFT Measurements

To establish empirical credibility we compare model predictions in the small-cluster Docker regime against the asynchronous-stream test of [15], with the production hardware multiplier disabled (1.0×). Results are in Table 6: the model

Table 5. Attack outcomes by malicious share p and penalty regime. TTD: time-to-detection; em-dash denotes consensus failure beyond the BFT threshold.

p	P(success) 1×	Cost 1× (M tok)	P(success) 2×	Cost 2× (M tok)	TTD 1× (mo)	TTD 2× (mo)
0.05	0.000	0.49	0.000	0.71	4.35	3.85
0.10	0.001	0.99	0.000	1.42	3.85	3.12
0.15	0.005	1.48	0.001	2.13	3.45	2.63
0.20	0.021	1.78	0.013	2.57	3.12	2.27
0.25	0.155	2.22	0.085	3.21	2.85	1.99
0.30	0.422	2.67	0.230	3.85	2.63	1.79
0.33	0.586	2.93	0.305	4.24	2.51	1.69
0.35	0.847	3.11	0.847	4.50	—	—
0.40	0.860	3.55	0.857	5.15	—	—

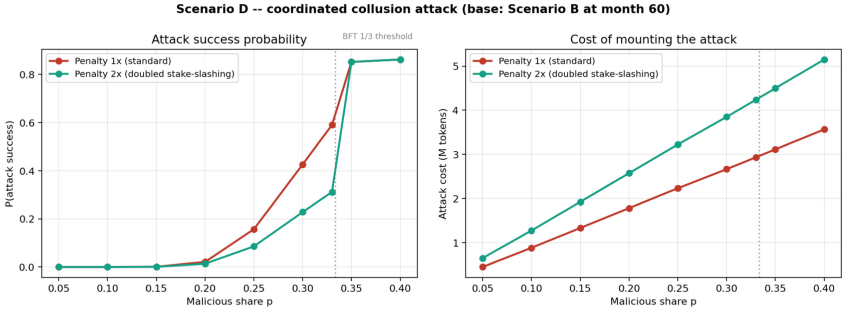

Fig. 2. Attack-success probability and attacker cost versus malicious share p, two penalty regimes, with the BFT 1/3 threshold marked.

agrees with measurement to within 21% for $n \leq 18$ and within 12% on average across that range. The 59% error at $n = 21$ reflects a regime our model does not simulate—at this cluster size the test bed reports a commit ratio of 0.28, indicating workload saturation. We note this limitation explicitly. For the cluster sizes used inside our federated regime (≈15-node clusters), the small-cluster validation is the relevant benchmark.

Table 6. Model versus published F-BFT measurements [15].

n	Measured CTPS [15]	Model CTPS	Relative error
6	1.23	1.14	7.3%
12	1.49	1.46	2.0%
15	1.72	1.40	18.6%
18	1.09	1.31	20.2%
21	0.81	1.29	59.3%

5.6 Sensitivity Analysis

We perturb the four most consequential Scenario B parameters by ±20% (or +50% for the attrition rate a_0) and re-run. End-of-horizon CTPS varies by less than 25% across the perturbation grid (Table 7); the qualitative ordering of regimes is preserved. The model is most sensitive to $N_{\max}$—networks pushed beyond the per-cluster optimum suffer disproportionately—which reinforces the design lesson that node-count expansion in a CBDC context must be paired with active scalability investment, not pursued as an end in itself.

Table 7. Sensitivity of Scenario B end-state to parameter perturbations.

Perturbation	Nodes at $t = 100$	CTPS at $t = 100$
Baseline B	577	67.7
r_0 +20%	603	71.6
r_0 −20%	532	71.8
$N_{\max}$ +20%	688	61.6
$N_{\max}$ −20%	458	83.7
a_0 +50%	559	79.0

5.7 Implications for Paradigm Reconstruction

The four scenarios connect Sect. 4's paradigm goals to AB2023 indicators on three concrete fronts. The Sect. 4.3 inclusion-versus-stability trade-off surfaces quantitatively as the contrast between Scenarios A and B: the retail-baseline regime preserves throughput stability with a Gini that barely moves over 100 months (0.89 → 0.80), while the wholesale-distributed regime drives Gini to 0.20 at the price of three years of throughput stagnation before sharding rescues performance. Section 4.4's cross-border functionality goal maps onto the post-sharding regime of Scenario B (CTPS 68, Gini 0.20, CII 0.11)—a candidate operating point for multi-CBDC bridges, where the transition path matters more than the endpoint: capacity expansions must be sequenced, not deferred, or the network spends years in the CTPS valley with elevated operational risk and political exposure. Section 4.2's trustworthiness goal cannot be satisfied by economic mechanisms alone—Scenario D shows doubling slashing penalties roughly halves sub-threshold attack probability but does nothing once p exceeds the BFT 1/3 threshold, so CBDC trust architectures need three layers: token-economic deterrence (effective for rational adversaries below BFT), technical Sybil resistance (necessary above the threshold), and AI anomaly detection (latency budget ≈2–4 months). The paradigm-reconstruction agenda becomes parameterised rather than aspirational, with the AB2023 model as the link.

6 Conclusion

This paper has examined the synergistic evolution of CBDC and blockchain through three integrated lenses: pain-point resolution in supply chain finance, comparative national approaches, and quantitative calibration via the AB2023 tokenomic model. Five findings stand out. Blockchain and CBDC address complementary frictions (information versus value flow [10,13]); neither alone resolves the structural credit and verification problems blocking SME participation in trade finance. National approaches diverge along stable institutional axes rather than transient technical choices: Singapore's design-driven, ecosystem-building strategy [8], the EU's privacy-first trajectory [6], the U.S. market-driven, regulation-first stance [9] are the principal stable points in the design space. The four-feature paradigm of Sect. 4 is a parameterisable specification—each feature maps to specific AB2023 indicators in Sect. 5. AB2023 calibration delivers operational findings (consortium throughput 28–68 CTPS by regime, Gini 0.89 → 0.20 under wholesale-distributed settings, hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold), with the throughput model agreeing with [15] to within 12% mean error for $n \leq 18$. Realising the paradigm requires joint advancement of technical standards (sharding, segmentation, post-quantum), legal frameworks (privacy regulation, cross-jurisdictional harmonisation), and institutional infrastructure (multi-CBDC platforms, deposit-cap policy); no single thread is sufficient and Sect. 5 quantifies the cost of each lagging behind.

6.1 Policy Implications

For jurisdictions designing CBDC–SCF integration, four recommendations follow. First, articulate policy objectives before architectural choices: Sect. 5 shows the cost of leaving the inclusion–efficiency–stability–cross-border trade-off implicit, as Scenarios A and B differ by a factor of more than two on Gini reduction (10% vs 77%) over the same horizon precisely because the policy goal differs [10]. Second, adapt rather than copy: PBM [8] suits jurisdictions whose central banks join protocol design; the EU privacy-first model [6] fits regimes where strong data-protection rules already constrain programmability. Third, invest in interoperability before fragmentation locks in—mBridge participation [7], ISO 20022 alignment, and Project Agorá architecture [11] all reduce later integration cost; deferring interoperability tends to make it permanently more expensive. Fourth, stress-test in simulation before deployment: AB2023 calibration surfaces non-obvious dynamics (the three-year throughput valley during cluster expansion, the BFT cliff in attack economics, the saturation of stake-slashing deterrence above $p \approx 0.33$) that descriptive analysis would miss. Live pilot data complement parametric simulation; they do not substitute for it.

6.2 Limitations and Future Research

Three limitations qualify the conclusions. First, the analysis is conceptual and computational, drawing on publicly available data and on a faithful but not exhaustive AB2023 re-implementation [14]; as more CBDC projects move from pilot to production, empirical validation against operational data will strengthen or refute simulation findings. Second, the AB2023 model does not capture the political-economy dimensions of CBDC adoption—distributional consequences of monetary digitalisation, redistribution of seigniorage between central and commercial banks, geopolitical implications of cross-border CBDC platforms all sit outside the parameter space. Third, the comparative analysis covers three high-income jurisdictions plus referenced examples from China, Nigeria, and the Bahamas; application to lower-income, infrastructure-constrained settings—where Sect. 4.3's inclusion paradigm matters most—requires recalibration of the user-growth and node-growth modules.

A natural research agenda follows. The most pressing item is empirical validation against pilot deployments such as e-CNY, eNaira, and Sand Dollar [1,5]. The second is the political economy of CBDC design choices, including who bears the cost of paradigm transitions when the throughput valley of Sect. 5 lasts three years. The third is technical: extension of AB2023 with cross-cluster latency modelling and a richer attack taxonomy, both consistent with the H-NET virtualisation work in [16]. CBDC–blockchain integration in supply chain finance has no single optimal design; appropriate choices emerge from the interplay of policy intent, institutional constraint, and technical capacity, with empirical validation of the AB2023 simulation against pilot deployments and cross-jurisdictional legal harmonisation as the principal open agenda items.

Acknowledgments. The authors acknowledge the Lyapunov Laboratory at SPBU for the computational environment used to run the simulation experiments. Any errors are our own.

Disclosure of Interests. The authors declare no competing interests.

References

1. World Bank: Global Payment Systems Survey 2024. World Bank, Washington, DC (2024)
2. Frizzo-Barker, J., Chow-White, P.A., Adams, P.R., Mentanko, J., Ha, D., Green, S.: Blockchain as a disruptive technology for business: a systematic review. Int. J. Inf. Manage. **51**, 102029 (2020). https://doi.org/10.1016/j.ijinfomgt.2019.10.014
3. Bank for International Settlements: Annual Economic Report 2023. BIS, Basel (2023)
4. Kosse, A., Mattei, I.: Making headway — results of the 2022 BIS survey on central bank digital currencies and crypto. BIS Papers No. 136. Bank for International Settlements, Basel (2023)
5. People's Bank of China: E-CNY White Paper (2024 update). People's Bank of China, Beijing (2024)
6. European Central Bank: Report on a Digital Euro – Second Progress Report. European Central Bank, Frankfurt (2023)
7. mBridge Project Consortium: Project mBridge – Phase III Technical Implementation Report. BIS Innovation Hub, Hong Kong SAR (2024)
8. Monetary Authority of Singapore: Project Orchid — Programmable Digital SGD and the Purpose Bound Money Protocol. MAS Whitepaper (2023)
9. Board of Governors of the Federal Reserve System: Money and Payments: The U.S. Dollar in the Age of Digital Transformation. Federal Reserve Discussion Paper. Federal Reserve Board, Washington, DC (2022)
10. Kumhof, M., Noone, C.: Central bank digital currencies – design principles for financial stability. Econ. Anal. Policy **83**, 553–572 (2024). https://doi.org/10.1016/j.eap.2024.07.012
11. BIS Innovation Hub: Project Agorá — Tokenisation of cross-border payments. BIS Innovation Hub Annual Report 2023. Bank for International Settlements, Basel (2023)
12. Harish, A.R., Liu, X.L., Zhong, R.Y., Huang, G.Q.: Blockchain for logistics 4.0: a systematic literature review. Transp. Res. Part E Logist. Transp. Rev. **201**, 104269 (2025). https://doi.org/10.1016/j.tre.2025.104269
13. Song, H., Yang, Y., Tao, Z.: The application of blockchain in enterprise financing. Nankai Bus. Rev. **25**(2), 34–48 (2022)
14. Khvatov, V., Bogdanov, A.: DGT Network Tokenomics — AB2023 Model Specification. Technical Documentation Rev. 21 (2023). https://dgtnetwork.gitbook.io/dgt-docs
15. Bogdanov, A., et al.: Testing and comparative analysis of the F-BFT-based DLT solution. In: Gervasi, O., et al. (eds.) ICCSA 2021. LNCS, vol. 12952, pp. 31–41. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86973-1_3
16. Bogdanov, A., Khvatov, V., Uteshev, A., Shchegoleva, N.: Virtual blockchain network for secure financial and industrial applications. In: Gervasi, O., Murgante, B., Garau, C., et al. (eds.) ICCSA 2024 Workshops. LNCS, vol. 14815, pp. 142–161. Springer, Cham (2024)

17. Hu, B., Bogdanov, A.V.: Blockchain as infrastructure for CBDC: technological foundations and challenges. In: O. Gervasi et al. (eds.) ICCSA 2026. LNCS 16759, pp. xx–yy. Springer, Cham (2026)
18. Li, D., Du, X.: Research status and prospects of blockchain technology in digital currency and supply chain finance. Chin. J. Manag. Sci. (2025). https://doi.org/10.16381/j.cnki.issn1003-207x.2023.1852
19. Tang, Y., Chen, G.: Blockchain technology empowering logistics supply chain resilience: theoretical mechanism, empirical examination, and improvement paths. Bus. Econ. Res. **7**, 82–86 (2026)

Orchestra-Paxos: A Revision of the Paxos with Dynamic Number of Quorum Participants

Mikhail Diagilev and Vladimir Korkhov(✉)

Saint Petersburg State University, 7-9 Universitetskaya Embankment, St Petersburg, Russia 199034
st072182@student.spbu.ru, v.korkhov@spbu.ru

Abstract. The present research focuses on applying consensus methods to ensure data consistency and fault tolerance in distributed systems, with particular attention to the Paxos consensus algorithm, widely used to coordinate decisions among distributed nodes. Despite its reliability, Paxos has practical limitations that affect system performance and availability in unpredictable network conditions. These include challenges in forming a quorum when nodes fail or become temporarily unavailable, as well as performance degradation caused by requiring participation from all acceptors in the consensus process. Message loss and network instability further reduce the efficiency of consensus mechanisms. To address these issues, the paper proposes a modified version of Paxos, called Orchestra-Paxos. This approach introduces partial consistency and improves fault tolerance by allowing consensus to be reached through an active subset of nodes, while others can be temporarily marked as unavailable. Unlike traditional Paxos, it does not require the involvement of a majority of accepting nodes in every replication cycle. The study uses comparative analysis between Orchestra-Paxos and Multi-Paxos, combining system modeling with experimental performance evaluation. Key metrics include average replication round time, total completion time, and the number of rounds required to accept all proposals. Results show that Orchestra-Paxos significantly improves performance under high message loss conditions (60%), reducing both the number of rounds and overall completion time by an average factor of 12 compared to Multi-Paxos.

Keywords: Computer science · Consensus algorithm · Distributed systems · Concurrent computing · Computing cluster architecture · Improved algorithm · Paxos · Dynamic quorum

1 Introduction

Modern distributed systems, data storage systems, and financial networks frequently require high availability, dependability, and security. The consensus

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 39–52, 2027.
https://doi.org/10.1007/978-3-032-30518-3_3

mechanism is a major component of such systems. This enables the system nodes to agree on conclusions about the sequence of transaction execution and the state of data, despite probable faults and delays. This is precisely why consensus algorithms are essential in achieving these goals.

However, as distributed systems become more complicated and scalable, classic consensus approaches like the Paxos algorithm [1] and the Raft algorithm [2] encounter a variety of issues. The primary issues are decreased performance, longer reaction times, and challenges maintaining consistency under high loads. Therefore, it is essential to modernize and optimize current consensus techniques, especially when dealing with extensive distributed systems consisting of dozens or even millions of nodes.

In this paper, the Paxos algorithm is selected for modernisation. Paxos is used in large distributed systems such as Megastore [3], Chubby [4], Cassandra [5] and MySQL [6]. It has a number of drawbacks. A major issue is the necessity of majority quorum in the voting phase. The system may shortly forfeit its capacity to achieve consensus for an unspecified duration, either due to an escalation in the response time of quorum members or as a consequence of failures. Such situations can lead to a loss of profit for the company and call its reputation into question, as was the case with Google in 2014 [7]. In view of this, this work is dedicated to developing an algorithm that will essentially build upon the concept of the Paxos algorithm [1]. On the other hand, at each stage of the voting process, gathering all quorum members will be optional. Any nodes that do not respond will be marked as temporarily unavailable. However, they will be able to resume decision-making to reach consensus once their operation has been restored. It is therefore anticipated that this approach can guarantee partial-incomplete system consistency and significantly increase the overall system performance.

In order to compare the developed algorithm with Paxos variants, namely Multi-Paxos [8], the following criteria will be used: the average time of a single replication round of the algorithm, the waiting time for the completion of all replication rounds, and the number of rounds required to fulfil all client requests [9].

The contribution to the field is that this is the first consensus algorithm that uses dynamical number of quorum participants based on the Paxos algorithm.

The topic of modernising consensus algorithms in large distributed systems and analysing failures within them is extremely relevant, as it demands the development of new solutions to improve the resilience and performance of modern systems while reducing the risks associated with various failures. This paper examines modern approaches to the modernisation of consensus algorithms.

2 Related Work

Consensus algorithms in the Paxos family have been extensively studied and optimized to address performance and scalability challenges in distributed systems. One of the most widely adopted variants is Multi-Paxos [8], which amortizes the cost of leader election by allowing a stable leader to propose multiple

values without repeating the full consensus protocol. This significantly reduces latency in steady-state operation and is commonly used in production systems.

Fast Paxos [10] improves upon classical Paxos by reducing the number of communication rounds required to reach consensus under certain conditions. By allowing proposers to bypass the leader in some cases, it achieves lower latency, though at the cost of increased complexity and a higher likelihood of conflicts.

Another algorithm is Flexible Paxos [11], in which communication does not involve all nodes or a fixed majority. Messages are sent only to a quorum of nodes, which can be chosen more flexibly. Flexible quorum configurations can improve availability and performance in distributed systems.

Building on these ideas, Fast Flexible Paxos [12] combines the reduced latency of Fast Paxos with the quorum flexibility of Flexible Paxos. This hybrid approach aims to achieve both lower latency and improved fault tolerance, although it introduces additional design and implementation complexity.

These variants highlight the ongoing effort to refine Paxos-based consensus mechanisms, balancing trade-offs between latency, fault tolerance, and system complexity. The algorithm proposed in this work builds upon these foundations, aiming to further modernize and optimize consensus in distributed environments.

3 Replicated State Machine

Large-scale distributed cloud systems often handle client data. Multiple computer nodes are frequently used to provide dependable data storage, modification, and deletion inside the system. Multiple nodes allow the system to tolerate failures and outages that impact only part of them. This enhances the issue of duplicating data over several nodes while preserving consistency. To ensure consistency among nodes, a technique for determining the order in which incoming messages are processed is required. Consensus acts as such a process. Various algorithms are used to reach consensus among nodes. Paxos [1] is one of the most well-known and commonly used.

It is important to bear in mind the FLP theorem of Fischer, Lynch and Paterson [13]. It states that it is impossible to achieve even a non-bounded consensus for $N > 2$ processes, even on a single bit, under the following conditions in a distributed system:

1. The system is asynchronous;
2. Any node in the system may experience a failure for an indefinite period of time;
3. The algorithm must terminate within a finite time;
4. The algorithm must be deterministic.

Thus, the theorem highlights the challenges of achieving reliable coordination in distributed systems in the presence of failures. Furthermore, in large cloud-based distributed systems, the order in which data arriving from the client is recorded may be important (for example, the order of financial transactions in banking systems), or it may be irrelevant (for example, non-sensitive data such as system process metrics).

4 Paxos

Paxos is a family of protocols designed to solve the consensus problem in distributed systems with unreliable computations. The creator of the family's leading algorithm, namely Paxos, is Leslie Lamport [1].

For simplicity, let us consider a record of a specific value in the system as the object for which the consensus problem is being solved. Each node in the distributed system can play one of three roles:

1. Proposer;
2. Acceptor;
3. Learner.

The proposer wants the value proposed by the client for the record to be selected in the system. The acceptors agree on the proposal and notify all learners of it. The learner wishes to know the accepted value for the record.

In the algorithm, the proposer attempts to put a proposal forward for approval by the acceptors. If another value has already been accepted, the proposer becomes aware of this and informs the learners. The proposal process is divided into two phases:

1. Preparation and promise:
   (a) The proposer receives a value v from the client and selects a unique round identifier p, then sends a message containing p to all acceptors to collect commitments;
   (b) Acceptors store p locally as the highest accepted proposal and respond to the proposer with a promise if they have not previously committed to any other proposer to participate in a round with a unique identifier $p' > p$. Otherwise, they refuse to participate in the vote with number p. If any of the acceptors has managed to accept value v' from a proposer with round number p', then the acceptor informs the proposer with round number p of a refusal with value v';
   (c) If a quorum (at least half of the acceptors) of promises is collected from the acceptors, the proposer proceeds to the second phase. Otherwise, they may attempt to select a unique round identifier greater than p and repeat the first phase.
2. Proposal and confirmation:
   (a) If the proposer has already received an accepted promise of participation in the voting from the acceptors, they must select the value with the highest unique identifier as v. If this has not occurred, the proposer is entitled to use their own proposal as v. The proposer sends a proposal with number p and v to the acceptors;
   (b) If an acceptor receives a proposal with number p, which is no less than the unique identifier of the round in which the acceptor promised to vote, the acceptor accepts the value v and notifies all learners. Otherwise, the acceptor rejects the proposal and responds to the proposer with the unique identifier $p' > p$ and the value v', which corresponds to p';

(c) Upon receiving a simple majority quorum (at least half of the acceptors) accepting the value from the acceptors, the proposer learns that the value has been accepted. Otherwise, the proposer may repeat the first phase with a unique identifier greater than p.

Paxos assures that once a value v is selected, no other value v' may be chosen. The procedure achieves consensus only if $[n/2]+1$ of the system's n nodes do not fail. The Paxos consensus algorithm addresses the issue of ending a consensus mechanism in the presence of malfunctioning nodes in an asynchronous system. Paxos, on the other hand, is a 'one-shot algorithm', which means that it does not expect to process more than one proposal from a client. This is precisely why, in real-world systems, Paxos is used to establish a consensus on the order of value acceptance inside cycles. Within these cycles, Paxos is employed to overcome the consensus problem. Generally, Paxos undergoes modifications in real-world systems and is not used in its classical form, as it is considered a theoretical algorithm [4].

Figure 1 illustrates an ideal implementation of Paxos for solving the consensus problem among system nodes. The following notation is used in Fig. 1: P—proposer, A—acceptor, L—learner, p—unique round identifier, v—value proposed by the proposer. The notes on Fig. 1 reflect the moments when each of the two phases began and ended. For better visual perception, only one learner is shown.

5 An Enhanced Paxos Algorithm – Orchestra-Paxos

Standard Paxos has a number of drawbacks. Two main issues stand out.

1. The inability to establish a quorum of acceptors in the first phase if more than half of the system's confirming nodes fail, as illustrated in Fig. 2;
2. Endless competition among several proposers in the first phase to assemble a quorum of acceptors, as shown in Fig. 3.

Whilst the second problem can be solved by configuring a random timeout following each failed attempt in the first phase, the first problem cannot be solved in a straightforward manner.

The Paxos algorithm has a number of variants, such as EPaxos [14], Flexible Paxos [11], Fast Flexible Paxos [12] and others. Each addresses a specific problem in Paxos, whilst also improving fault tolerance or performance.

One of the main modernised algorithms is Multi-Paxos. The difference from the classical version is that the proposer becomes the master after successfully passing the first phase. All subsequent proposals from the system's client go exclusively through this master. However, the drawback of this approach lies in the selection of a new master in the event of a failure of the current one. Furthermore, Paxos originally positions itself as an algorithm that does not depend on a master. Furthermore, performance studies of systems using Multi-Paxos demonstrate that the system's throughput is determined by the speed of

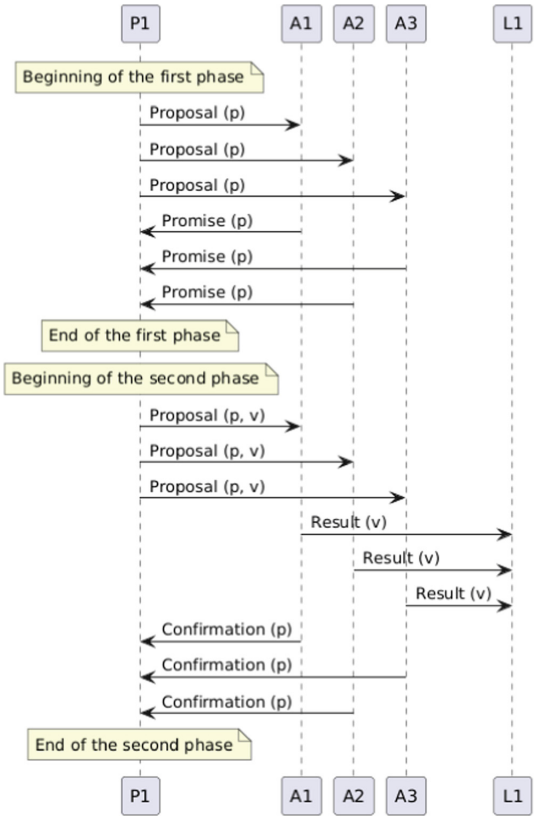

Fig. 1. Paxos's perfect execution

the slowest node. This, in turn, is a drawback of the classic Paxos algorithm as well.

Thus, two main problems of the standard Paxos algorithm are identified. It is for this reason that this paper proposes a modernised version of the Paxos algorithm—Orchestra-Paxos—which resolves the issues described above and is applied to the problem of replicated state machine in large distributed systems for values where the order of writing is not strictly mandatory. It also consists of two phases:

1. Preparation and promise:
   (a) The proposer selects a unique round identifier p and sends a message containing p to all acceptors to collect commitments. A timer is started to wait for commitments from the acceptors;
   (b) Acceptors store p locally as the highest accepted proposal and respond to the proposer with a promise if they have not previously committed to any other proposer to participate in a round with a unique identifier $p' > p$.

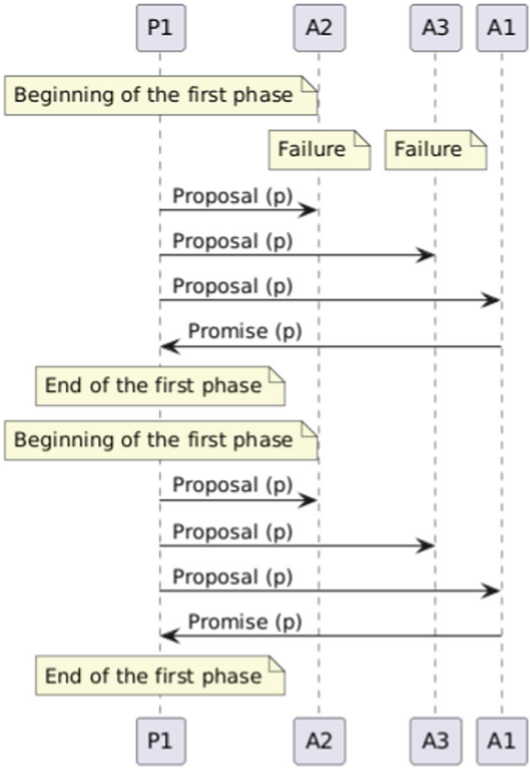

Fig. 2. Failure to gather a quorum of acceptors

Otherwise, they refuse to participate in the vote with number p and reply to the proposer with p' and the value v', if it was proposed earlier;

(c) The proposer, in the absence of replies refusing to participate in their vote, collects a set of acceptors (active acceptors) upon the expiry of the timer and proceeds to the second phase. Otherwise, the proposer may attempt to select a unique round identifier greater than p and repeat the first phase.

2. Proposal and confirmation:
   (a) The proposer sends an offer with number p, v and set of active acceptors to all members of the set of active acceptors. A timer is started to wait for the acceptors to accept the values;
   (b) If an acceptor receives a proposal with number p, which is no less than the unique identifier of the round in which the acceptor promised to vote, the acceptor takes the value v and notifies all learners about the acceptance of the value, and also transmits the set of active acceptors. The acceptor then notifies the proposer of the confirmation of the value. Otherwise, the acceptor rejects the proposal and replies to the proposer with a unique identifier $p' > p$ and a value v' corresponding to p';

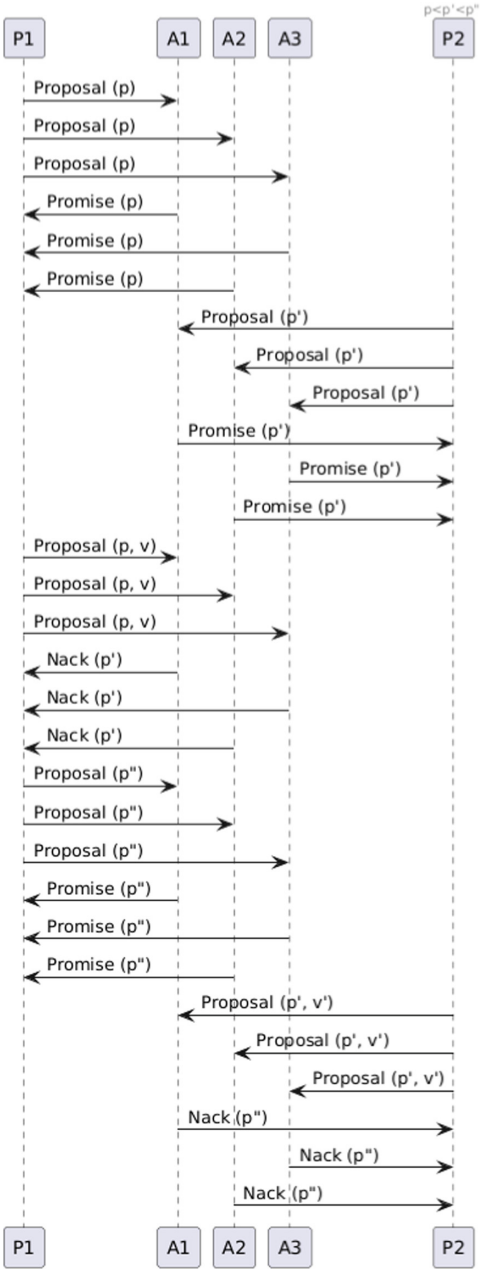

Fig. 3. Endless competition between several proposers

(c) Upon receiving a simple majority quorum (at least half of the acceptors) from the set of acceptors obtained in the first phase regarding the acceptance of the value, the proposer learns that the value has been accepted. If a proposer does not obtain a quorum of confirmations or the timer expires, they may repeat the first phase with a unique identifier after the expiry of a random timer that is greater than p. The set of active acceptors also becomes empty.

Thus, this algorithm resolves the problem of the impossibility of gathering a quorum in the first phase. Should such a situation arise, it is optionally possible to additionally notify the system participants responsible for collecting node survivability metrics. Furthermore, timers in both phases help to avoid waiting for lagging nodes, which has a direct impact on system performance. Thus, Orchestra-Paxos achieves consensus provided that $[n'/2] + 1$ nodes have not failed, where $n' \in N$ and N is the set of all nodes in the system. Thus, even if the set of active acceptors contains only a single node, it is possible to solve the consensus problem in the system. Furthermore, the second phase uses the classic Paxos algorithm, which also guaranties that consensus is reached.

Orchestra-Paxos is not leader-oriented due to the fact that each proposer's value has the exact same probability of being chosen in the round of replication.

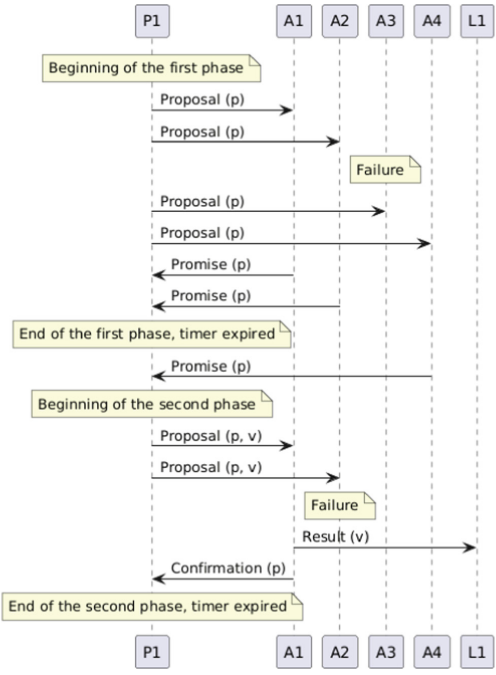

Fig. 4. Successful execution of Orchestra-Paxos

Figure 4 illustrates the execution of Orchestra-Paxos under the condition of failure and slowdown of some of the system's confirming nodes. In the first phase, it is shown that proposer $P1$ adds only $A1$ and $A2$ to the set of active acceptors. Acceptor $A3$ does not make it into the set of active acceptors because it has failed. Confirming node $A4$ is not included in the set of active acceptors, as it did not manage to respond with a confirmation to proposer $P1$ within the timer duration of the first phase. In the second phase, proposer $P1$ sends the proposal only to members of the set of active acceptors. Acceptor $A1$ successfully accepts the value v and reports this to the learner $L1$, then confirms its acceptance by sending a message to proposer $P1$. Acceptor $A2$ experiences a failure and therefore does not accept proposal v. Upon expiry of the second-phase timer, proposer $P1$ gathers a simple majority quorum (at least half of the acceptors), and thus successfully completes the round.

6 Preliminary Assessment

The Orchestra-Paxos algorithm was implemented in the *Go* programming language. To test the algorithm, the TCP/IP network protocol was used as the transport layer between the system's nodes. Furthermore, the following transport parameters were emulated:

1. The channel bandwidth was 1.25 megabits per second;
2. The round-trip time for each message was 20 ms;
3. The probability of message loss between nodes was configured using a parameter.

Each experiment used $3N$ nodes, where N is the number of proposers, acceptors and learners. Each proposer sent a new proposal once per second for 120 s. The timer for each phase of the algorithm was 100 ms. Furthermore, if the first phase needed to be repeated, the proposer waited between 15 and 20 s after an unsuccessful current phase. In the second phase the same method was used. This timer was chosen because the standard Go programming language library uses 10 connection recovery attempts at 1-s intervals. After 120 s had elapsed, the algorithm checked every 5 s whether the number of messages sent matched the number of messages received. If the proposals matched, the experiment was considered successfully completed, and the elapsed time was recorded. Messages from clients and proposers were always delivered. The experiments were conducted on a machine running *macOS*, with 48 GB of RAM and 14 cores.

Similar experiments were also conducted using Multi-Paxos. The master proposed N (N – number of acceptors and learners) proposals within 120 s.

Orchestra-Paxos was compared to Multi-Paxos on the following metrics:

1. Average time per round of the algorithm;
2. Waiting time for the completion of all replication rounds;
3. Number of rounds required to fulfill all client requests.

Table 1. Results of the experiments with three participants in each role

Algorithm	Probability of message loss %	Average duration of a single replication round sec	The time taken to complete all replication rounds sec	The number of rounds required to process all client requests
Orchestra-Paxos	0	0.35	213.277	603
Multi-Paxos	0	**0.3**	**121.276**	**399**
Orchestra-Paxos	0.01	0.39	245.27	629
Multi-Paxos	0.01	**0.33**	**130.730**	**401**
Orchestra-Paxos	5	0.36	255.313	702
Multi-Paxos	5	**0.29**	**135.323**	**468**
Orchestra-Paxos	30	0,58	826.239	**1432**
Multi-Paxos	30	**0,23**	**480.367**	1580
Orchestra-Paxos	60	0.53	**3003.12**	**5690**
Multi-Paxos	60	**0.21**	9452.631	45293

Table 2. Results of the experiments with five participants in each role

Algorithm	Probability of message loss %	Average duration of a single replication round sec	The time taken to complete all replication rounds sec	The number of rounds required to process all client requests
Orchestra-Paxos	0	0,19	420.332	2210
Multi-Paxos	0	**0.16**	**104.388**	**666**
Orchestra-Paxos	0.01	0.19	420.273	2197
Multi-Paxos	0.01	**0.16**	**120.282**	**779**
Orchestra-Paxos	5	0.2	445.321	2178
Multi-Paxos	5	**0.14**	**120.317**	**838**
Orchestra-Paxos	30	0,29	1030.845	3497
Multi-Paxos	30	**0.10**	**215.11**	**2078**
Orchestra-Paxos	60	0.51	**4072.325**	**7962**
Multi-Paxos	60	**0.09**	14997.527	178129

The results of experiments using Orchestra-Paxos and Multi-Paxos are shown in Table 1, with the number of proposers (except for the Multi-Paxos algorithm, where there is always a single proposer—the master node), acceptors and learners set to 3. Similarly, for 5 participants in each role, the results of the experiments are shown in Table 2. The results of the testing are averaged over three runs of each experiment. Moreover, the scalability of the algorithm is not the primary

focus of this work, so experiments were conducted on three and five participants in each of the roles in the system, which is generally a common replication factor.

The results of the experiment show that Multi-Paxos outperforms Orchestra-Paxos by a factor of three on average across the specified measures, assuming minimal message loss probability. Under high message loss probability, Orchestra-Paxos surpasses Multi-Paxos by a factor of twelve in terms of the number of rounds required and overall completion time. The improved technique is originally developed for a more flexible role allocation and, with more modifications, may perform better in scalable and fault-tolerant complicated network topologies.

An implementation of the Orchestra-Paxos algorithm in the Go programming language, along with test scenarios for verifying it and a comparison with the Multi-Paxos algorithm, are available on the GitHub version control system in the Orchestra-paxos repository [15].

7 The Disadvantages and the Future of the Algorithm

There are some disadvantages of the Orchestra-Paxos:

1. In the current implementation of the modernized algorithm, the system is in risk of segmentation;
2. The performance results of Orchestra-Paxos leave much to be desired, when the probability of message loss is minimal.

Future work will focus on extending the current algorithm by adding a group of replicated leaders that will check the availability of the system's nodes through a heartbeating mechanism, orchestrate the number of dynamic quorum participants and replicate data on revived nodes. The development direction aims to address the issue of system segmentation and ensure full consistency across the entire system. Moreover, adjusting the timer durations for each phase may positively impact the experimental results, potentially achieving performance comparable to Multi-Paxos or even surpassing it when the probability of message loss is minimal.

Lastly, after upgrading the algorithm, information about strict consistency and formal guaranties of the algorithm's correctness will be added to future research. In the present article, strict consistency is out of scope.

8 Conclusion

The research presented in this paper introduced Orchestra-Paxos, an enhanced consensus algorithm that replaces the fixed majority quorum requirements of traditional Paxos with a dynamic set of active participants. By utilizing timers to identify and exclude unresponsive nodes during the replication cycle, the algorithm ensures that consensus can be reached even when a majority of acceptors are temporarily unavailable. This approach effectively addresses the problem of

stalled progress in unstable networks, providing a robust foundation for high-availability distributed systems where the strict requirement for a fixed majority often leads to performance degradation or total system failure.

The experimental evaluation demonstrates that Orchestra-Paxos significantly enhances system resilience under adverse network conditions. Specifically, in environments with 60% message loss, the algorithm reduced the total completion time and the number of required replication rounds by an average factor of 12 compared to Multi-Paxos. Although Multi-Paxos maintains a performance advantage in stable network conditions due to its leader-oriented optimizations, Orchestra-Paxos proves to be a superior solution for large-scale, unpredictable environments. These results confirm that the flexibility of dynamic quorums allows for continued operations in scenarios where classic consensus models would otherwise reach a deadlock.

The Orchestra-Paxos algorithm has substantial stability features and optimisation potential, making it a compelling option for systems that demand high fault tolerance and predictability. A new generation of scalable, fault-tolerant, and high-performance consensus algorithms may now be developed thanks to dynamic quorum size, carefully chosen algorithm parameters, and pragmatic system architecture.

References

1. Lamport, L.: The part-time parliament. ACM Trans. Comput. Syst. **16**(2), 133–169 (1998)
2. Ongaro, D., Ousterhout, J.: In search of an understandable consensus algorithm. In: 2014 USENIX Conference on USENIX Annual Technical Conference (USENIX ATC'14), pp. 305–320. USENIX Association, Philadelphia, PA, USA (2014)
3. Baker, J., et al.: Megastore: providing scalable, highly available storage for interactive services. In: 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011), Proceedings Conference Innovative Data Syst Res (CIDR), Asilomar, CA, USA, pp. 224–234 (2011)
4. Chandra, T., Griesemer, R., Redstone, J.: Paxos made live: an engineering perspective. In: 26th ACM Symposium on Principles of Distributed Computing (PODC 2007), pp. 398–407. Association for Computing Machinery, New York, NY, USA (2007)
5. Lakshman, A., Malik, P.: Cassandra: a decentralized structured storage system. ACM SIGOPS Operating Syst. Rev. **44**(2), 35–40 (2010)
6. MySQL Blog Archive. https://dev.mysql.com/blog-archive/the-king-is-dead-long-live-the-king-our-homegrown-paxos-based-consensus/. Accessed 10 Mar 2026
7. Gunawi, H.-S., et al.: Why does the cloud stop computing? Lessons from hundreds of service outages. In: 7th ACM Symposium on Cloud Computing (SoCC 2016), pp. 1–16. Association for Computing Machinery, Santa Clara, CA, USA (2016)
8. Camargos, L.-J., Schmidt, R.-M., Pedone, F.: Multicoordinated Paxos. In: 26th ACM Symposium on Principles of Distributed Computing (PODC 2007), pp. 316–317. Association for Computing Machinery, New York, NY, USA (2007)
pg

9. Moraru, I., Andersen, D.-G., Kaminsky, M.: There is more consensus in egalitarian parliaments. In: 24th ACM Symposium on Operating Systems Principles (SOSP 2013), pp. 358–372. Association for Computing Machinery, New York, NY, USA (2013)
10. Lamport, L.: Fast Paxos. Distrib. Comput. **19**(2), 79–103 (2006)
11. Howard, H., Malkhi, D., Spiegelman, A.: Flexible Paxos: quorum intersection revisited. In: 20th International Conference on Principles of Distributed Systems (OPODIS 2016), pp. 25:1–25:14. Dagstuhl Publishing, Madrid, Spain (2016)
12. Howard, H., Charapko, A., Mortier, R.: Fast flexible Paxos: relaxing quorum intersection for fast Paxos. In: 22nd International Conference on Distributed Computing and Networking (ICDCN 2021), pp. 186–190. ACM Digital Library, Nara, Japan (2021)
13. Fischer, M.-J., Lynch, N.-A., Paterson, M.-S.: Impossibility of distributed consensus with one faulty process. J. ACM **32**(2), 374–382 (1985)
14. Ryabinin, F., Gotsman, A., Sutra, P.: Making democracy work: fixing and simplifying egalitarian Paxos. In: 29th International Conference on Principles of Distributed Systems (OPODIS 2025), pp. 22:1–22:19. Dagstuhl Publishing, Iaşi, Romania (2025)
15. Orchestra-Paxos code repository. https://github.com/ACE-777/orchestra-paxos. Accessed 11 Feb 2026

Computational and Applied Mathematics (CAM 2026)

A Mach-Number-Weighted Nonlinear Multiscale Viscosity Method for the Compressible Navier–Stokes Equations

Elaine Bernine(✉), Isaac P. Santos, and Lucia Catabriga

Federal University of Espírito Santo, Vitória, Brazil
elaine.bernine@edu.ufes.br, {isaac.santos,lucia.catabriga}@ufes.br

Abstract. In this work, the nonlinear multiscale viscosity finite element method (NMV-YZβ_ω) is proposed for the numerical solution of the compressible Navier–Stokes equations. The formulation is based on the decomposition of the problem into macro and microscales within the framework of Variational Multiscale Stabilized Methods (VMS). A nonlinear artificial viscosity operator is introduced at both discretization scales, constructed from the YZβ discontinuity-capturing parameter and incorporating a Mach-number-dependent weighting. The method is validated through classical benchmark problems, including spatial convergence analyses and evaluation of numerical accuracy. The results demonstrate robust convergence and optimal rates in the L^2 and H^1 Sobolev norms, with solutions comparable to those obtained using the NMV-YZβ and NMV-CAU methods. Computational performance is also evaluated under different flow regimes. The proposed method exhibits performance comparable to that of the NMV-YZβ formulation, which in turn outperforms standard stabilized variational methods such as SUPG+YZβ and CAU, as well as the multiscale stabilized variational method NMV-CAU.

Keywords: Finite element · Compressible Navier–Stokes equations · Variational Multiscale Stabilized (VMS) · Nonlinear Multiscale Viscosity

1 Introduction

The compressible Navier–Stokes equations constitute the fundamental mathematical model for fluid flows with density variations and arise in many engineering and scientific applications. Due to their nonlinear and coupled structure, analytical solutions are rarely available, making numerical methods essential for their solution. In this context, the Finite Element Method (FEM) has been widely employed to obtain approximate solutions, particularly when combined with stabilization strategies capable of handling convection-dominated regimes.

This work was partially funded by CAPES and FAPES.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 55–72, 2027.
https://doi.org/10.1007/978-3-032-30518-3_4

Standard Galerkin finite element formulations are generally inadequate for compressible flows because they may exhibit numerical instabilities, such as spurious oscillations, loss of stability at high Reynolds number regimes, and difficulties in capturing sharp gradients associated with boundary layers and shock waves.

To mitigate these challenges, several variational stabilization techniques have been developed over the last four decades. Among the most widely used approaches are the stabilized variational methods *Consistent Approximate Upwind* (CAU) [1] and *Streamline Upwind Petrov-Galerkin* (SUPG) method with the YZβ discontinuity-capturing operator (SUPG+YZβ) [17–19]. In addition, multiscale stabilized variational methods have been proposed, including NMVC/NMV-CAU [4,5] and NMV-YZβ [2,3,5].

The CAU method was originally developed by Galeão and Do Carmo [9] for the advection-diffusion equation and incorporates a nonlinear discontinuity-capturing operator that acts along the gradient direction of the approximate solution, improving stability in regions with strong gradients. Later, the formulation was extended to the compressible Euler and Navier–Stokes equations using entropy variables [1], and subsequently reformulated in conservative variables [8].

The SUPG+YZβ formulation for compressible flows was introduced in [17–19] to reduce oscillations in regions with large gradients. In this approach, an adaptive artificial viscosity operator dynamically adjusts the diffusion level increasing stabilization in regions with steep gradients while preserving accuracy in smooth regions.

An alternative framework for stabilized finite element formulations is the *Variational Multiscale Stabilized* (VMS) method proposed by Hughes [12]. In this approach, the solution is decomposed into resolved (macroscale) and unresolved (microscale) components, allowing the influence of the unresolved scales to be incorporated into the numerical formulation. Applications of the VMS approach to compressible flows can be found in [13].

Within this multiscale framework, the *Nonlinear Multiscale Viscosity* (NMV) methods introduce nonlinear artificial viscosity operators at both discretization scales. Formulations based on the YZβ discontinuity-capturing parameter were proposed in [2,3], and later applied to the compressible Navier–Stokes equations in [5]. Bernine et al. [4] proposed the NMV-CAU method, extending the NMV framework to compressible Navier-Stokes by defining artificial viscosity from the CAU discontinuity-capturing parameter.

Numerical studies reported in [2–5] indicate that NMV-YZβ and NMV-CAU produce results comparable to those obtained with SUPG+YZβ and CAU methods, respectively, while exhibiting improved computational performance in terms of GMRES iterations and CPU time for both Euler and compressible Navier-Stokes problems. Motivated by these results, this work proposes a new formulation, denoted NMV-YZβ_ω, belonging to the class of NMV methods. The method introduces a nonlinear artificial viscosity operator at both discretization scales, constructed from the YZβ discontinuity-capturing parameter and incorporating a Mach-number-dependent weighting.

The remainder of this paper is organized as follows. Section 2 presents the governing equations of the compressible Navier-Stokes system. Section 3 introduces the proposed stabilized multiscale formulations, namely NMV-YZβ and NMV-YZβ_ω. Section 4 describes the temporal discretization procedure. Section 5 presents numerical experiments for subsonic, sonic, and transonic flow regimes, analyzing spatial convergence, accuracy, and computational performance. Finally, Sect. 6 summarizes the main conclusions.

2 Mathematical Formulation

The governing compressible Navier–Stokes equations are formulated from the conservation laws of mass, linear momentum, and energy, given respectively by

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i}\left(\rho u_i\right) = 0, \quad \text{in } \Omega \times (0, t_f], \tag{1}$$

$$\frac{\partial (\rho u_i)}{\partial t} + \frac{\partial}{\partial x_j}\left(\rho u_i u_j + p\,\delta_{ij} - \sigma_{ij}\right) = f_i, \quad \text{in } \Omega \times (0, t_f], \tag{2}$$

$$\frac{\partial E_t}{\partial t} + \frac{\partial}{\partial x_i}\left[(E_t + p)u_i - \sigma_{ij}u_j + q_i\right] = 0, \quad \text{in } \Omega \times (0, t_f], \tag{3}$$

where ρ, u_i, p, σ_{ij}, f_i, E_t, and q_i, with $i, j = 1, 2$, denote, respectively, the density, the components of the velocity vector, the pressure, the components of the viscous stress tensor, the components of the body force per unit volume, the total energy and the components of the heat flux vector.

The viscous stress tensor is defined as $\sigma_{ij} = \lambda \frac{\partial u_k}{\partial x_k}\delta_{ij} + \mu\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)$, with $i, j, k = 1, 2$, where $\mu = \mu_0 \left(\frac{T}{T_0}\right)^{3/2} \frac{(T_0+S)}{(T+S)}$ is the dynamic viscosity coefficient, with S being the Sutherland constant, T the temperature and μ_0 and T_0 denoting reference values, and $\lambda = -\frac{2\mu}{3}$ is the second viscosity coefficient, considering the ideal gas (Stokes hypothesis).

When the fluid is considered as a perfect gas, and following the definitions presented in [16], we have $E_t = \rho\left(e_{in} + \frac{1}{2}u_i u_i\right) = \rho E$, which denotes the total energy per unit volume. Here, E is the total specific energy, $e_{in} = c_v T$ is the specific internal energy, $p = \rho R T$ is the pressure, $k = \gamma R \mu/(\gamma - 1)Pr$ is the thermal conductivity, c_v is the specific heat at constant volume, c_p is the specific heat at constant pressure, $\gamma = c_p/c_v$ is the adiabatic exponent, $R = c_p - c_v$ is the specific gas constant and Pr is the Prandtl number.

The compressible Navier–Stokes equations (1), (2) and (3) can be rewritten in vector form with coordinates $(x_1, x_2) \in \Omega \subset \mathbb{R}^2$ and time variable $t \in (0, t_f]$ as follows:

$$\frac{\partial \boldsymbol{U}}{\partial t} + \frac{\partial \boldsymbol{F}_i}{\partial x_i} - \frac{\partial \boldsymbol{G}_i}{\partial x_i} = \boldsymbol{S}, \tag{4}$$

where $\boldsymbol{U} = [\rho, \rho u_i, \rho E]^T$ is the vector of conservative variables and $\boldsymbol{S} = [0, f_i, 0]^T$ represents the body force per unit volume, with $i = 1, 2$. The Euler flux vectors,

$\boldsymbol{F}_i$ (advective transport terms), and the viscous flux vectors, $\boldsymbol{G}_i$ (diffusive transport terms), for $i = 1, 2$, are defined as:

$$\boldsymbol{F}_i = \begin{bmatrix} \rho u_i \\ \rho u_i u_1 + p\delta_{i1} \\ \rho u_i u_2 + p\delta_{i2} \\ (\rho E + p)\, u_i \end{bmatrix} \quad \text{and} \quad \boldsymbol{G}_i = \begin{bmatrix} 0 \\ \sigma_{1i} \\ \sigma_{2i} \\ \sigma_{ik} u_k - q_i \end{bmatrix},$$

where δ_{ij} is the Kronecker delta and q_i denotes the components of the heat flux vector. Furthermore, Eq. (4) can be rewritten in the quasi-linear form as

$$\frac{\partial \boldsymbol{U}}{\partial t} + \boldsymbol{A}_i \frac{\partial \boldsymbol{U}}{\partial x_i} - \frac{\partial}{\partial x_i}\left(\boldsymbol{K}_{ij} \frac{\partial \boldsymbol{U}}{\partial x_j}\right) = \boldsymbol{S}, \quad \text{in } \Omega \times (0, t_f], \tag{5}$$

where $\boldsymbol{A}_i = \dfrac{\partial \boldsymbol{F}_i}{\partial \boldsymbol{U}}$ denotes the Jacobian matrices associated with the advective transport in the x_i-direction, and $\boldsymbol{K}_{ij}$ represents the diffusion matrices such that $\boldsymbol{K}_{ij} \dfrac{\partial \boldsymbol{U}}{\partial x_j} = \boldsymbol{G}_i$, for $i, j = 1, 2$.

Related to Eq. (5), we consider the following initial and Dirichlet boundary conditions: $B\boldsymbol{U} = \boldsymbol{Z}$, on $\Gamma \times (0, t_f]$ and $\boldsymbol{U}(x_i, 0) = \boldsymbol{U}_0(x_i)$, with $i = 1, 2$, where Γ is the boundary of Ω, B is the boundary operator, and $\boldsymbol{Z}$ and $\boldsymbol{U}_0$ are known vector-valued functions.

3 Nonlinear Multiscale Variational Formulations

To apply the finite element method, we consider a disjoint partition of the domain Ω, denoted by $\mathcal{T}^h = \{\Omega^e\}_{e=1}^{nel}$, into nel linear triangular elements. For simplicity, only prescribed (Dirichlet) boundary conditions are considered. Let $\mathcal{S}^h$ denote the space of admissible functions and $\mathcal{V}^h$ the space of test functions, defined as follows:

$$\mathcal{S}^h = \{\boldsymbol{U}^h \in \left[H^1(\Omega)\right]^4; \quad \boldsymbol{U}^h|\Omega^e \in \left[\mathbb{P}^1(\Omega^e)\right]^4, \quad \boldsymbol{U}^h = \boldsymbol{Z} \text{ on } \Gamma\},$$
$$\mathcal{V}^h = \{\boldsymbol{W}^h \in \left[H^1(\Omega)\right]^4; \quad \boldsymbol{W}^h|\Omega^e \in \left[\mathbb{P}^1(\Omega^e)\right]^4, \quad \boldsymbol{W}^h = \boldsymbol{0} \text{ on } \Gamma\},$$

where $H^1(\Omega)$ denotes the Hilbert space of square-integrable functions whose first-order derivatives are also square-integrable, and $\mathbb{P}^1(\Omega^e)$ represents the set of linear polynomials defined on Ω^e.

The class of methods known as *Nonlinear Multiscale Variational* (NMV) formulations is characterized by the introduction of an additional term aimed at controlling the artificial viscosity of the system. The enriched solution and test spaces are defined as $\mathcal{S}^E = \mathcal{S}^h \oplus \mathcal{S}^b$ and $\mathcal{V}^E = \mathcal{V}^h \oplus \mathcal{V}^b$, where $\mathcal{S}^h$ and $\mathcal{V}^h$ denote the standard finite element spaces and $\mathcal{S}^b$ and $\mathcal{V}^b$ represent the bubble-function spaces, defined as follows:

$$\mathcal{S}^b = \{\boldsymbol{U}^b \in [H^1(\Omega)]^4;\ \boldsymbol{U}^b|_{\Omega^e} \in [\text{span}(\phi^b)]^4,\ \forall \Omega^e \in \mathcal{T}^h\},$$
$$\mathcal{V}^b = \{\boldsymbol{W}^b \in [H^1(\Omega)]^4;\ \boldsymbol{W}^b|_{\Omega^e} \in [\text{span}(\phi^b)]^4,\ \forall \Omega^e \in \mathcal{T}^h\}.$$

The cubic bubble function ϕ^b is defined on each element Ω^e as $\phi^b = 27N_1^e N_2^e N_3^e$, where $N_j^e = N_j^e(x, y)$ are the local basis functions associated with the space $\mathbb{P}^1(\Omega^e)$.

The NMV formulation for the compressible Navier–Stokes equations consists of finding $\boldsymbol{U}^E = \boldsymbol{U}^h + \boldsymbol{U}^b \in \mathcal{S}^E$, with $\boldsymbol{U}^h \in \mathcal{S}^h$ and $\boldsymbol{U}^b \in \mathcal{S}^b$, such that

$$\int_\Omega \left[\boldsymbol{W}^E \cdot \left(\frac{\partial \boldsymbol{U}^E}{\partial t} + \boldsymbol{A}_i^h \frac{\partial \boldsymbol{U}^E}{\partial x_i} \right) + \frac{\partial \boldsymbol{W}^E}{\partial x_i} \cdot \left(\boldsymbol{K}_{ij}^h \frac{\partial \boldsymbol{U}^E}{\partial x_j} \right) \right] d\Omega + \\ + \sum_{e=1}^{nel} \int_{\Omega^e} \delta_\nu(\boldsymbol{U}^h) \left(\frac{\partial \boldsymbol{W}^E}{\partial x_i} \cdot \frac{\partial \boldsymbol{U}^E}{\partial x_i} \right) d\Omega^e = \int_\Omega \boldsymbol{W}^E \cdot \boldsymbol{S}^E d\Omega, \tag{6}$$

for all $\boldsymbol{W}^E \in \mathcal{V}^E$, where $\boldsymbol{W}^E = \boldsymbol{W}^h + \boldsymbol{W}^b$, with $\boldsymbol{W}^h \in \mathcal{V}^h$ and $\boldsymbol{W}^b \in \mathcal{V}^b$. The term $\delta_\nu(\boldsymbol{U}^h)$ denotes a nonlinear parameter depending on $\boldsymbol{U}^h$, which regulates the amount of artificial viscosity introduced into the system.

3.1 Nonlinear Multiscale Variational Method NMV-YZβ

The NMV-YZβ method was initially introduced in [2,3] for the compressible Euler equations. Subsequently, in [5], its application was extended to the compressible Navier–Stokes equations. The variational formulation of this method is defined by Eq. 6, in which the artificial viscosity parameter is defined by the YZβ discontinuity-capturing parameter, that is, $\delta_\nu = \delta_{yz\beta}$, given by

$$\delta_{yz\beta}(\boldsymbol{U}^h) = \left\| \boldsymbol{Y}^{-1} \boldsymbol{R}_e(\boldsymbol{U}^h) \right\| \left(\sum_{i=1}^{2} \left\| \boldsymbol{Y}^{-1} \frac{\partial \boldsymbol{U}^h}{\partial x_i} \right\|^2 \right)^{\frac{\beta}{2}-1} \left\| \boldsymbol{Y}^{-1} \boldsymbol{U}^h \right\|^{1-\beta} \left(\frac{h_{yz\beta}}{2} \right)^{\beta}, \tag{7}$$

where $\boldsymbol{Y}$ is a diagonal matrix whose components are defined as $Y_r = U_r|_{\text{ref}}$, with $r = 1, 2, 3, 4$. The quantities $U_r|_{\text{ref}}$ represent the reference values of the conservative variables, assumed in this work as the initial condition of the problem. Moreover, $h_{yz\beta} = 2 \left(\sum_{a=1}^{3} |\mathbf{j} \cdot \nabla N_a| \right)$ denotes the local characteristic length of each element. Here, N_a represents the finite element shape function associated with node a, and $\mathbf{j} = \frac{\nabla \rho}{\|\nabla \rho\|}$ is the unit vector in the direction of the density gradient.

When considering $\beta = 1$ in the parameter $\delta_{yz\beta}$ defined in Eq. (7), the discontinuity-capturing operator introduces artificial viscosity into the system of equations with greater emphasis on regions with smooth gradients. On the other hand, when $\beta = 2$, the artificial viscosity acts predominantly in regions with steep gradients. Thus, in order to balance these effects, the parameter $\delta_{yz\beta}$ is defined as $\delta_{yz\beta}(\boldsymbol{U}^h) = \frac{1}{2} \left(\delta_{yz\beta}|_{\beta=1} + \delta_{yz\beta}|_{\beta=2} \right)$.

3.2 Nonlinear Multiscale Variational Method NMV-YZβ_ω

After studying the YZβ discontinuity-capturing parameter and the nonlinear artificial viscosity parameter dependent on the global Mach number, defined

in [3] for the Euler equations, the NMV-YZβ_ω method is proposed in this work for the compressible Navier–Stokes equations. This method consists of adding a nonlinear artificial viscosity term dependent on the Mach number in both scales of the problem. The variational formulation of the method is given by Eq. (6), with the artificial viscosity parameter defined as $\delta_\nu = \delta_\omega$, given by

$$\delta_\omega(\boldsymbol{U}^h) = (1 - \omega(M))\, \delta_{yz\beta}|_{\beta=1} + \omega(M)\delta_{yz\beta}|_{\beta=2},$$

where

$$\omega(M) = \frac{1}{4}\left(1 + \frac{1}{1 + \exp(-40\,(M-2))} + \frac{1}{1 + \exp(-40\,(M-0.3))}\right). \tag{8}$$

The function $\omega(M)$ is introduced to weight the action of the artificial viscosity added to the system, establishing a combination between the YZβ discontinuity-capturing parameters defined in Eq. (7) for $\beta = 1$ and $\beta = 2$. Its formulation aims to modulate the intensity of the stabilization according to the flow regime in each element of the domain. For $M < 0.3$, the function intensifies the contribution associated with regions with smooth gradients. Conversely, for $M > 2$, it increases the contribution related to regions characterized by steeper gradients. In the intermediate interval $0.3 \leq M \leq 2$, the weighting function promotes a smooth transition between these two behaviors, as shown in Fig. 1.

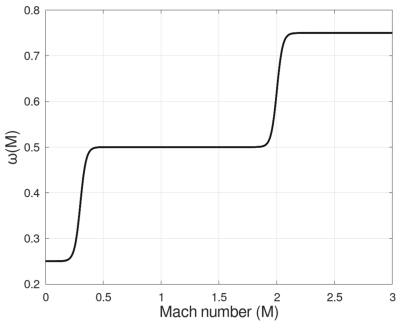

Fig. 1. Weighting function $\omega(M)$ defined in Eq. (8), considering $M \in [0, 3]$.

This weighting provides a smooth transition of the artificial viscosity across the different flow regimes. Consequently, the artificial viscosity parameter varies continuously with respect to M, ensuring a gradual transition between the different stabilization regimes and avoiding abrupt variations of the artificial viscosity parameter.

3.3 Local Matrices of the Variational Formulation

By applying the finite element approximations in the variational formulations of the NMV-YZβ_ω and NMV-YZβ methods, considering a triangular mesh and linear polynomial approximation, a local nonlinear system of differential-algebraic

equations is obtained, defined as follows:

$$\begin{bmatrix} \boldsymbol{M}^e_{hh} & \boldsymbol{M}^e_{hb} \\ \boldsymbol{M}^e_{bh} & \boldsymbol{M}^e_{bb} \end{bmatrix} \begin{bmatrix} \dot{\boldsymbol{\mathcal{U}}}_h \\ \dot{\boldsymbol{\mathcal{U}}}_b \end{bmatrix} + \begin{bmatrix} \boldsymbol{N}^e_{hh} & \boldsymbol{N}^e_{hb} \\ \boldsymbol{N}^e_{bh} & \boldsymbol{N}^e_{bb} \end{bmatrix} \begin{bmatrix} \boldsymbol{\mathcal{U}}_h \\ \boldsymbol{\mathcal{U}}_b \end{bmatrix} = \begin{bmatrix} \boldsymbol{S}_h \\ \boldsymbol{S}_b \end{bmatrix}, \tag{9}$$

where $\boldsymbol{\mathcal{U}}_h$ and $\boldsymbol{\mathcal{U}}_b$ are the vectors of nodal values associated with $\boldsymbol{U}^h$ and $\boldsymbol{U}^b$, respectively, for each element Ω^e, whereas $\dot{\boldsymbol{\mathcal{U}}}_h$ and $\dot{\boldsymbol{\mathcal{U}}}_b$ denote the corresponding vectors of time derivatives. The vectors $\boldsymbol{S}_h$ and $\boldsymbol{S}_b$ represent the nodal values of the source term components associated with the macro and microscales, respectively. The element mass matrices associated with Eq. (9) are defined as

$$\boldsymbol{M}^e_{pq} : \int_{\Omega^e} \boldsymbol{W}^p \cdot \frac{\partial \boldsymbol{U}^q}{\partial t}\, d\Omega^e, \qquad p, q \in \{h, b\},$$

and the element stiffness matrices are given by

$$\boldsymbol{N}^e_{pq} : \int_{\Omega^e} \left[\boldsymbol{W}^p \cdot \left(\boldsymbol{A}^h_i \frac{\partial \boldsymbol{U}^q}{\partial x_i} \right) + \frac{\partial \boldsymbol{W}^p}{\partial x_i} \cdot \left(\boldsymbol{K}^h_{ij} \frac{\partial \boldsymbol{U}^q}{\partial x_j} \right) + \delta_\nu(\boldsymbol{U}^h) \left(\frac{\partial \boldsymbol{W}^p}{\partial x_i} \cdot \frac{\partial \boldsymbol{U}^q}{\partial x_i} \right) \right] d\Omega^e,$$

with $p, q \in \{h, b\}$. Finally, the element load vectors are defined as

$$\boldsymbol{S}^e_p : \int_{\Omega^e} \boldsymbol{W}^p \cdot \boldsymbol{S}^p\, d\Omega^e, \qquad p \in \{h, b\}.$$

4 Time Integration Scheme

To advance the numerical solution in time for the compressible Navier–Stokes equations, given by Eq. (9), a time integration scheme is used based on the implicit predictor-multicorrector method described in [11]. This scheme adopts the formulation of the α-method, which establishes the relationships between the vectors of unknowns and their time discretized derivatives.

Using the relations of the generalized trapezoidal rule associated with the α-method, we obtain

$$\begin{aligned} \dot{\boldsymbol{U}}^{n+1,i+1}_p &= \dot{\boldsymbol{U}}^{n+1,i}_p + \Delta\dot{\boldsymbol{U}}^{n+1,i+1}_p, \\ \boldsymbol{U}^{n+1,i+1}_p &= \boldsymbol{U}^{n+1,i}_p + \alpha\Delta t \Delta\dot{\boldsymbol{U}}^{n+1,i+1}_p, \qquad p \in \{h, b\}. \end{aligned}$$

Thus, the system of Eqs. (9) can be rewritten in the form

$$\boldsymbol{M}_{hh}\dot{\boldsymbol{U}}^{n+1,i+1}_h + \boldsymbol{M}_{hb}\dot{\boldsymbol{U}}^{n+1,i+1}_b + \boldsymbol{N}_{hh}\boldsymbol{U}^{n+1,i+1}_h + \boldsymbol{N}_{hb}\boldsymbol{U}^{n+1,i+1}_b = \boldsymbol{S}^{n+1}_h, \tag{10}$$

$$\boldsymbol{M}_{bh}\dot{\boldsymbol{U}}^{n+1,i+1}_h + \boldsymbol{M}_{bb}\dot{\boldsymbol{U}}^{n+1,i+1}_b + \boldsymbol{N}_{bh}\boldsymbol{U}^{n+1,i+1}_h + \boldsymbol{N}_{bb}\boldsymbol{U}^{n+1,i+1}_b = \boldsymbol{S}^{n+1}_b. \tag{11}$$

Using the α-method relations in Eqs. (10) and (11), we obtain, respectively,

$$\boldsymbol{M}_1\Delta\dot{\boldsymbol{U}}^{n+1,i+1}_h + \boldsymbol{N}_1\Delta\dot{\boldsymbol{U}}^{n+1,i+1}_b = -\boldsymbol{R}_1, \tag{12}$$

$$\boldsymbol{M}_2\Delta\dot{\boldsymbol{U}}^{n+1,i+1}_h + \boldsymbol{N}_2\Delta\dot{\boldsymbol{U}}^{n+1,i+1}_b = -\boldsymbol{R}_2, \tag{13}$$

where

$$\boldsymbol{R}_1 = \left(\boldsymbol{M}_{hh}\dot{\boldsymbol{U}}_h^{n+1,i} + \boldsymbol{M}_{hb}\dot{\boldsymbol{U}}_b^{n+1,i}\right) + \left(\boldsymbol{N}_{hh}U_h^{n+1,i} + \boldsymbol{N}_{hb}U_b^{n+1,i}\right) - \boldsymbol{S}_h^{n+1},$$
$$\boldsymbol{R}_2 = \left(\boldsymbol{M}_{bh}\dot{\boldsymbol{U}}_h^{n+1,i} + \boldsymbol{M}_{bb}\dot{\boldsymbol{U}}_b^{n+1,i}\right) + \left(\boldsymbol{N}_{bh}U_h^{n+1,i} + \boldsymbol{N}_{bb}U_b^{n+1,i}\right) - \boldsymbol{S}_b^{n+1},$$
$$\boldsymbol{M}_1 = \boldsymbol{M}_{hh} + \alpha\Delta t\boldsymbol{N}_{hh}, \quad \boldsymbol{M}_2 = \boldsymbol{M}_{bh} + \alpha\Delta t\boldsymbol{N}_{bh},$$
$$\boldsymbol{N}_1 = \boldsymbol{M}_{hb} + \alpha\Delta t\boldsymbol{N}_{hb}, \quad \boldsymbol{N}_2 = \boldsymbol{M}_{bb} + \alpha\Delta t\boldsymbol{N}_{bb}.$$

Performing static condensation on Eqs. (13) yields

$$\Delta\dot{\boldsymbol{U}}_b^{n+1,i+1} = -\boldsymbol{N}_2^{-1}\left(\boldsymbol{R}_2 + \boldsymbol{M}_2\Delta\dot{\boldsymbol{U}}_h^{n+1,i+1}\right), \tag{14}$$

where $\boldsymbol{N}_2$ is an invertible 4×4 matrix at the element level. Substituting Eq. (14) into Eqs. (12) leads to the condensed system

$$\boldsymbol{M}^*\Delta\dot{\boldsymbol{U}}_h^{n+1,i+1} = \boldsymbol{F}^*,$$

where $\boldsymbol{M}^* = \left(\boldsymbol{M}_1 - \boldsymbol{N}_1\boldsymbol{N}_2^{-1}\boldsymbol{M}_2\right)$ and $\boldsymbol{F}^* = -\boldsymbol{R}_1 + \boldsymbol{N}_1\boldsymbol{N}_2^{-1}\boldsymbol{R}_2$.

Algorithm 1 describes the steps of the implicit predictor-multicorrector method applied to the system of equations obtained from the NMV formulations. In this context, Δt is the time step, α denotes the time integration parameter, the subscript n corresponds to the time step of the solution, i is the nonlinear iteration index, i_{MAX} is the maximum number of nonlinear iterations, t_f is the final time, and $\Delta\dot{\boldsymbol{U}}_h$ represents the increment of the time derivative during the nonlinear iteration. Since linear triangular elements are employed, the element matrices are computed explicitly, avoiding the need for numerical quadrature. In this work, $\alpha = 0.5$ is adopted in all experiments, leading to a second-order accurate time discretization for both the macro- and microscales.

5 Numerical Results

In this section, classical numerical experiments of fluid flows governed by the compressible Navier–Stokes equations are presented in order to analyze spatial convergence, evaluate the accuracy of the numerical approximation, and validate the implementation of the NMV-YZβ_ω method. To this end, the spatial convergence rate is investigated using the planar Couette problem, which admits a known analytical solution [14], considering the $L^2(\Omega)$ and $H^1(\Omega)$ Sobolev norms. In addition, numerical results are presented for the following benchmark problems: planar Couette flow, the Carter flat plate problem, and the flow over the NACA0012 airfoil, considering both transonic flow with $M_0 = 0.85$ and subsonic flow with $M_0 = 0.5$. The results obtained with the proposed NMV-YZβ_ω method are compared with those of the NMV-YZβ method in order to assess the effectiveness of the proposed approach.

In all experiments, the fluid is assumed to be a compressible gas with specific heat at constant volume $c_v = 717.5$ J/(kg·K), specific heat at constant pressure

Algorithm 1. NMV Predictor-Multicorrector Algorithm

Input: $\boldsymbol{U}_h^0$, $\dot{\boldsymbol{U}}_h^0$, $\boldsymbol{U}_b^0$ e $\dot{\boldsymbol{U}}_b^0$;
$t = 0$;
repeat ▷ Prediction phase
$t = t + \Delta t$;
$\boldsymbol{U}_h^{n+1,0} = \boldsymbol{U}_h^n + (1-\alpha)\,\Delta t\dot{\boldsymbol{U}}_h^n$; $\quad \dot{\boldsymbol{U}}_h^{n+1,0} = 0$;
$\boldsymbol{U}_b^{n+1,0} = \boldsymbol{U}_b^n + (1-\alpha)\,\Delta t\dot{\boldsymbol{U}}_b^n$; $\quad \dot{\boldsymbol{U}}_b^{n+1,0} = 0$;
$i = 0$;
repeat ▷ Multicorrection phase
▷ Residual calculation

$$\boldsymbol{R}_1^{n+1,i} = \left(\boldsymbol{M}_{hh}\dot{\boldsymbol{U}}_h^{n+1,i} + \boldsymbol{M}_{hb}\dot{\boldsymbol{U}}_b^{n+1,i}\right) + \left(\boldsymbol{N}_{hh}\boldsymbol{U}_h^{n+1,i} + \boldsymbol{N}_{hb}\boldsymbol{U}_b^{n+1,i}\right) - \boldsymbol{S}_h^{n+1};$$

$$\boldsymbol{R}_2^{n+1,i} = \left(\boldsymbol{M}_{bh}\dot{\boldsymbol{U}}_h^{n+1,i} + \boldsymbol{M}_{bb}\dot{\boldsymbol{U}}_b^{n+1,i}\right) + \left(\boldsymbol{N}_{bh}\boldsymbol{U}_h^{n+1,i} + \boldsymbol{N}_{bb}\boldsymbol{U}_b^{n+1,i}\right) - \boldsymbol{S}_b^{n+1};$$

$$\boldsymbol{M}_1 = \boldsymbol{M}_{hh} + \alpha\Delta t\boldsymbol{N}_{hh}; \qquad \boldsymbol{M}_2 = \boldsymbol{M}_{bh} + \alpha\Delta t\boldsymbol{N}_{bh};$$
$$\boldsymbol{N}_1 = \boldsymbol{M}_{hb} + \alpha\Delta t\boldsymbol{N}_{hb}; \qquad \boldsymbol{N}_2 = \boldsymbol{M}_{bb} + \alpha\Delta t\boldsymbol{N}_{bb};$$
$$M^* = \boldsymbol{M}_1 - \boldsymbol{N}_1\boldsymbol{N}_2^{-1}\boldsymbol{M}_2; \qquad F^* = -\boldsymbol{R}_1 + \boldsymbol{N}_1\boldsymbol{N}_2^{-1}\boldsymbol{R}_2;$$
$M^*\Delta\dot{\boldsymbol{U}}_h^{n+1,i+1} = F^*$; ▷ Preconditioned GMRES solver
▷ Correction phase

$$\boldsymbol{U}_h^{n+1,i+1} = \boldsymbol{U}_h^{n+1,i} + \alpha\Delta t\Delta\dot{\boldsymbol{U}}_h^{n+1,i+1};$$
$$\dot{\boldsymbol{U}}_h^{n+1,i+1} = \dot{\boldsymbol{U}}_h^{n+1,i} + \Delta\dot{\boldsymbol{U}}_h^{n+1,i+1};$$
$$\Delta\dot{\boldsymbol{U}}_b^{n+1,i+1} = -\boldsymbol{N}_2^{-1}\left(\boldsymbol{R}_2^{n+1,i} + \boldsymbol{M}_2\Delta\dot{\boldsymbol{U}}_h^{n+1,i+1}\right);$$
$$\boldsymbol{U}_b^{n+1,i+1} = \boldsymbol{U}_b^{n+1,i} + \alpha\Delta t\Delta\dot{\boldsymbol{U}}_b^{n+1,i+1};$$
$$\dot{\boldsymbol{U}}_b^{n+1,i+1} = \dot{\boldsymbol{U}}_b^{n+1,i} + \Delta\dot{\boldsymbol{U}}_b^{n+1,i+1};$$

$i = i + 1$;
until $i < i_{MAX}$
$n = n + 1$;
until $t < t_f$
Output: $\boldsymbol{U}_h^{t_f}$, $\dot{\boldsymbol{U}}_h^{t_f}$, $\boldsymbol{U}_b^{t_f}$ e $\dot{\boldsymbol{U}}_b^{t_f}$;

$c_p = 1004.5$ J/(kg · K), gas constant $R = 287$ J/(kg · K), adiabatic exponent $\gamma = 1.4$, and Prandtl number $Pr = 0.72$. The numerical simulations are performed using the GMRES solver with a maximum of 100 cycles, 30 vectors in the Krylov basis, a tolerance of 10^{-5}, three nonlinear corrections iterations (i_{MAX}), a BlockDiag preconditioner, and matrix-vector products computed using the element-by-element (EBE) storage structure, as considered in [8].

The reference values $U_i|_{\text{ref}}$, used in the construction of the YZβ discontinuity-capturing operator, are taken in all experiments as the values prescribed at the inlet boundary Γ_1, as indicated in the figures of the following sections.

The code was developed in the C programming language within the framework of a finite element library under continuous development at the Laboratory of Optimization and Computational Modeling of the Federal University of

Espírito Santo (LabOtiM/UFES[1]). Compilation and execution were performed using GCC (GNU Compiler Collection) version 11.4.0, with the -Ofast flag to maximize performance through aggressive optimizations, and the -Wall flag to enable all compiler warnings, thereby increasing the efficiency and robustness of the code. The numerical experiments were carried out on a machine equipped with an Intel® Core™ i5-1035G1 @ 1.00GHz processor, 8 GB of RAM, and Ubuntu 20.04.6 LTS operating system.

5.1 Spatial Convergence Analysis

To evaluate the spatial convergence rates, the two-dimensional planar Couette flow problem [14] is considered on the domain $\Omega = (0,1) \times (0,1)$. This configuration consists of a stationary lower plate and an upper plate moving with uniform velocity, as illustrated in Fig. 2. For this case, the fluid viscosity is defined as $\mu = 0.000018$ and the thermal conductivity as $\kappa = 0.0251125\ \mathrm{W\,m^{-1}\,K^{-1}}$.

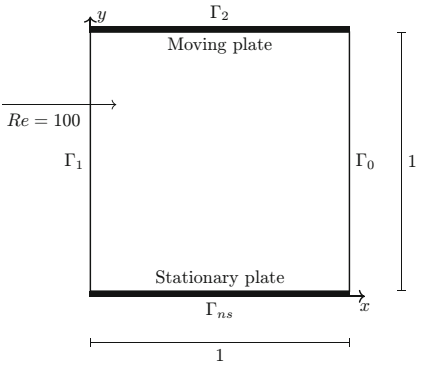

Fig. 2. Description of the planar Couette flow problem.

The boundary conditions are specified as follows: on Γ_1, the conservative variables are prescribed according to the exact solution of the problem, given by $\rho = T^{-1}$, $\rho u_1 = yT^{-1}$, $\rho u_2 = 0$, $\rho E = 2.5 + 0.5y^2T^{-1}$ and $T = 0.8 + 0.2y + \frac{\mu y^3}{2\kappa}(1-y)$; on Γ_2, the variables are prescribed as $\{U_1, U_2, U_3, U_4\} = \{1, 1, 0, 3\}$; on Γ_{ns}, all variables are prescribed as $\{U_1, U_2, U_3, U_4\} = \{1.25, 0, 0, 2.5\}$; and on Γ_0, a homogeneous Neumann (zero-flux) boundary condition is imposed.

Figures 3a and 3b present the spatial convergence rates of the density in the planar Couette problem, measured in the $L^2(\Omega)$ and $H^1(\Omega)$ norms using the NMV-YZβ_ω and NMV-YZβ methods, respectively. The simulations are performed on structured meshes with triangular elements of sizes $h \in$

[1] All experiment codes are available at https://github.com/mod-comp-ufes/FEM_CODES.

$\{2^{-5}, 2^{-6}, 2^{-7}, 2^{-8}\}$. The error norms between the exact and approximate solutions are evaluated at the final time $t_f = 10^{-3}$ seconds, considering a single time step of $\Delta t = 10^{-3}$, in order to eliminate the influence of the temporal error in the estimation of the spatial error. The results show that both the NMV-YZβ_ω and NMV-YZβ methods achieve optimal convergence rates close to 2 in the $L^2(\Omega)$ norm and 1 in the $H^1(\Omega)$ norm. These results confirm the expected asymptotic behavior of the methods and demonstrate their accuracy and numerical stability.

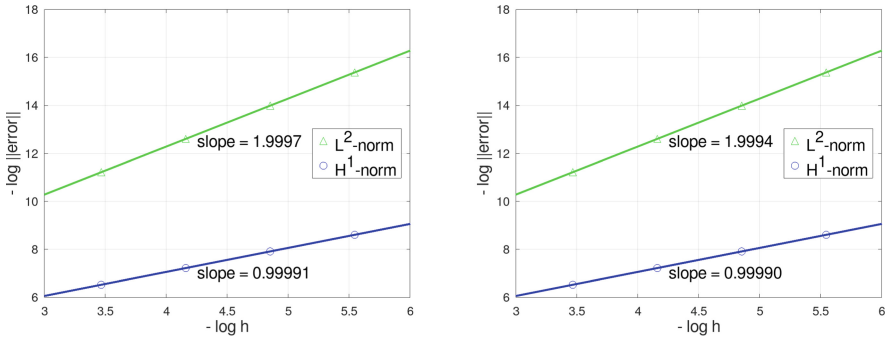

(a) NMV-YZβ_ω. (b) NMV-YZβ.

Fig. 3. Spatial convergence rates of the density error in the Couette problem measured in the L^2 and H^1 norms.

5.2 Carter Flat Plate Problem

The Carter flat plate problem is considered in over a two-dimensional domain $\Omega = (-0.2, 1.2) \times (0, 0.8)$ with Reynolds number $Re = 1000$ and Mach number $M_0 = 3$. The flat plate extends over the interval $0 \leq x \leq 1.2$, generating a boundary layer and a shock structure near the leading edge, as shown in Fig. 4. The thermal conductivity is taken as $\kappa = 1.395138\ \mathrm{W\,m^{-1}\,K^{-1}}$, and the temperature-dependent viscosity is modeled using Sutherland's law, $\mu = \dfrac{0.0906T^{1.5}}{T + 0.0001406}$, as described in [6,15].

The boundary conditions are prescribed as follows: on Γ_1, all variables are prescribed as $\{U_1, U_2, U_3, U_4\} = \{1, 1, 0, 0.698675\}$, with reference temperature $T_0 = 2.769 \times 10^{-4}$; on Γ_2, homogeneous Neumann boundary conditions are applied, with $U_3 = 0$; on Γ_{ns}, a no-slip condition is imposed, that is, $U_2 = U_3 = 0$, and a stagnation temperature $T_{\mathrm{stag}} = 7.7532 \times 10^{-4}$ is prescribed; and on Γ_0, a zero-flux boundary condition is imposed.

The numerical solutions were obtained at the final time $t_f = 3.5$ s using a structured mesh with 11421 nodes and 22400 triangular elements of size $h = 0.01$. Figures 5a and 5b show the density distributions obtained with the NMV-YZβ_ω and NMV-YZβ methods, respectively. The obtained results are consistent with

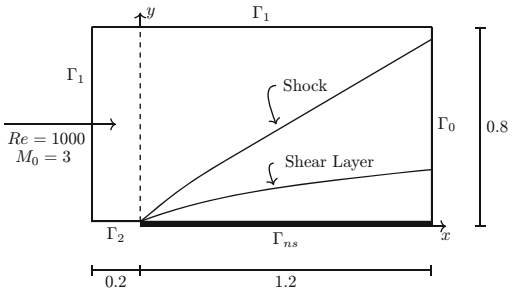

Fig. 4. Description of the Carter flat plate problem.

those reported in [1], which employed the CAU stabilization method, and in [4,5], where the nonlinear variational multiscale method NMV-CAU was used.

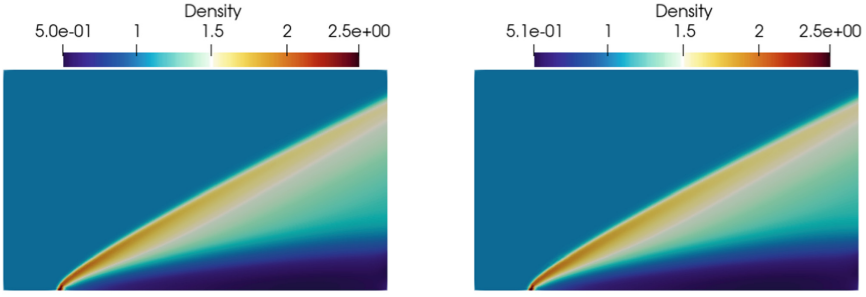

(a) NMV-YZβ_ω. (b) NMV-YZβ.

Fig. 5. Numerical results for density distributions in the Carter flat plate problem.

Figure 6a presents the density profiles for the Carter flat plate problem at $x = 1$, considering the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods. The numerical results are compared with those reported in [7], obtained using the finite difference scheme of [10] and a mesh with triangular elements of size $h = 0.015$. The solutions obtained with the NMV-YZβ_ω and NMV-YZβ methods exhibit strong similarity, highlighting the consistency of the NMV-YZβ_ω multiscale formulation. Furthermore, these solutions are closer to the reference solution reported in [7] than the NMV-CAU solution.

Figure 6b shows the time evolution of the Euclidean norm of the density residuals for the Carter flat plate problem, considering the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods. The NMV-YZβ_ω and NMV-YZβ methods present residuals with similar magnitude, with the NMV-YZβ_ω method achieving a greater reduction during the initial time steps. After approximately 1800 time steps, the NMV-YZβ method becomes more effective, producing smaller residu-

als. In contrast, the NMV-CAU method presents the largest residual magnitudes for this case.

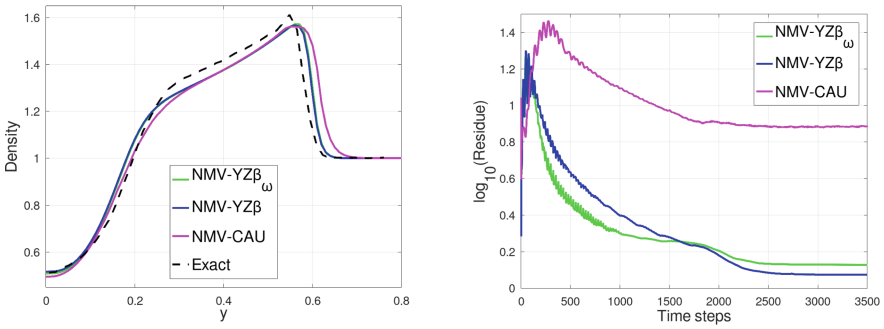

(a) Density profile at $x = 1$. (b) Temporal variation of the residual.

Fig. 6. Numerical results for the Carter flat plate problem.

The computational performance results are presented in Table 1 for the Carter flat plate problem, using the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods. The table summarizes the number of GMRES iterations and the CPU time spent for each method. The results show that the NMV-YZβ_ω method achieved the best computational performance, requiring 1198 fewer GMRES iterations than the NMV-YZβ method, while the NMV-CAU method presented the largest number of iterations among the methods analyzed. However, the CPU time is similar for all methods, being slightly lower for the NMV-YZβ_ω method.

Table 1. Computational performance of the Carter flat plate problem.

Methods	GMRES Iterations	CPU Time (s)
NMV-YZβ_ω	63051	1334.443
NMV-YZβ	64249	1346.938
NMV-CAU	67561	1368.057

5.3 Viscous NACA0012 Airfoil

The viscous NACA0012 airfoil problem is a two-dimensional case with a domain defined by a circle of radius 15, centered at the origin, in which the flow passes over the airfoil with an angle of attack of 0°, as illustrated in Fig. 7a. The aerodynamic profile considered in this study is defined by

$$y_t(x) = \pm 5t \left[0.2969\sqrt{x} - 0.1260x - 0.3516x^2 + 0.2843x^3 - 0.1015x^4\right], \quad (15)$$

where $x \in [0, 1]$ denotes the normalized position along the chord, and $t = 0.12$ represents the maximum thickness-to-chord ratio for the NACA0012 airfoil, as illustrated in Fig. 7b.

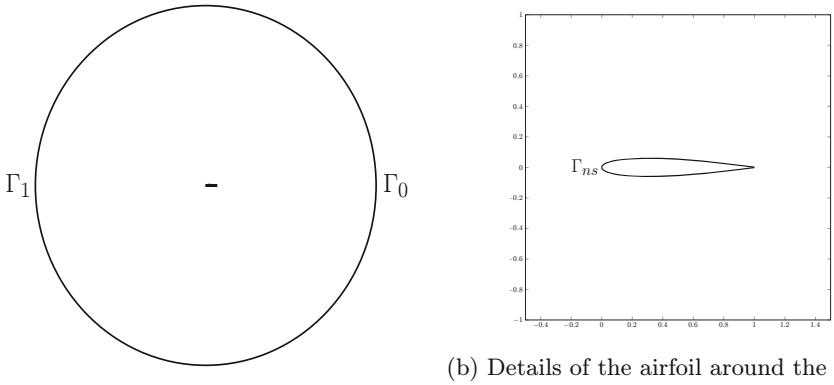

(a) General description.

(b) Details of the airfoil around the profile defined by Eq. (15).

Fig. 7. Description of the NACA0012 airfoil problem.

The boundary conditions are specified as follows: on Γ_1, corresponding to the region with $x \leq 0$ and $x^2 + y^2 = 15^2$, the variables $\{U_2, U_3, U_4\}$ are prescribed as $\{1, 0, c_v^* + 0.5\}$; on Γ_{ns}, defined by $y_t(x)$ in Eq. (15), no-slip conditions are imposed with $U_2 = U_3 = 0$ and stagnation temperature $T_{\text{stag}} = T_0 \left(1 + \frac{\gamma-1}{2} M_0^2\right)$; on Γ_0, corresponding to the region with $x > 0$ and $x^2 + y^2 = 15^2$, only the density is prescribed as $\rho = 1\ kg/m^3$. The dimensionless specific heat at constant volume depends on the inlet Mach number and is given by $c_v^* = \left(0.56 M_0^2\right)^{-1}$.

The numerical results for the viscous flow around the NACA0012 airfoil are presented for both the transonic regime with $M_0 = 0.85$ and the subsonic regime with $M_0 = 0.5$. The numerical solutions were obtained at the final time $t_f = 20$ seconds using an unstructured mesh composed of 2982 nodes and 5631 triangular elements.

Figures 8a and 8b present the numerical solutions for the viscous NACA0012 airfoil in the transonic regime with $M_0 = 0.85$, showing the density distributions obtained using the NMV-YZβ_ω and NMV-YZβ methods, respectively. The solutions obtained with both methods are similar. These numerical results align closely with those reported in [5,15].

Figure 10a shows the temporal evolution of the Euclidean norm of the density residuals for the viscous transonic NACA0012 airfoil problem with $M_0 = 0.85$, considering the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods. The NMV-YZβ method presents the smallest residual magnitudes, whereas the NMV-CAU method produces the largest residuals. Overall, all three methods produce residuals of small magnitude.

The computational performance results for the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods in the viscous transonic flow over the NACA0012 airfoil

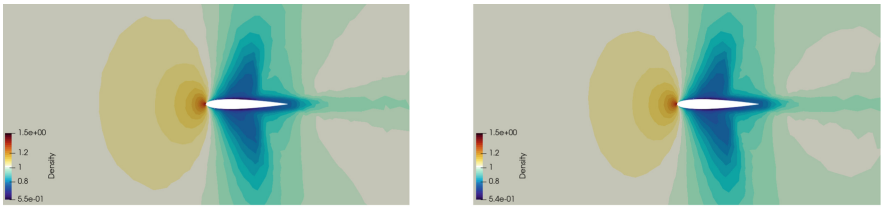

(a) NMV-YZβ_ω. (b) NMV-YZβ.

Fig. 8. Numerical results for the density distributions in the transonic NACA0012 airfoil problem with $M_0 = 0.85$.

are presented in Table 2. The results show that the NMV-YZβ method required 397 fewer GMRES iterations than the NMV-YZβ_ω method, while the NMV-CAU method presented the largest number of iterations among the analyzed methods. However, the CPU time for the NMV-YZβ_ω method is the lowest, being approximately 2.5% lower than that obtained with the NMV-CAU method, which presented the highest CPU time.

Table 2. Computational performance for the transonic NACA0012 airfoil problem.

Methods	GMRES Iterations	CPU Time (s)
NMV-YZβ_ω	359999	2023.721
NMV-YZβ	359602	2044.655
NMV-CAU	360378	2076.091

The Figs. 9a and 9b show the density distributions for the viscous NACA0012 airfoil in the subsonic regime with $M_0 = 0.5$, obtained with the NMV-YZβ_ω and NMV-YZβ methods, respectively. These solutions are similar and are also in good agreement with those reported in [5,15].

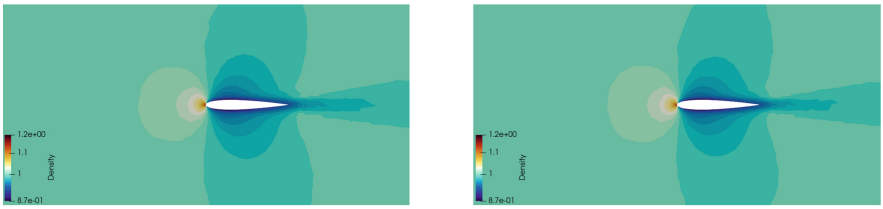

(a) NMV-YZβ_ω. (b) NMV-YZβ.

Fig. 9. Computed density distributions for the subsonic NACA0012 airfoil problem with $M_0 = 0.5$.

Figure 10b shows the temporal evolution of the Euclidean norm of the density residuals for the viscous subsonic NACA0012 airfoil problem using the NMV-YZβ_ω and NMV-YZβ methods. The NMV-YZβ method achieves the smallest residuals, although both methods reach similar orders of magnitude.

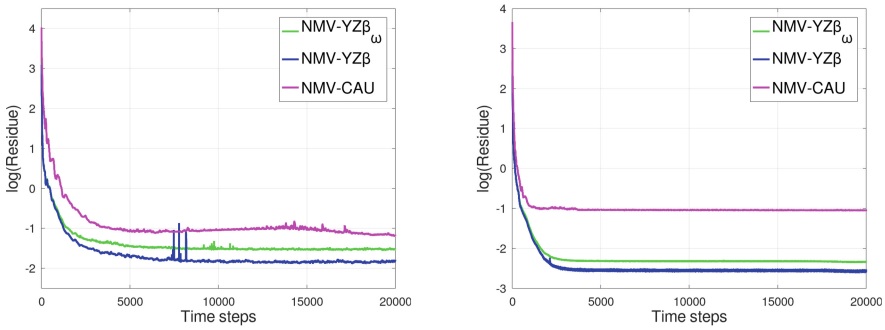

(a) Transonic - $M_0 = 0.85$. (b) Subsonic - $M_0 = 0.5$.

Fig. 10. Temporal evolution of the Euclidean norm of the density residual for the viscous NACA0012 airfoil problem.

The computational performance results for the viscous subsonic flow over the NACA0012 airfoil, using the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods, are presented in Table 3. For this case, the NMV-YZβ_ω and NMV-YZβ methods required the same number of GMRES iterations. The Mach number takes values in the interval $[0, 0.58]$, including values $M < 0.3$, which led the NMV-YZβ_ω method to intensify the action of the parameter $\beta = 1$ in those regions. On the other hand, the NMV-CAU method presented the highest computational cost. However, the CPU times are similar among the analyzed methods, with the NMV-YZβ_ω method being approximately 1.5% faster than the NMV-CAU method.

Table 3. Computational performance for the subsonic NACA0012 airfoil problem.

Methods	GMRES Iterations	CPU Time (s)
NMV-YZβ_ω	359999	2025.696
NMV-YZβ	359999	2045.060
NMV-CAU	360455	2056.188

6 Conclusions

This work presented the multiscale variational method NMV-YZβ_ω for the solution of the compressible Navier–Stokes equations. This method decomposes

the problem of interest into two subproblems, corresponding to the macro and microscales, and introduces artificial viscosity at both discretization scales through the weighting of the YZβ discontinuity-capturing parameter.

The NMV-YZβ_ω, NMV-YZβ and NMV-CAU methods were applied to several benchmark problems governed by the compressible Navier–Stokes equations. The spatial convergence of the methods was investigated using the planar Couette problem, evaluating their behavior in the $L^2(\Omega)$ and $H^1(\Omega)$ norms. The accuracy of the formulations was further assessed through the Carter flat plate problem by comparing their numerical solutions with the reference solutions. In addition, numerical results were obtained for the NACA0012 airfoil problem under different flow regimes. The obtained solutions are consistent with those reported in the literature, confirming the validity of the proposed NMV-YZβ_ω method and its implementation.

Computational performance was also analyzed by comparing the NMV-YZβ_ω, NMV-YZβ, and NMV-CAU methods in terms of GMRES iterations and CPU time. The results indicate that the NMV-YZβ_ω method provides solutions comparable to those of the NMV-YZβ method while maintaining similar computational efficiency. Moreover, both approaches demonstrate superior performance with respect to the NMV-CAU method in terms of solution quality and computational cost.

Future work will focus on extending the proposed formulation to more challenging flow configurations, including three-dimensional problems and higher Reynolds number regimes. In addition, further investigations will consider the application of the NMV-YZβ_ω method to other compressible flow problems and the development of more efficient iterative solution strategies for large-scale simulations.

Acknowledgment. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001 - and by the Fundação de Amparo à Pesquisa e Inovação do Espírito Santo (FAPES), Grant Notice No. 02/2026.

Declaration of Generative AI and AI-assisted technologies in the writing process. During the preparation of this work the authors used ChatGPT/OpenAI in order to improve language and readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

1. Almeida, R.C., Galeão, A.C.: An adaptive Petrov–Galerkin formulation for the compressible Euler and Navier–Stokes equations. Comput. Methods Appl. Mech. Eng. **129**(1–2), 157–176 (1996)
2. Bento, S.S., de Lima, L.M., Sedano, R.Z., Catabriga, L., Santos, I.P.: A nonlinear multiscale viscosity method to solve compressible flow problems. In: Gervasi, O., Murgante, B., Misra, S., Rocha, A.M.A.C., Torre, C., Taniar, D., Apduhan, B.O.,

Stankova, E., Wang, S. (eds.) ICCSA 2016. LNCS, vol. 9786, pp. 3–17. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-42085-1_1
3. Bento, S.S., Barbosa, P.W., Santos, I.P., de Lima, L.M., Catabriga, L.: A nonlinear finite element formulation based on multiscale approach to solve compressible Euler equations. In: Gervasi, O., et al. (eds.) ICCSA 2017. LNCS, vol. 10409, pp. 735–743. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-62407-5_55
4. Bernine, E., Santos, I.P., Catabriga, L.: A nonlinear variational multiscale stabilized method for compressible Navier-Stokes equations. In: Gervasi, O., et al. (eds.) Computational Science and Its Applications – ICCSA 2024 Workshops. LNCS, vol. 14650, pp. 251–267. Springer, Cham (2024)
5. Bernine, E., Bento, S. S., de Lima, L. M., Santos, I. P., Catabriga, L.: Nonlinear stabilized multiscale methodologies for compressible Navier–Stokes equations. SSRN Electron. J. (2025)
6. Bonilla de Toro, J.: Implementation of finite element solvers for the compressible Navier–Stokes equations. PhD thesis, Universitat Politècnica de Catalunya (2015)
7. Carter, J.E.: Numerical solutions of the Navier-Stokes equations for the supersonic laminar flow over a two-dimensional compression corner. NASA Technical Report R-385 (1972)
8. Catabriga, L., Coutinho, A.L.G.A.: Implicit SUPG solution of Euler equations using edge-based data structures. Comput. Methods Appl. Mech. Eng. **191**(32), 3477–3490 (2002)
9. Galeão, A. C., do Carmo, E. G. D.: A consistent approximate upwind Petrov–Galerkin method for convection-dominated problems. Comput. Methods Appl. Mech. Eng. **68**(1), 83–95 (1988)
10. Godunov, S.K., Semendyaev, K.A.: Difference methods for the numerical solution of problems in gas dynamics. Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki **2**(1), 3–14 (1962)
11. Hughes, T.J.R., Tezduyar, T.E.: Finite element methods for first-order hyperbolic systems with particular emphasis on the compressible Euler equations. Comput. Methods Appl. Mech. Eng. **45**(1–3), 217–284 (1984)
12. Hughes, T.J.R.: Multiscale phenomena: green's functions, the Dirichlet-to-Neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods. Comput. Methods Appl. Mech. Eng. **127**(1–4), 387–401 (1995)
13. Koobus, B., Farhat, C.: A variational multiscale method for the large eddy simulation of compressible turbulent flows on unstructured meshes-application to vortex shedding. Comput. Methods Appl. Mech. Eng. **193**(15–16), 1367–1383 (2004)
14. Couette, M.: Études sur le frottement des liquides. Annales de Chimie et de Physique **21**, 433–510 (1890)
15. Parada, S., Codina, R., Baiges, J.: A VMS-based fractional step technique for the compressible Navier-Stokes equations using conservative variables. J. Comput. Phys. **459**, 111137 (2022)
16. Quarteroni, A.: Numerical Models for Differential Problems. Springer, Milan (2009)
17. Tezduyar, T.E.: Determination of the stabilization and shock-capturing parameters in SUPG formulation of compressible flows. In: Proceedings of the European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS (2004)
18. Tezduyar, T.E., Senga, M.: Stabilization and shock-capturing parameters in SUPG formulation of compressible flows. Comput. Methods Appl. Mech. Eng. **195**(13–16), 1621–1632 (2006)
19. Tezduyar, T.E., Senga, M.: SUPG finite element computation of inviscid supersonic flows with YZβ shock-capturing. Comput. Fluids **36**(1), 147–159 (2007)

Dynamics Between Order and Chaos for a Family of Two-Dimensional Cellular Automata

Ana Maria Leitão[1,2,3,4(✉)], Ricardo Severino[2(✉)], and Maria João Alves[1,3,4(✉)]

[1] FMH - Faculty of Human Kinetics, Lisbon University, Lisbon, Portugal
[2] CMAT - Centre of Mathematics, University of Minho, Braga, Portugal
ricardo@math.uminho.pt
[3] INET-md - Institute of Ethnomusicology - Center for Studies in Music and Dance, branch of FMH, University of Lisbon, Lisbon, Portugal
[4] CIPER - Interdisciplinary Centre for the Study of Human Performance, branch of FMH, Lisbon University, Lisbon, Portugal
ana-leitao1@campus.ul.pt, mjalves@edu.ulisboa.pt

Abstract. We study a family of finite two-dimensional Boolean cellular automata with a four-neighbour von Neumann neighbourhood and fixed null boundary conditions. With the goal of identifying local rules that generate dynamical regimes intermediate between order and chaos, we present results that allow us to identify rules whose dynamics are extremely simple. We then present computational results showing that the dynamics of another set of local rules is chaotic. This provides a framework for the search for cellular automata whose behaviour can be classified as lying between order and chaos. We conclude by presenting examples of cellular automata whose dynamics should be classified as between order and chaos. Some of these examples show that Wolfram's original idea of classifying dynamical regimes between order and chaos as complex captures only part of what these very simple dynamical systems can exhibit.

Keywords: Finite cellular automata · 2D cellular automata · complex dynamics

1 Dynamics Between Order and Chaos

The study of elementary cellular automata (ECA) occupies a central role in understanding how complex behavior can emerge from extremely simple deterministic rules. Despite their minimal definition, these systems serve as a fundamental testing ground for investigating the transition between order and chaos – a hallmark of complex dynamics in discrete media. SpaceâĂŞtime diagrams generated by ECA vividly demonstrate that structured and disordered patterns can coexist within the same rule, and that transitions between regularity and irregularity naturally arise from local interactions. Such observations reveal that

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 73–83, 2027.
https://doi.org/10.1007/978-3-032-30518-3_5

even the simplest algorithmic prescriptions are capable of producing emergent structures whose richness far exceeds their computational simplicity.

The search for dynamics between order and chaos need not remain confined to this limited class of elementary automata. The small size of the rule set – comprising only 256 local transition rules – may partially explain the absence of other intermediate dynamical forms in Wolfram's original classification. Moreover, the frequent adoption of periodic boundary conditions further constrains the range of observable behaviors, occasionally suppressing nontrivial regimes. As demonstrated in [2], certain one-dimensional cellular automata display transitions between order and disorder that differ substantially from those reported by Wolfram.

Building on this insight, subsequent studies by various researchers have extended the analysis to two-dimensional lattices and to systems with slightly expanded neighborhoods, see [4], revealing an even broader spectrum of emergent phenomena. This collective body of evidence supports the hypothesis that dimensionality plays a critical role in the development of complexity, suggesting that rich dynamical behavior is not an exclusive feature of simple one-dimensional rules but expands with the spatial structure underlying discrete systems.

In what follows, we turn our attention to the family of two-dimensional peripheral cellular automata with fixed null boundary conditions. By examining their dynamical evolution, we aim to identify regimes that may be interpreted as lying between order and chaos, and to determine the structural features that give rise to such intermediate behaviors. Through this approach, we seek to contribute to a broader understanding of how spatial constraints and boundary conditions influence the emergence of complexity in discrete dynamical systems.

2 A Family of Two-Dimensional Finite Four-Neighbor Cellular Automata

Two-dimensional finite cellular automata (FCA) are discrete computational models defined on a grid of cells, each capable of adopting one of a finite number of possible states. At every time step, the state of each cell is updated synchronously according to a prescribed local rule that depends on the configuration of its neighboring cells. In rectangular lattices, two standard neighborhood schemes have been extensively studied: the von Neumann neighborhood, consisting of the four orthogonally adjacent cells, and the Moore neighborhood, which includes the eight surrounding cells. Even when limiting the analysis to the simplest of these choices, the number of admissible local transition rules is so large that an exhaustive exploration becomes computationally unfeasible, highlighting the need to restrict attention to a simpler family amenable to systematic study.

We consider finite square $n \times n$ Boolean synchronous cellular automata with a four neighbor peripheral neighborhood, that is, a neighborhood of von Neumann type with the center cell not considered. If we denote by

$$\mathsf{C}(t) = c_{ij}(t), \qquad i, j = 1, \cdots, n,$$

the system state configuration at time t, then the state of the cell c_{ij} at time $t+1$ is given by the value that a Boolean function ϕ of four variables – the local update rule – takes on the 4-tuple consisting of the states of the up, left, right, and down neighbors of that cell at the previous time t:

$$c_{ij}(t+1) \;=\; \phi\left(c_{i-1,j}(t),\, c_{i,j+1}(t),\, c_{i+1,j}(t),\, c_{i,j-1}(t)\right), \quad i,j = 1,\ldots,n.$$

Since we have a finite system, we know we have to choose some boundary conditions: in this case we choose fix null boundary conditions. Hence, we have

$$\begin{aligned} c_{0k}(t) \;&=\; c_{n+1,k}(t) \;=\; 0 \qquad k = 1,\ldots,n \\ c_{k0}(t) \;&=\; c_{k,n+1}(t) \;=\; 0 \end{aligned}$$

For convenience, we will follow the usual procedure of associating an integer code number $\mathcal{N}_\phi$ with each automaton ϕ. The numbering scheme considered here is the one used in [1] and is defined as follows: to each neighborhood state $(\mathsf{u},\, \mathsf{l},\, \mathsf{r},\, \mathsf{d})$, with $\mathsf{u},\mathsf{l},\mathsf{r},\mathsf{d} \in \{0,1\}$, we associate the integer number whose binary representation is $(\mathsf{u}\,\mathsf{l}\,\mathsf{r}\,\mathsf{d})_2$. Hence, we have 16 different neighborhood states $\mathsf{n}_k;\ k = 0,\ldots,15$, ordered according to these numbers. The graphical representation of such distinct neighborhood states can be given as

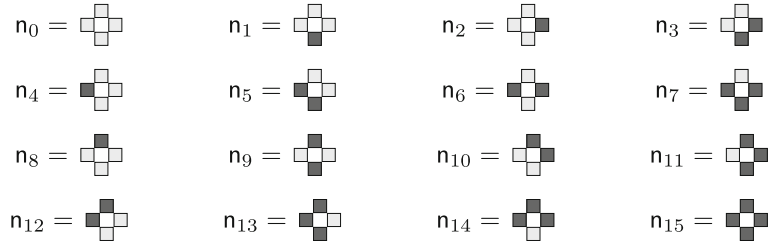

The integer code $\mathcal{N}_\phi$, corresponding to the rule ϕ, is then given by

$$\mathcal{N}_\phi = \sum_{k=0}^{15} \phi(\mathsf{n}_k)\, 2^k.$$

In what follows, we will indistinctly refer to a cellular automaton either by the associated boolean function ϕ or the integer code $\mathcal{N}_\phi$. Let us illustrate with an example the concepts of Wolfram integer code and local transition rule.

Example 1. Consider the transition rule $\mathcal{N}_\phi = 38\,617 = (1001011011010111)_2$, whose graphical representation can be given as For a system comprising 5×5 elements, we computed the system configurations at time instances $T = 1,\ldots,24$, starting from a randomly chosen configuration, as shown in Table 1. Upon examining the set of configurations, we can readily identify the dynamics attractor, a cycle with a period of 8, emerging after a transient of 15 time steps.

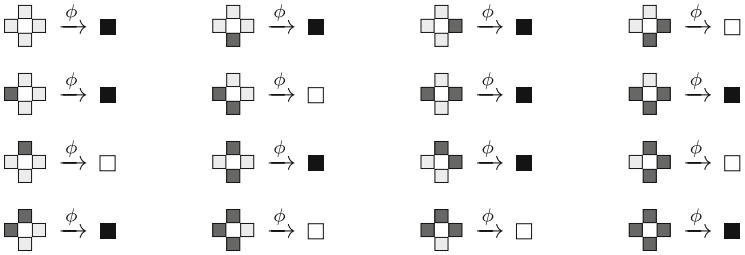

Table 1. The dynamics observed for transition rule $\mathcal{N}_\phi = 38\,617$, starting from a randomly chosen initial configuration, for $T = 0$ to 24. Upon careful observation, we can easily discern the emergence of a final cycle with a period of 8, occurring after a transient of 15 time steps.

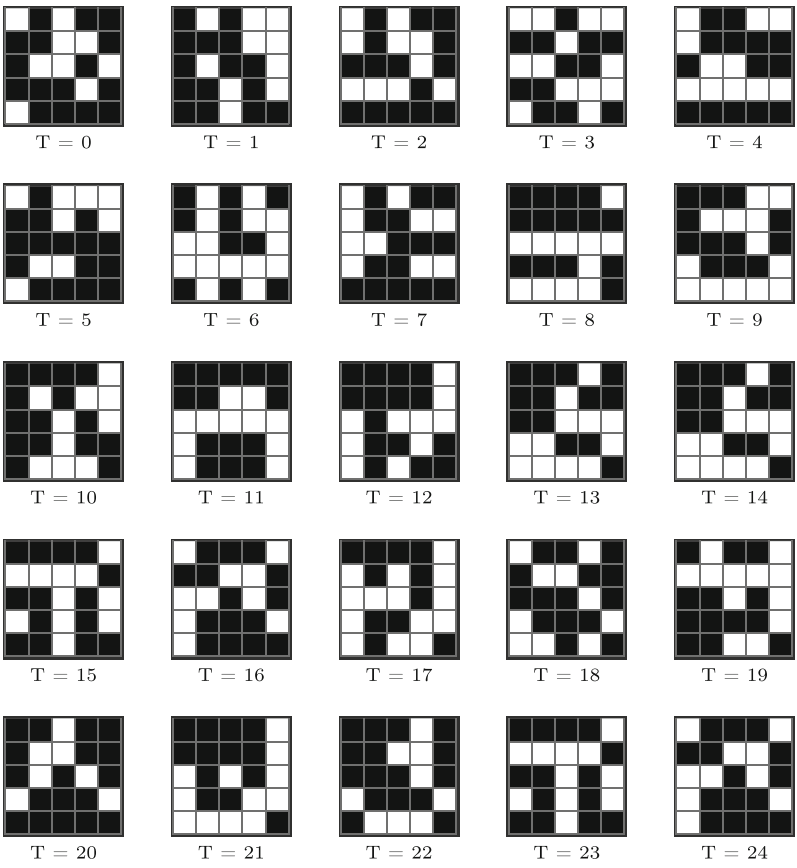

Since our aim is to study exhaustively the dominant dynamics of the family of cellular automata described above, it is essential to reduce their number. Taking

advantage of the results presented in [3], and taking into account the differences in the choice of boundary conditions – the conjugacy symmetry is no longer valid, since we are considering systems with fix null boundary conditions – we can establish that the following result holds.

Proposition 1. *There are 9 616 dynamically nonequivalent square Boolean peripheral cellular automata.*

We believe that this number may be sufficiently small to allow a detailed study of the dynamics of this type of two-dimensional cellular automata, in a manner similar to what has been done for the one-dimensional elementary cellular automata.

Next, we will identify the cellular automata for which a homogeneous configuration is dynamically relevant.

3 Homogeneous Cellular Automata

Wolfram's classification of one-dimensional elementary cellular automata, based on an analysis of the patterns produced by their temporal evolution, places in Class I those automata whose behavior is characterized by convergence to homogeneous final states [1]. The special configurations referred to by Wolfram are precisely the homogeneous ones, that is, those in which all cells of the system are in either the zero state or the one state. Next, we identify the subset of local transition rules whose dynamics are characterized by convergence to homogeneous final configurations.

Proposition 2. *Consider the cellular automaton*

$$\mathcal{N}_\phi = (d_{15}\, d_{14}\, d_{13}\, d_{12}\, d_{11}\, d_{10}\, d_9\, d_8\, d_7\, d_6\, d_5\, d_4\, d_3\, d_2\, d_1\, d_0)_2.$$

If at least one of the following conditions holds:

$d_0 = d_1 = d_2 = d_3 = 0$;
$d_0 = d_1 = d_4 = d_5 = 0$;
$d_0 = d_4 = d_8 = d_{12} = 0$;
$d_0 = d_2 = d_8 = d_{10} = 0$,

then its dynamics will have as fixed point the homogeneous null configuration.

Proof. Consider a system with $n \times n$ elements and let $\mathsf{C}^T = (c_{ij}^T)$ denote the configuration of the system at time T. Each of the conditions listed previously corresponds to evaluating the system's dynamics from one of its four corners. Here we describe the reasoning for the upper-left corner under the first condition; the remaining three cases follow analogously.

Assume that $d_0 = d_1 = d_2 = d_3 = 0$. This ensures that, no matter the initial configuration chosen, $c_{11}^T = 0$, for $T > 0$. Likewise, we obtain $c_{12}^T = c_{21}^T = 0$, for $T > 1$. Iterating the argument yields $c_{1n}^T = c_{n1}^T = 0$, for $T > n-1$. Consequently, the same reasoning allows us to conclude that, after a sufficiently large time T, $\mathsf{C}^T = (c_{ij}^T)$ will be the homogeneous zero configuration, as required. □

The next two results characterise those cellular automata whose dynamics eventually settle either into a fixed point consisting of the homogeneous configuration 1, or into a 2-cycle consisting of both homogeneous configurations.

Lemma 1. *The cellular automaton* $\mathcal{N}_\phi = 65\,535 = (1111111111111111)_2$ *is the only local transition rule whose dynamics have as fixed point the homogeneous one configuration.*

Lemma 2. *There are no cellular automata whose dynamics have as attractor the homogeneous 2-cycle.*

From these results, we may conclude that there are 1 706 dynamically nonequivalent square Boolean peripheral homogeneous cellular automata with fixed null boundary conditions whose dynamics end in a fixed point consisting of a homogeneous configuration.

As the preceding results indicate, the adoption of fixed null boundary conditions for cellular automata yields a number of local transition rules with trivial homogeneous dynamics that is greater than one might initially expect. This observation naturally raises the question of whether there exist cellular automata whose dynamics genuinely display chaotic behavior.

In the following, we examine in detail the typical dynamics of a cellular automaton and show that they exhibit features commonly associated with chaos.

4 An Example of a Chaotic Cellular Automaton

We have already discussed Wolfram's classification of cellular automata in the context of trivial homogeneous Class I dynamics. We now show that there exist square Boolean peripheral cellular automata with fixed null boundary conditions whose dynamics exhibit characteristics akin to those associated with chaotic Class III cellular automata.

Table 2. Configurations of the system during its evolution from a randomly chosen initial configuration on a 30×30 lattice for the cellular automaton $\mathcal{N}_\phi = 38\,617$. The sequence illustrates the temporal development of the dynamics and suggests the absence of any clear tendency toward regularity.



T = 10 000 000 T = 20 000 000 T = 30 000 000

Consider the cellular automaton $\mathcal{N}_\phi = 38\,617$ introduced in Example 1. We argue that its dynamics are associated with chaotic behavior. To substantiate this claim, we first present several configurations of the system at different time steps, showing that no tendency toward regularity can be discerned. In other words, one does not observe a growth of local patterns in the images corresponding to the system's configurations at large times (Table 2).
Another way to identify chaotic dynamics is to study how the period of their attractors varies as the number of elements in the system increases. In what follows, we present the results of a statistical analysis of the period of the final cycle of the dynamics, obtained from 1 000 randomly selected initial configurations, for systems of varying sizes. The results are summarized in the following tables.

from 1 000 initial 6×6 configurations

period	n.configs	period	n.configs	period	n.configs
400	500	3 600	12	23 600	415
472	68	4 248	5		

from 1 000 initial 7×7 configurations

period	n.configs	period	n.configs	period	n.configs
96	134	11 208	11	68 182	4
894	41	14 304	145	98 070	9
934	76	28 954	1	119 428	1
1 636	17	31 290	10	171 780	1
3 360	12	39 264	97	417 498	144
3 576	8	44 832	170	731 292	50
7 008	2	65 262	4	764 012	62
9 816	1				

from 1 000 initial 8×8 configurations

period	n.configs	period	n.configs	period	n.configs
8 496	366	1 231 920	23	6 380 496	453
13 518	141	1 960 110	16		

from 1 000 initial 9×9 configurations

period	n.configs	period	n.configs	period	n.configs
522	83	64 214	155	16 759 854	243
58 240	74	15 200 640	167	1 869 911 680	269
60 428	1	15 771 708	7	9 700 808 980	2

If we focus on the growth of the largest attractor periods as a function of the number of elements in the system, we claim that the data obtained allow us to conclude that the cellular automaton $\mathcal{N}_\phi = 38\,617$ belongs to Wolfram's Class III.

As a final remark in this section, we note that the data above reveal something notably different from what is known about Class III elementary cellular automata: some dynamics have small attractor periods. In other words, this cellular automaton exhibits a kind of coexistence of dynamics. This is more apparent in the data obtained for systems with 9×9 elements, but it is also evident in the other cases. The nature and implications of this coexistence of dynamics merit closer examination in future work.

5 Ordered and Chaotic Cellular Automata

Following Wolfram's approach, we now aim to identify the cellular automata belonging to Class II – those whose typical dynamics display very short transients and small attractor periods. To this end, we computed the evolution of 1 000 random 30×30 initial configurations and selected the cellular automata whose dynamics reached a final state within 100 time steps. According to this criterion, 1 761 cellular automata were classified as belonging to Class II. The Wolfram integer code of those cellular automata as well as the homogeneous cellular automata belonging to Class I, will be provided in an appendix file.

Having identified both Class I and Class II cellular automata, we next focused on detecting those belonging to Class III, characterized by chaotic dynamics and long-lived, irregular transients. For this purpose, we computed the evolution of 1 000 random 8×8 initial configurations and selected the cellular automata whose dynamics never reached a final state within 1 000 time steps. According to this criterion, 321 cellular automata were classified as belonging to Class III. Similarly, an appendix will include the Wolfram integer codes corresponding to these cellular automata.

We acknowledge that the criterion adopted in our search for Class III cellular automata may seem rather stringent, particularly in view of our observation that some cellular automata exhibit a coexistence of dynamical behaviours. Given the analysis presented above, it is therefore unsurprising that cellular automaton 28 617 fails to satisfy this criterion. Thus, although this cellular automaton was used as an example of chaotic behaviour, it will not be listed among the chaotic cellular automata of this family. Nevertheless, we deem it appropriate to maintain this more restrictive criterion until the mechanisms underlying such coexistence are better understood.

In conclusion, we say that there are 7 114 dynamically nonequivalent square Boolean peripheral cellular automata with fixed null boundary conditions that may exhibit dynamics between order and chaos. The Wolfram integer code of those cellular automata will be provided in an appendix file.

Our proposal to study this family of square cellular automata with fixed null boundary conditions, aimed at identifying dynamics that lie between order and chaos, would be meaningless if their existence were not already established. We now proceed to present an example of a cellular automaton with these characteristics, whose dynamics are clearly positioned between order and disorder.

6 Cellular Automata Between Order and Chaos

The search for cellular automata whose typical dynamics exhibit behaviors clearly lying between order and disorder is extremely difficult, since a priori there is no criterion comparable to those used to identify local transition rules belonging either to Class II or to Class III. In this sense, it is necessary to carry out an exhaustive investigation of the typical dynamics within the set of cellular automata identified as having the potential to exhibit this type of behavior.

Considering that what we are proposing is only the beginning of a study that requires more time – it is clearly a work in progress – we would like to introduce a class of cellular automata whose typical dynamics display unique characteristics that set them apart from both Classes II and III.

Consider the cellular automaton $\mathcal{N}_\phi = 1\,241 = (0000010011011001)_2$. We argue that its typical dynamics simultaneously exhibit features of both order and disorder. Let us consider a system of 30×30 elements with an initial configuration chosen at random. After 2 000 time steps, the dynamics appear as an information flow from left to right, in which small geometric structures are translated nearly perfectly from the left side of the system to the right, as illustrated by the configurations shown in Table 3.

The dynamic behavior illustrated by the configurations presented above is noteworthy in two respects. First, it involves an apparent information flow in a system with fixed null boundary conditions, which is highly counterintuitive, since the boundary states are fixed and therefore cannot act as sources of incoming information. Second – and equally importantly – the translation of small geometric structures is evident yet imperfect, ultimately leading to an attractor with an exceptionally long period of 28 764. Furthermore, since, in this particular case, the transient phase has length 1 361, it follows that all configurations shown above already belong to the final cycle.

In this section, we have proposed the identification of cellular automata with dynamics between order and chaos through the empirical observation of selected configurations. Measures aimed at detecting Wolfram's Class IV cellular automata have been introduced in the past, but only after the characteristic features shared by such dynamics had been recognized. Similarly, we believe that a parameter measuring the notion of flow may be introduced for the cellular automata considered here. Nevertheless, such a parameter is likely to be meaningful only within this particular class of systems, whose dynamics lie between order and chaos. For this reason, in our opinion, any attempt to classify these systems should first rely on a careful empirical analysis of their temporal evolution, from which the common dynamical features of this intermediate class may then be extracted.

7 Conclusions

In this work, we investigated a class of finite two-dimensional Boolean cellular automata defined by a four-neighbour von Neumann neighbourhood and fixed

Table 3. Configurations of the system during its evolution from a randomly chosen initial configuration on a 30×30 lattice, for the cellular automaton 1 241. The sequence illustrates the temporal development of the dynamics and suggests that, after some instants of time, there is a flow of information from the left to the right.

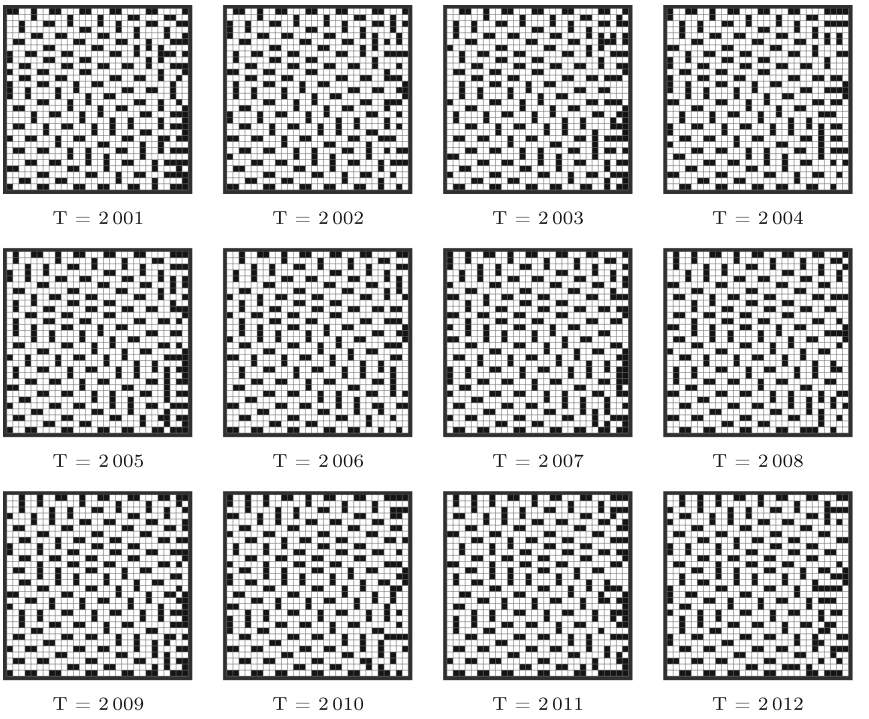

null boundary conditions. Although this setting gives rise to a large number of possible local rules, namely 65 536, the finite character of the system makes a direct computational exploration feasible. This allowed us to perform a systematic characterization of the dynamical behaviours generated within this family, rather than focusing only on isolated examples.

Our exploration revealed a broad range of behaviours. Some local rules give rise to remarkably simple dynamics, rapidly converging to fixed points or short-period attractors, whereas others produce evolutions that are clearly chaotic in character. Between these two extremes, we identified rules whose dynamics display more subtle and structured forms of evolution, suggesting the presence of intermediate regimes between order and chaos. In this sense, the present study provides not only a catalogue of representative behaviours, but also a practical route for locating and analysing cellular automata whose dynamics do not fit neatly into the usual ordered/chaotic dichotomy.

The examples discussed throughout the paper show that the notion of complexity underlying Wolfram's classification captures only part of the phenomenology exhibited by these systems. In particular, the two-dimensional character of the configurations plays a decisive role in the emergence of spatial struc-

tures, apparent information flows, coexistence of dynamical behaviours, and long-period attractors.

One of the main contributions of this work is therefore to show that finite two-dimensional cellular automata with simple local rules and fixed null boundary conditions already provide a rich enough framework for studying the transition between order and chaos. The results presented here indicate that such transitions can be explored empirically through the temporal evolution of configurations, and that this empirical analysis may serve as a first step towards the definition of more refined quantitative measures, possibly tailored to specific families of cellular automata.

Overall, our results reinforce the view that finite two-dimensional cellular automata constitute an effective and insightful setting for studying discrete dynamical systems. At the same time, they suggest that the classification of their behaviours requires criteria sensitive not only to convergence or apparent randomness, but also to spatial organization, propagation phenomena, and the coexistence of distinct dynamical regimes.

Acknowledgements. Research at CMAT was partially financed by Portuguese funds through FCT - Fundação para a Ciência e a Tecnologia within the project UID/00013/2025.

Research at INET - md was supported by the FCT - Fundação para a Ciência e a Tecnologia under grant number 2021.07216.BD.

Ana Maria Leitão was supported by a 2021 PhD scholarship, at the Faculty of Human Kinetics, University of Lisbon, entitled "Generative Dance as a Self-Organized Dynamic System: A Study of Choreographic Emergence and the Phenomenon of Collectivity," funded by the FCT - Fundação para a Ciência e a Tecnologia under grant number 2021.07216.BD

(https://doi.org/10.54499/2021.07216.BD)

The authors are very grateful to the referees, whose careful reading and valuable comments led to improvements in the exposition of the paper.

References

1. Wolfram, S.: Universality and complexity in cellular automata. Physica D **10**, 1–35 (1984)
2. Aguiar, I., Severino, R.: Two elementary cellular automata with a new kind of dynamic. Complex Syst. **24**(2) (2015)
3. Freitas, J.A., Severino, R.: 2D elementary cellular automata with four neighbors. Int. J. Bifurcat. Chaos **23**(4) (2013)
4. Packard, N., Wolfram, S.: Two-dimensional cellular automata. J. Stat. Phys. **38**, 901–946 (1985)
5. Wolfram, S.: A New Kind of Science. Wolfram Media, (2002)
6. Ilachinski, A.: Cellular automata: a discrete universe. World Scientific (2001)
7. Gács, P.: Reliable cellular automata with self-organization. J. Stat. Phys. **103**(1–2), 45–267 (2001)
8. Cook, M.: Universality and computation in CA. Complex Syst. **15**(1), 1–40 (2004)
9. Frisch, U., Hasslacher, B., Pomeau, Y.: Lattice-gas automata for the Navier-Stokes equation. Phys. Rev. Lett. **56**, 1505–1508 (1986)

Appell-Vietoris Polynomials: Structural Properties and Zero Bounds

Isabel Cação[2(✉)], M. Irene Falcão[1], Helmuth R. Malonek[2], and Graça Tomaz[3,4]

[1] Centro de Matemática, Universidade do Minho, Braga, Portugal
mif@math.uminho.pt
[2] CIDMA, Universidade de Aveiro, Aveiro, Portugal
{isabel.cacao,hrmalon}@ua.pt
[3] CISE - Electromechatronic Systems Research Centre, Universidade da Beira Interior, Covilhã, Portugal
gtomaz@ipg.pt
[4] Polytechnic of Guarda, Guarda, Portugal

Abstract. The paper focuses on a sequence of real polynomials forming an Appell sequence whose coefficients are built from the well-known Vietoris sequence. A three-term recurrence relation and the corresponding Sturm—Liouville equation are established as well as hypergeometric and Jacobi representations. The distribution of real zeros is characterized via the Hilbert—Klein theory for Jacobi polynomials with arbitrary parameters. Finally, quantitative bounds for all complex zeros are provided by applying variants of the Eneström—Kakeya theorem.

Keywords: Appell polynomials · Vietoris sequence · Jacobi polynomials · Zero bounds · Eneström-Kakeya bounds

1 Introduction

> *This Table has truly admirable and extraordinary properties. Beyond what I have already shown of the mystery of combinations hidden within it, it is known to those skilled in the deeper parts of geometry that the most important secrets of the rest of mathematics lie concealed within it.*
>
> Jacob Bernoulli in *Ars Conjectandi* (1713)

The table that Bernoulli had in mind here is, of course, the table of binomial coefficients, as we call its entries today. Generalizations of Pascal's triangle of various kinds and arising from different areas of theoretical or applied mathematics appear frequently. Nevertheless, further generalizations continue to emerge, since their diversity seems almost inexhaustible and depends largely on the perspective of the researcher.

Perhaps unexpectedly, our perspective on Bernoulli's table was motivated by problems in hypercomplex analysis [8,9]. In fact, while searching for an adequate

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 84–98, 2027.
https://doi.org/10.1007/978-3-032-30518-3_6

definition of Appell polynomials in several hypercomplex variables with respect to a hypercomplex derivative operator introduced in [11], the dominant role of a monotonically non-increasing sequence of rational numbers became apparent. These numbers are

$$c_k := \frac{1}{2^k}\binom{k}{\lfloor k/2 \rfloor}, \qquad k = 0, 1, 2, \ldots. \tag{1}$$

They form a strictly positive pairwise occurring sequence that first appeared in Vietoris' work on positive trigonometric sums [21]. Since then, the sequence – now known as the *Vietoris sequence* – has been shown to satisfy interesting combinatorial identities and to play a structural role in hypercomplex analysis [2,3,8]. Combining it with the well-known Appell property [1], led to the introduction of the *Appell-Vietoris polynomials* in [5],

$$\mathcal{V}_n(x) = \sum_{k=0}^{n} \binom{n}{k} c_{n-k}\, x^k, \qquad n = 0, 1, 2, \ldots, \tag{2}$$

where qualitative features of these polynomials were highlighted, namely:

(i) for each fixed degree n, the coefficient row $\{\binom{n}{k} c_{n-k}\}_{k=0}^{n}$ is unimodal and log-concave (see [19,22] for definitions); and
(ii) based on inherited structural properties and numerical evidence, we observe (without full proof) that for even n the polynomials are strictly positive on $\mathbb{R}$ (hence have no real zeros), while for odd n they exhibit exactly one sign change with the unique real zero in $(-1, 0)$.

These observations motivate a deeper structural study linking the Appell nature of $\mathcal{V}_n(x)$ to classical special functions and to precise zero distribution results. The paper is organized as follows: Sect. 2 recalls the Vietoris numbers and the definition of the Appell-Vietoris polynomials, including the coefficient triangle and its unimodality/log-concavity. Section 3 establishes the three-term recurrence and the Sturm–Liouville equation, and derives hypergeometric and Jacobi representations together with Rodrigues formula. Using the Hilbert–Klein theory for Jacobi polynomials with arbitrary parameters, Sect. 4 determines the location of the real zeros, and Sect. 5 provides quantitative bounds for all complex zeros based on variants of the Eneström–Kakeya Theorem.

2 Appell-Vietoris Polynomials Revisited

In this section we recall the main definitions and structural properties of the Appell-Vietoris polynomials that will be used throughout the paper.

Let $(a_k(n))_{k\geq 0}$ denote the triangular array given by

$$a_k(n) = \binom{n}{k} c_{n-k}, \;\; n = 0, 1, \ldots; \;\; k = 0, 1, \ldots n, \tag{3}$$

where c_k are the Vietoris numbers defined in (1). Then the Appell-Vietoris polynomials (2) can be rewritten as

$$\mathcal{V}_n(x) = \sum_{k=0}^{n} a_k(n) x^k, \ n = 0, 1, 2 \ldots.$$

By construction, the sequence $(\mathcal{V}_n(x))_{n\geq 0}$ forms an Appell family and therefore satisfies the relation

$$\mathcal{V}_n'(x) = n\mathcal{V}_{n-1}(x), \qquad n \geq 1. \tag{4}$$

Since the polynomials $\mathcal{V}_n(x)$ are defined through the coefficients $a_k(n)$, it is also useful to recall some of their structural properties. The coefficients $a_k(n)$ are positive rationals and form a triangular array whose rows inherit several properties from the Vietoris sequence; see [4,5] for further details.

Table 1 displays the first rows of the triangle (3) along with the corresponding Appell-Vietoris polynomials $\mathcal{V}_n(x), \ n = 0, \ldots, 6$.

Table 1. Triangle $a_k(n)$: first rows and corresponding Appell-Vietoris polynomials

n	$a_k(n)$	$\mathcal{V}_n(x)$
0	1	1
1	$\frac{1}{2} \quad 1$	$\frac{1}{2} + \mathrm{x}$
2	$\frac{1}{2} \quad 1 \quad 1$	$\frac{1}{2} + \mathrm{x} + x^2$
3	$\frac{3}{8} \quad \frac{3}{2} \quad \frac{3}{2} \quad 1$	$\frac{3}{8} + \frac{3}{2}\,\mathrm{x} + \frac{3}{2}x^2 + x^3$
4	$\frac{3}{8} \quad \frac{3}{2} \quad 3 \quad 2 \quad 1$	$\frac{3}{8} + \frac{3}{2}x + 3x^2 + 2x^3 + x^4$
5	$\frac{5}{16} \quad \frac{15}{8} \quad \frac{15}{4} \quad 5 \quad \frac{5}{2} \quad 1$	$\frac{5}{16} + \frac{15}{8}\,\mathrm{x} + \frac{15}{4}x^2 + 5x^3 + \frac{5x^4}{2} + x^5$
6	$\frac{5}{16} \quad \frac{15}{8} \quad \frac{45}{8} \quad \frac{15}{2} \quad \frac{15}{2} \quad 3 \quad 1$	$\frac{5}{16} + \frac{15}{8}\,\mathrm{x} + \frac{45}{8}x^2 + \frac{15}{2}x^3 + \frac{15}{2}x^4 + 3x^5 + x^6$

Furthermore, for each fixed $n \geq 1$, the finite sequence

$$a_0(n), a_1(n), \ldots, a_n(n)$$

is both log-concave and unimodal. More precisely, it satisfies the log-concavity property

$$(a_k(n))^2 \geq a_{k-1}(n) a_{k+1}(n), \ k = 1, \ldots, n-1,$$

and the unimodality condition

$$a_0(n) \leq a_1(n) \leq \cdots \leq a_{s-1}(n) \leq a_s(n) \geq a_{s+1}(n) \geq \cdots \geq a_n(n), \tag{5}$$

where $s = \lfloor \frac{n+1}{2} \rfloor$ (cf. [5, Th. 2]).

These properties will play a significant role in Sect. 5, where we study bounds for the zeros of the Appell-Vietoris polynomials.

3 Analytic Framework: Recurrence, Sturm–Liouville Equation, and Jacobi Polynomials

Polynomial sequences often satisfy recurrence relations with polynomial coefficients. They are important tools to derive several properties of the considered polynomial sequence and provide an easy way to construct explicitly the sequence.

Theorem 1. *The Appell–Vietoris polynomials satisfy, for $n \geq 2$, the recurrence relation*

$$(n+1)\mathcal{V}_n(x) = (1+2nx)\mathcal{V}_{n-1}(x) + (n-1)(1-x^2)\mathcal{V}_{n-2}(x), \tag{6}$$

with initial conditions $\mathcal{V}_0(x) := 1$ and $\mathcal{V}_1(x) := x + \frac{1}{2}$.

Proof. Define

$$\mathcal{W}_n(x) := (1+2nx)\mathcal{V}_{n-1}(x) + (n-1)(1-x^2)\mathcal{V}_{n-2}(x).$$

We show that $\mathcal{W}_n(x) = (n+1)\mathcal{V}_n(x)$ by comparing coefficients. Write

$$\mathcal{W}_n(x) = \sum_{k=0}^{n} A_{k,n} x^k.$$

Using the explicit representation of $\mathcal{V}_m(x)$ (and the convention $\binom{m}{r} = 0$ for $r < 0$ or $r > m$, and $c_{-j} = 0$, $j \geq 1$), a direct computation yields

$$A_{k,n} = (2n+1-k)\tbinom{n-1}{k-1}c_{n-k} + \tbinom{n-1}{k}c_{n-1-k} + (n-1)\tbinom{n-2}{k}c_{n-2-k}.$$

Now use the recurrence (see [4])

$$(j+2)c_{j+1} = c_j + jc_{j-1}, \qquad j \geq 1, \qquad c_0 = 1,\ c_1 = \tfrac{1}{2}, \tag{7}$$

with $j = n-1-k$, i.e.

$$c_{n-1-k} = (n+1-k)c_{n-k} - (n-1-k)c_{n-2-k}.$$

Substituting into $A_{k,n}$ and grouping the terms in c_{n-k} and c_{n-k-2} gives

$$A_{k,n} = \Big((2n+1-k)\tbinom{n-1}{k-1} + (n+1-k)\tbinom{n-1}{k}\Big)c_{n-k} + \Big((n-1)\tbinom{n-2}{k} - (n-1-k)\tbinom{n-1}{k}\Big)c_{n-2-k}.$$

Using $\binom{n-1}{k} = \binom{n-2}{k} + \binom{n-2}{k-1}$, the contributions to the coefficient of c_{n-k-2} cancel pairwise, hence this term disappears. Finally, applying the standard binomial identities $\binom{n-1}{k-1} = \frac{k}{n}\binom{n}{k}$, $\binom{n-1}{k} = \frac{n-k}{n}\binom{n}{k}$, $\binom{n-2}{k-2} = \frac{k(k-1)}{n(n-1)}\binom{n}{k}$, one obtains

$$A_{k,n} = (n+1)\binom{n}{k}c_{n-k}.$$

Therefore $\mathcal{W}_n(x) = (n+1)\mathcal{V}_n(x)$, corresponding to (6). □

Special-value evaluations of the Appell—Vietoris polynomials reveal connections with classical number sequences, and these links play a central role in Sect. 4. In particular, evaluating at $x = 0$ recovers the Vietoris numbers:

$$\mathcal{V}_n(0) = c_n, \tag{8}$$

which can be given explicitly by (1) or, equivalently, by the recurrence obtained from Theorem 1,

$$(n+1)c_n = c_{n-1} + (n-1)c_{n-2}, \quad n \geq 2, \qquad c_0 = 1,\ c_1 = \tfrac{1}{2},$$

(see also [6, Th. 6]). The special value (8) illustrates one of several connections between the Appell—Vietoris polynomials and classical number sequences. Additional links of this kind will be derived in the corollary that follows.

Corollary 1. *Let $\mathcal{C}_n$ and R_n denote the Catalan numbers and the Riordan numbers, respectively. Then*

$$\mathcal{V}_n(-1) = (-2)^{-n}\mathcal{C}_n \qquad \text{and} \qquad \mathcal{V}_n\big(-\tfrac{1}{2}\big) = (-2)^{-n}R_n.$$

Proof. Evaluating the recurrence relation (6) at the special points $x = -1$ and $x = -\frac{1}{2}$ and applying the specified normalization conditions yields

1. The sequence $(\mathcal{C}_n)_{n\geq 0}$ satisfies

$$\mathcal{C}_n = \frac{4n-2}{n+1}\,\mathcal{C}_{n-1}, \qquad n \geq 1, \qquad \mathcal{C}_0 = 1,$$

   and therefore coincides with the Catalan numbers (see, e.g., [13, Sect. 5]).
2. The sequence $(R_n)_{n\geq 0}$ satisfies

$$(n+1)R_n = (n-1)\big(2R_{n-1} + 3R_{n-2}\big), \quad n \geq 2, \quad R_0 = 1,\ R_1 = 0, \tag{9}$$

   which reproduces the Riordan numbers (see e.g. [7], Th. 4.1).

□

Another consequence of the recursion (6), together with the Appell property of the polynomials $\mathcal{V}_n(x)$, is the following result.

Theorem 2 (Differential equation). *The Appell-Vietoris polynomial $\mathcal{V}_n(x)$ is a solution of the second order linear homogeneous differential equation*

$$(1-x^2)y''(x) + (1+2nx)y'(x) - n(n+1)y(x) = 0, \tag{10}$$

which can be rewritten, for $x \in (-1,1)$, in the Sturm–Liouville form

$$\frac{d}{dx}\left(p(x)\frac{dy}{dx}\right) + q(x)y(x) = 0, \tag{11}$$

where

$$p(x) = (1-x)^{-\frac{1}{2}-n}(1+x)^{\frac{1}{2}-n} \tag{12}$$

and

$$q(x) = -n(n+1)(1-x)^{-\frac{3}{2}-n}(1+x)^{-\frac{1}{2}-n}. \tag{13}$$

Proof. To prove that $\mathcal{V}_n(x)$ is a solution of (10) it is sufficient to use the recurrence relation (6) together with the Appell property (4). The Sturm–Liouville form can be obtained by using standard techniques: normalise (10) to $y''(x) + P(x)\,y'(x) + Q(x)\,y(x) = 0$ with $P(x) = (1 + 2nx)/(1 - x^2)$ and $Q(x) = -n(n+1)/(1-x^2)$; multiply the resulting equation by the integrating factor $p(x) = \exp\left(\int P(x)\,dx\right) = (1-x)^{-1/2-n}(1+x)^{1/2-n}$, which is strictly positive on $(-1, 1)$ to obtain (11) with (12), (13). □

The next result shows that the Appell—Vietoris polynomials belong to a certain subclass of Jacobi polynomials. Jacobi polynomials can be defined in several ways; for our purposes, here we define Jacobi polynomials $P_n^{(\alpha,\beta)}(x)$ as solutions of the differential equation

$$\left(1 - x^2\right) y'' + (\beta - \alpha - (\alpha + \beta + 2)x)y' + n(n + \alpha + \beta + 1)y = 0, \tag{14}$$

where α, β are arbitrary complex parameters. The cases where $\alpha > -1$ and $\beta > -1$ are the most common in virtue of their classical orthogonality. For these restrictions on the parameters, it is well known that the polynomial solutions of (14) are of the form $\kappa P_n^{(\alpha,\beta)}$, where κ is a constant (cf. [20, Th. 4.2.2]). In the general case of arbitrary parameters, Theorem 4.23.2 in [20] establishes that the differential equation has two linearly independent polynomial solutions only in the following two scenarios:

(1) α, β negative integers, $\alpha \geq -n,\ \beta \geq -n,\ \alpha + \beta \leq -n - 1,\ n \geq 1$;

(2) α, β negative integers, $\alpha < -n,\ \beta < -n,\ n \geq 0$.

The following result shows the relationship between Appell-Vietoris polynomials $\mathcal{V}_n(x)$ and Jacobi polynomials $P_n^{(\alpha,\beta)}(x)$, with parameters α and β as negative half-integers that vary linearly with the degree n. This dependence is illustrated by adopting henceforth the notation (α_n, β_n) for the parameters.

Theorem 3 (Jacobi polynomials). *The Appell-Vietoris polynomials can be expressed in terms of Jacobi polynomials as*

$$\mathcal{V}_n(x) = \frac{(-2)^n}{n+1} P_n^{(\alpha_n,\beta_n)}(x), \tag{15}$$

where

$$\alpha_n = -n - \tfrac{3}{2} \qquad \textit{and} \qquad \beta_n = -n - \tfrac{1}{2}. \tag{16}$$

Proof. Observe that (16) leads to

$$\alpha_n + \beta_n = -2n - 2 \qquad \text{and} \qquad \beta_n - \alpha_n = 1,$$

and (14) reduces to (10). This means that both $\mathcal{V}_n(x)$ and $P_n^{(\alpha_n,\beta_n)}(x)$ are polynomial solutions of (14). Since α_n and β_n are not negative integers, we are not in either of the exceptional cases (1)–(2) above. It follows that (14) admits at

most one polynomial solution (up to a multiplicative constant), and thus there exists a constant κ_n such that

$$\mathcal{V}_n(x) = \kappa_n P_n^{(\alpha_n,\beta_n)}(x).$$

To determine κ_n, we use the values at $x = 1$. From the recurrence relation (6) we obtain

$$\mathcal{V}_n(1) = \frac{(2n+1)!!}{(n+1)!},$$

and from [20, p. 65] we have

$$P_n^{(\alpha_n,\beta_n)}(1) = \binom{n+\alpha_n}{n} = \binom{-\frac{3}{2}}{n} = (-1)^n \frac{(2n+1)!!}{2^n\, n!}.$$

Therefore,

$$\kappa_n = \frac{\mathcal{V}_n(1)}{P_n^{(\alpha_n,\beta_n)}(1)} = (-1)^n \frac{2^n}{n+1},$$

which yields (15). □

The following results are straightforward, given the relationship between Appell-Vietoris polynomials and Jacobi polynomials as outlined in Theorem 3.

Corollary 2 (Rodrigues formula). *For every $n \geq 0$,*

$$\mathcal{V}_n(x) = \frac{1}{(n+1)!}(1-x)^{\frac{3}{2}+n}(1+x)^{\frac{1}{2}+n}\frac{d^n}{dx^n}\left((1-x)^{-\frac{3}{2}}(1+x)^{-\frac{1}{2}}\right).$$

Corollary 3 (Connection to hypergeometric function). *For every $n \geq 0$,*

$$\mathcal{V}_n(x) = \frac{(2n+1)!!}{(n+1)!}\ {}_2F_1\left(-n, -(n+1); -(n+\tfrac{1}{2}); \tfrac{1-x}{2}\right).$$

4 Real Zeros of Appell-Vietoris Polynomials

The coefficients of the Appell-Vietoris polynomials are all strictly positive; therefore, none of their real zeros can be positive. In [5] it was observed that $\mathcal{V}_n(x)$ exhibit specific behavior regarding their real zeros: they have no real zeros when n is even, and exactly one real zero when n is odd.

Relying on Theorem 3, it is enough to study the real zeros of the Jacobi polynomials $P_n^{(\alpha_n,\beta_n)}(x)$, for $\alpha_n = -n - \frac{3}{2}$ and $\beta_n = -n - \frac{1}{2}$. In the case of the classical Jacobi polynomials with parameters $\alpha > -1$ and $\beta > -1$, their orthogonality implies that all the zeros are real and simple and lie in the interval $(-1, 1)$. For more general parameters this statement is no longer valid: the classical orthogonality fails and the zeros can be complex and multiple. However, known results on the number of real zeros of the corresponding Jacobi polynomials, developed by Hilbert, Klein and Stieltjes (mentioned by Szegö in [20], Th.

6.72), can be applied. We start by considering the symbol used by F. Klein in [12]

$$E(u) = \begin{cases} 0, & \text{if } u \leq 0 \\ \lfloor u \rfloor, & \text{if } u > 0, u \text{ is not integer} \\ u-1, & \text{if } u = 1, 2, 3, \ldots \end{cases}.$$

Taking into account that $-1, 1$ and ∞ are not zeros of $\mathcal{V}_n(x)$ and, consequently, of the Jacobi polynomials $P_n^{(\alpha_n,\beta_n)}(x)$, with $\alpha_n = -n - \frac{3}{2}$, $\beta_n = -n - \frac{1}{2}$, the following theorem can be applied to establish the number and location of the real zeros of the Appell-Vietoris polynomials.

Theorem 4. [20, Th. 6.72] *Let α, β be arbitrary real values, and set*

$$X = X(\alpha,\beta) = E\{\tfrac{1}{2}(|2n+\alpha+\beta+1| - |\alpha| - |\beta| + 1)\},$$
$$Y = Y(\alpha,\beta) = E\{\tfrac{1}{2}(-|2n+\alpha+\beta+1| + |\alpha| - |\beta| + 1)\},$$
$$Z = Z(\alpha,\beta) = E\{\tfrac{1}{2}(-|2n+\alpha+\beta+1| - |\alpha| + |\beta| + 1)\}.$$

Denoting by N_1, N_2 and N_3 the number of zeros of $P_n^{(\alpha,\beta)}(x)$ in $(-1,1)$, $(-\infty,-1)$, and $(1,+\infty)$, respectively, these values are given by

$$N_1 = N_1(\alpha,\beta) = \begin{cases} 2\lfloor \frac{X+1}{2} \rfloor, & \text{if } (-1)^n \binom{n+\alpha}{n}\binom{n+\beta}{n} > 0 \\ 2\lfloor \frac{X}{2} \rfloor + 1, & \text{if } (-1)^n \binom{n+\alpha}{n}\binom{n+\beta}{n} < 0 \end{cases},$$
$$N_2 = N_2(\alpha,\beta) = \begin{cases} 2\lfloor \frac{Y+1}{2} \rfloor, & \text{if } \binom{2n+\alpha+\beta}{n}\binom{n+\beta}{n} > 0 \\ 2\lfloor \frac{Y}{2} \rfloor + 1, & \text{if } \binom{2n+\alpha+\beta}{n}\binom{n+\beta}{n} < 0 \end{cases},$$
$$N_3 = N_3(\alpha,\beta) = \begin{cases} 2\lfloor \frac{Z+1}{2} \rfloor, & \text{if } \binom{2n+\alpha+\beta}{n}\binom{n+\alpha}{n} > 0 \\ 2\lfloor \frac{Z}{2} \rfloor + 1, & \text{if } \binom{2n+\alpha+\beta}{n}\binom{n+\alpha}{n} < 0 \end{cases}.$$

We can now state the following theorem concerning the real zeros of the Appell-Vietoris polynomials.

Theorem 5. *The Appell-Vietoris polynomials $\mathcal{V}_n(x)$ have no real zeros when n is even, and exactly one real zero when n is odd, located in the interval $[-\frac{1}{2}, 0)$.*

Proof. In a first stage, due to relation (15), we count the number of real zeros in the interval $(-1,1)$ of the Jacobi polynomials $P_n^{(\alpha_n,\beta_n)}(x)$, with $\alpha_n = -n - \frac{3}{2}$ and $\beta_n = -n - \frac{1}{2}$, applying Theorem 4. For that purpose, we compute the numbers N_1, N_2 and N_3 for this special class of Jacobi polynomials:

1. $X = E(-n) = 0$ and

$$(-1)^n \binom{n+\alpha_n}{n}\binom{n+\beta_n}{n} = (-1)^n \binom{-3/2}{n}\binom{-1/2}{n} = (-1)^n \times \text{positive factor};$$

therefore

$$N_1 = \begin{cases} 2\lfloor \frac{X+1}{2} \rfloor = 0, & \text{if } n \text{ is even} \\ 2\lfloor \frac{X}{2} \rfloor + 1 = 1, & \text{if } n \text{ is odd} \end{cases}.$$

2. $Y = E(1/2) = 0$ and

$$\binom{2n+\alpha_n+\beta_n}{n}\binom{n+\beta_n}{n} = \binom{-2}{n}\binom{-1/2}{n} > 0\,;$$

therefore

$$N_2 = 2\lfloor \tfrac{Y+1}{2} \rfloor = 0\,.$$

3. $Z = E(-1/2) = 0$ and

$$\binom{2n+\alpha_n+\beta_n}{n}\binom{n+\alpha_n}{n} = \binom{-2}{n}\binom{-3/2}{n} > 0\,;$$

therefore

$$N_3 = 2\lfloor \tfrac{Z+1}{2} \rfloor = 0\,.$$

We have just proved that the Appell—Vietoris polynomials $\mathcal{V}_n(x)$ possess no real zeros when n is even, and exactly one real zero when n is odd, which lies in the interval $(-1, 1)$. Moreover, this interval can be further narrowed. Taking into account the connection between $\mathcal{V}_n(-\frac{1}{2})$ and the Riordan numbers established in Corollary 1, together with the nonnegativity of the Riordan sequence (cf. (9)), it follows that $\mathcal{V}_n(-\frac{1}{2}) \leq 0$ for odd n, with equality only when $n = 1$. Since $\mathcal{V}_n(0) = c_n > 0$ (cf. (8)), there is a change of sign, and hence the unique real zero for odd n lies in the interval $[-\frac{1}{2}, 0)$. □

To illustrate Theorem 5, we present in Fig. 1 the graphs of the first Appell-Vietoris polynomials, grouped by the parity of their degrees.

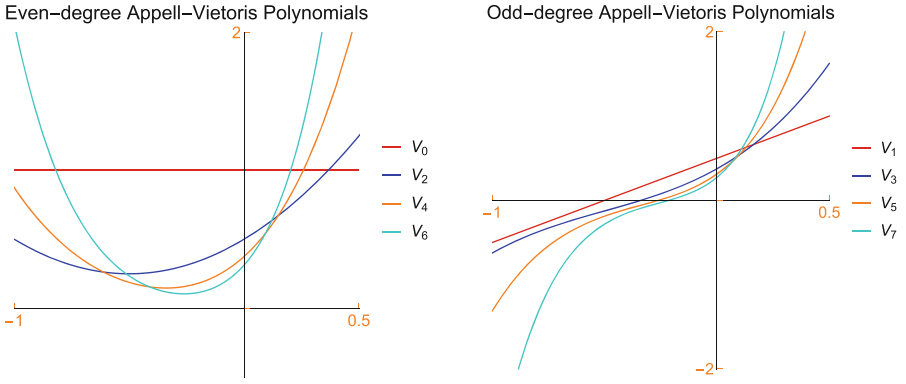

Fig. 1. First Appell-Vietoris Polynomials

5 Eneström-Kakeya Bounds

The literature provides numerous results on the localization of polynomial zeros, including Gershgorin-type bounds, root-radius estimates, and refinements of the classical Cauchy and Fujiwara inequalities [15,17]. In this work, we derive bounds for the zeros of the Appell-Vietoris polynomials $\mathcal{V}_n(x)$ by applying two variants of the Eneström–Kakeya Theorem and some of its generalizations.

We first recall the following result from [10, Th. 4.1], which applies to polynomials with positive coefficients forming an unimodal sequence.

Theorem 6. *Let $p(z) = \sum_{j=0}^{n} a_j z^j$ be a polynomial of degree n with positive real coefficients. If*

$$a_0 \le a_1 \le \cdots \le a_k \ge a_{k+1} \ge \cdots \ge a_n,$$

for some $k \in \{0, \ldots, n\}$, then all zeros of p satisfy

$$|z| \le \frac{2a_k}{a_n} - 1.$$

Since, for each degree n, the coefficient sequences (3) of the Appell-Vietoris polynomials $\mathcal{V}_n(x)$ are positive and unimodal (see (5)), we obtain the following corollary.

Corollary 4. *All zeros of $\mathcal{V}_n(x)$ lie in the disk $|x| \le R^{(n)}$, where*

$$R^{(n)} = 2\binom{n}{s} c_{n-s} - 1, \tag{17}$$

with $s = \lfloor (n+1)/2 \rfloor$ and c_{n-s} denoting the Vietoris numbers defined in (1).

Table 2 reports, for each degree $n = 0, \ldots, 10$, the bound $R^{(n)}$ given by (17) and the quantity

$$z_{\max}^{(n)} = \max_{0 \le i \le n} |z_i^{(n)}|,$$

where $z_i^{(n)}$ are the zeros of $\mathcal{V}_n(x)$. We note that these bounds are very conservative and become increasingly loose as n grows, in contrast with the observed distribution of the zeros.

In order to obtain sharper bounds for the class of polynomials under consideration, we apply now the following form of the Eneström-Kakeya Theorem [10, Th. 1.4].

Theorem 7. *Let $p(z) = \sum_{j=0}^{n} a_j z^j$ be a polynomial of degree n with positive real coefficients. All zeros of p lie within the annulus*

$$\min_{0 \le k \le n-1} \frac{a_k}{a_{k+1}} \le |z| \le \max_{0 \le k \le n-1} \frac{a_k}{a_{k+1}}. \tag{18}$$

Applying this result to the Appell-Vietoris polynomials yields the following corollary.

Table 2. Bounds and zeros of $\mathcal{V}_n(x)$ obtained from Corollary 4.

n	$z_{\max}^{(n)}$	$R^{(n)}$	n	$z_{\max}^{(n)}$	$R^{(n)}$
1	0.5	1.	6	2.05	14.
2	0.71	1.	7	2.37	25.25
3	1.05	2.	8	2.69	51.5
4	1.39	5.	9	3.01	93.5
5	1.72	9.	10	3.32	156.5

Corollary 5. *Let $\mathcal{V}_n(x)$ denote the Appell-Vietoris polynomial of fixed degree n. All zeros of $\mathcal{V}_n(x)$ lie in the annulus*

$$R_1^{(n)} \leq |x| \leq R_2^{(n)},$$

where

$$R_1^{(n)} = \begin{cases} \dfrac{1}{n}, & \text{if } n \text{ is even} \\ \dfrac{1}{n+1}, & \text{if } n \text{ is odd} \end{cases} \qquad \text{and} \qquad R_2^{(n)} = \frac{n}{2}. \tag{19}$$

Proof. Define, for $k = 0, \ldots, n-1$, the finite sequence

$$b_k(n) = \frac{a_k(n)}{a_{k+1}(n)},$$

where $a_k(n)$ are the Appell-Vietoris coefficients (3). Observing (1), the relation between consecutive terms of the Vietoris sequence is given by

$$c_k = \begin{cases} \frac{k}{k+1} c_{k-1}, & \text{if } k \text{ odd} \\ c_{k-1}, & \text{if } k > 1 \text{ even} \\ 1 & \text{if } k = 0 \end{cases},$$

and, consequently

$$b_k(n) = \begin{cases} \dfrac{k+1}{n-k}, & \text{if } n-k \text{ is even} \\ \dfrac{k+1}{n-k+1}, & \text{if } n-k \text{ is odd} \end{cases}.$$

Since

$$\frac{b_{k+1}(n)}{b_k(n)} = \begin{cases} \dfrac{k+2}{k+1} > 1, & \text{if } n-k \text{ is even} \\ \dfrac{(k+2)(n-k+1)}{(k+1)(n-k-1)} > 1, & \text{if } n-k \text{ is odd} \end{cases},$$

it follows at once that $(b_k(n))_k$ is a strictly increasing sequence, and therefore

$$b_0 \leq b_k \leq b_{n-1}.$$

Since $b_0 = \frac{1}{n}$, for n even, $b_0 = \frac{1}{n+1}$ for n odd and $b_{n-1} = \frac{n}{2}$, the conclusion follows directly from (18). □

Table 3 reports, for several values of n, the bounds $R_1^{(n)}$ and $R_2^{(n)}$ given by (19) together with the moduli of the smallest and largest zeros. Figure 2 provides a geometric illustration of these annular bounds, showing the zeros and the corresponding Eneström–Kakeya annuli for degrees $n = 2, \ldots, 6$.

Table 3. Bounds and zeros of $\mathcal{V}_n(x)$ obtained from Corollary 5.

n	$R_1^{(n)}$	$z_{\min}^{(n)}$	$z_{\max}^{(n)}$	$R_2^{(n)}$
1	0.5	0.5	0.5	0.5
2	0.5	0.707	0.707	1.
3	0.25	0.339	1.052	1.5
4	0.25	0.440	1.391	2.
5	0.167	0.261	1.723	2.5
6	0.167	0.328	2.049	3.
7	0.125	0.214	2.371	3.5
8	0.125	0.265	2.689	4.
9	0.1	0.183	3.006	4.5
10	0.1	0.224	3.321	5.
⋮	⋮	⋮	⋮	⋮
20.	0.05	0.129	6.438	10.

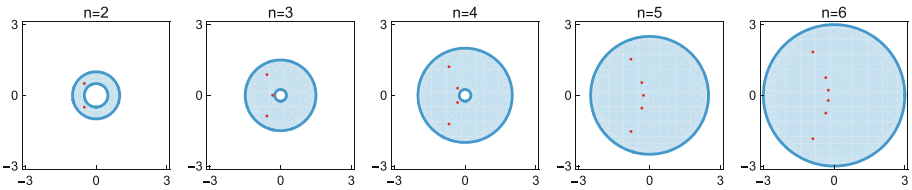

Fig. 2. Bounds for the zeros of the Appell-Vietoris polynomials of degrees $n = 2, \ldots, 6$.

Several results from the literature were tested for the localization of the zeros of the Appell—Vietoris polynomials, with Corollary 5 consistently yielding the sharpest bounds. Nevertheless, Corollary 4 has been retained as a complementary reference: it allows the exploitation of the unimodality property of the coefficients and serves as a baseline for comparison, highlighting the improvement achieved by the sharper classical Eneström–Kakeya estimates. Together, these results provide a comprehensive picture of the zero distribution and illustrate the relative effectiveness of different bounding approaches available in the literature.

6 Concluding Remarks

In this paper, we revisited some fundamental properties of the coefficients of the Appell-Vietoris polynomials that were used to derive structural properties as a three-term recurrence relation and a Sturm–Liouville differential equation for this family. We subsequently proved that the Appell-Vietoris polynomials form a subclass of Jacobi polynomials with negative, non-integer parameters depending linearly on the degree. Polynomials of this type play an important role in Chebyshev spectral methods for boundary-value problems (cf. [18]).

The last two sections were dedicated to the zeros of the Appell-Vietoris polynomials. Their real zeros were counted and localized by combining known results for Jacobi polynomials with specific properties of the Appell-Vietoris family. We also carried out a comparative study of bounds for the full zero set, relying on the Eneström–Kakeya estimates. Other approaches could in principle be pursued, for instance by invoking results on the zero distribution of hypergeometric polynomials or of Jacobi polynomials with varying parameters; see Corollary 3 and Theorem 3, respectively. Concerning the hypergeometric route, we conducted an extensive literature search but, to the best of our knowledge, the specific parameter configuration arising here has not been treated, even after applying standard linear and quadratic transformations. Regarding Jacobi polynomials with varying parameters the zero distribution has been investigated by some authors under conditions on the parameter sequences. In particular, Finkelshtein et al. [16] analyse the zero behaviour for the class of Jacobi polynomials in the parameter range where classical orthogonality fails. They considered $P_n^{(\alpha_n,\beta_n)}(x)$ such that

$$\lim_{n\to\infty} \frac{\alpha_n}{n} = A \qquad \text{and} \qquad \lim_{n\to\infty} \frac{\beta_n}{n} = B, \qquad A, B \in \mathbb{R},$$

and distinguish five cases according to the values of A and B. The case of the Jacobi polynomials considered in the present work, namely $A + B = -2$, corresponds to a transition between these five cases and is not covered by their analysis. As Kuijlaars et al. observed (cf. [14, p. 198]), this case "presents additional difficulties".

The results obtained, together with our numerical experiments, indicate that the zero sets exhibit a specific underlying structure. A more detailed analysis of their zero distribution will be pursued in future work, where other possible approaches will be explored.

Acknowledgments. Research at CIDMA has been supported by FCT (Portuguese Foundation for Science and Technology) under Projects UID/04106/2025 https://doi.org/10.54499/UID/04106/2025 and UID/PRR/04106/2025 https://doi.org/10.54499/UID/PRR/04106/2025. Research at CMAT was partially financed by Portuguese funds through FCT within the Project UID/00013/2025 (https://doi.org/10.54499/UID/00013/2025). Research at CISE has been financed by FCT under Projects UID/4131/2025 (https://doi.org/10.54499/UID/04131/2025), UID/PRR/4131/2025. (https://doi.org/10.54499/UID/PRR/04131/2025), and UID/PRR2/4131/2025. (https://doi.org/10.54499/UID/PRR2/04131/2025).

References

1. Appell, P.: Sur une classe de polynômes. Ann. Sci. École Norm. Sup. **9**(2), 119–144 (1880)
2. Cação, I., Falcão, M.I., Malonek, H.R.: Hypercomplex polynomials, Vietoris' rational numbers and a related integer numbers sequence. Complex Anal. Oper. Theory **11**, 1059–1076 (2017)
3. Cação, I., Falcão, M.I., Malonek, H.R.: On generalized Vietoris' number sequences. Discrete Appl. Math. **269**, 77–85 (2019)
4. Cação, I., Falcão, M.I., Malonek, H.R., Miranda, F., Tomaz, G.: Remarks on the Vietoris sequence and corresponding convolution formulas. In: Gervasi, O., et al. (eds.) ICCSA 2023. LNCS, vol. 14104, pp. 677–692. Springer, Cham (2023)
5. Cação, I., Falcão, M.I., Malonek, H.R., Miranda, F., Tomaz, G.: On Appell-Vietoris polynomials. In: Gervasi, O., et al. (eds.) ICCSA 2024. LNCS, vol. 14815, pp. 302–316. Springer, Cham (2024)
6. Cação, I., Falcão, M.I., Malonek, Tomaz, G.: A Sturm-Liouville equation on the crossroads of continuous and discrete hypercomplex analysis. Math. Meth. Appl. Sci **47**(10), 7962–7987 (2024)
7. Chen, W.Y.C., Deng, E.Y.P., Yang, L.L.M.: Riordan paths and derangements. Discrete Math. **308**(11), 2222–2227 (2008)
8. Falcão, M.I., Malonek, H.R.: Generalized exponentials through Appell sets in $\mathbb{R}^{n+1}$ and Bessel functions. In: AIP Conference Proceedings, vol. 936, pp. 738–741 (2007)
9. Falcão, M.I., Cruz, J.F., Malonek, H.R.: Remarks on the generation of monogenic functions. In: Gürlebeck, K., Könke, C. (eds.) Proceedings of the 17th International Conference on on the Application of Computer Science and Mathematics on Architecture and Civil Engineering (IKM), Weimar, Germany (2006)
10. Gardner, R.B., Govil, N.K.: Eneström-kakeya theorem and some of its generalizations. In: Joshi, S., Dorff, M., Lahiri, I. (eds.) Current Topics in Pure and Computational Complex Analysis. Trends in Mathematics, pp. 171–199 (2014)
11. Gürlebeck, K., Malonek, H.: A hypercomplex derivative of monogenic functions in $\mathbb{R}^{n+1}$ and its applications. Complex Variables Theory Appl. **39**, 199–228 (1999)
12. Klein, F.: Über die Nullstellen der hypergeometrischen Reihe. Math. Ann. **37**, 573–590 (1890)
13. Koshy, T.: Catalan Numbers with Applications. Oxford University Press, New York (2009)
14. Kuijlaars, A.B.J., Martínez-Finkelshtein, A.: Strong asymptotics for Jacobi polynomials with varying nonstandard parameters. J. Anal. Math. **94**, 195–234 (2004)
15. McNamee, J.M.: Numerical Methods for Roots of Polynomials–Part I, Studies in Computational Mathematics, vol. 14, Elsevier (2007)
16. Martínez-Finkelshtein, A., Martínez-González, P., Orive, R.: Zeros of Jacobi polynomials with varying non-classical parameters. In: Special Functions (Hong Kong, 1999), pp. 98–113. World Scientific Publishing, River Edge (2000)
17. Rahman, Q.I., Schmeisser, G.: Analytic Theory of Polynomials. Oxford University Press, Oxford (2002)
18. Shen, J., Wang, L.: Legendre and Chebyshev dual Petrov-Galerkin methods for hyperbolic equations. Comput. Methods Appl. Mech. Eng. **196**(37–40), 3785–3797 (2007)
19. Stanley, R.P.: Log-concave and unimodal sequences in algebra, combinatorics, and geometry. Ann. N. Y. Acad. Sci. **576**, 500–535 (1989)

20. Szegö, G.: Orthogonal Polynomials, vol. 23, 4th edn. Amer. Math. Soc, Providence, RI (1975)
21. Vietoris, L.: Über das Vorzeichen gewisser trigonometrischer Summen. Sitzungsber. Österr. Akad. Wiss **167**, 125–135 (1958)
22. Wilf, H.: Generatingfunctionology, 2nd edn. Academic Press, San Diego (1994)

Interpolation Methods and Cryptographic Primitives in Secret Sharing

Thays Rocha Neri Ferreira(✉) and Fábio Borges

Laboratório Nacional de Computação Científica, Petrópolis, RJ, Brazil
academicthaysrocha@gmail.com
http://www.lncc.br/borges/

Abstract. Shamir Secret Sharing (SSS) relies fundamentally on polynomial interpolation for secret reconstruction, making the choice of interpolation method a decisive factor for computational performance. Since reconstruction is the core operational phase of the scheme, differences in interpolation strategies directly affect efficiency. This work presents a systematic and comparative evaluation of interpolation techniques applied to SSS, including Lagrange, Newton, barycentric Lagrange, and NTT-based methods. Furthermore, the impact of distinct cryptographic primitives—Pedersen Commitments, Paillier, ElGamal, and RSA—is analyzed within the framework of Pedersen Verifiable Secret Sharing. All approaches were implemented in Python within a unified experimental environment to ensure consistent performance evaluation. The analysis considers execution time and memory usage as the primary evaluation metrics. The experimental results demonstrate that interpolation methods constitute the dominant factor influencing both temporal performance and memory usage. In particular, FFT/NTT-based methods outperform the alternatives for larger threshold values, benefiting from their asymptotic efficiency. Conversely, classical Lagrange interpolation consistently exhibits the lowest memory consumption among the evaluated techniques. With respect to verification mechanisms, RSA-based schemes achieve the best overall performance across the tested scenarios. Collectively, these findings offer concrete and practical guidance for designing efficient and scalable Secret Sharing deployments, especially in contexts where computational resources and performance constraints must be carefully balanced.

Keyword: Shamir secret sharing, pedersen verifiable secret sharing, polynomial interpolation, cryptographic primitives

1 Introduction

According to the guidelines of the National Institute of Standards and Technology (NIST) [11], recent technological advances increasingly highlight the urgent need to strengthen and modernize currently used cryptographic protocols through the adoption of decentralized strategies capable of increasing

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 99–116, 2027.
https://doi.org/10.1007/978-3-032-30518-3_7

the resilience, robustness, and security of systems against threats. Among the decentralized cryptographic protocols, Secret Sharing schemes stand out, allowing multiple parties to collaborate on a project using their private inputs without the need for undue disclosure of confidential information.

Shamir Secret Sharing (SSS), proposed in [15], stands out among Secret Sharing schemes for offering theoretical information security and a simple mathematical foundation, as cited in [3]. Its construction is based on classical principles of polynomial interpolation over finite fields. These characteristics explain the widespread adoption of SSS in both theoretical and practical applications, and it also serves as a basis for more sophisticated extensions, such as verifiable schemes and applications in secure computing and distributed systems.

However, several polynomial interpolation methods are available in the literature, of which only a limited number of studies systematically investigate their application in the context of the Shamir Secret Sharing scheme. Alternatives have been explored with the aim of improving aspects such as computational efficiency, scalability, and memory usage. The most comprehensive study identified to date is presented in [7], which performs a comparative analysis of the performance of seven different interpolation methods, including Newton, barycentric Lagrange, and the Fast Fourier Transform (FFT) applied to the SSS scheme. Unlike previous studies that focus only on interpolation techniques in Shamir Secret Sharing, this work separates the cost of reconstruction from the cost of verification and provides a joint evaluation of interpolation methods and cryptographic primitives in Verifiable Secret Sharing under a unified experimental framework.

In the context of Verifiable Secret Sharing (VSS) schemes, an additional requirement arises, related not only to the confidentiality of the secret but also to the need to guarantee the integrity and consistency of distributed actions, even in the presence of dishonest participants. As described in [16], these schemes extend Shamir's Secret Sharing by incorporating cryptographic mechanisms that allow participants to efficiently verify whether received shares are valid and consistent with a single underlying polynomial, without revealing the secret.

Despite the practical relevance of VSS schemes, the literature is still limited regarding the systematic analysis of the methods employed in verification. To the best of our knowledge, there are few studies that comparatively evaluate the impact of different cryptographic primitives on the performance of Verifiable Secret Sharing schemes.

Thus, with the aim of addressing the gaps related to the choice of interpolation methods in Shamir Secret Sharing and the cryptographic primitives employed in Verifiable Secret Sharing schemes, this article presents a comprehensive analysis of the main polynomial interpolation techniques, including Lagrange, Newton, barycentric Lagrange, and the Fast Fourier Transform, as well as the cryptographic primitives and commitment schemes including Pedersen commitments and public-key systems such as Paillier, ElGamal, and RSA. All approaches considered were implemented in Python, allowing for a uniform and comparable evaluation of their computational performance.

In general, this article proposes the following main research questions:

(i) Which method is most execution-time efficient for secret reconstruction in large-scale scenarios?
(ii) Which method has the highest memory consumption for secret reconstruction in large-scale scenarios?
(iii) Which cryptographic primitive shows the best time performance in verifying participants' honesty in the Pedersen Verifiable Secret Sharing scheme?
(iv) Which cryptographic primitive exhibits the highest memory consumption in the Pedersen Verifiable Secret Sharing scheme?

These questions motivated and guided the research, forming the basis of this work. Thus, this article contributes primarily to the following:

(i) Systematic comparison of four interpolation techniques for secret reconstruction.
(ii) Evaluation of four cryptographic primitives in Pedersen-based VSS.
(iii) Experimental analysis isolating interpolation cost from verification cost.

The remainder of the article is organized as follows: Sect. 2 presents the essential foundations, covering algebraic theory, Secret Sharing, the considered cryptographic primitives, and the analyzed polynomial interpolation methods. Section 3 describes the experimental setup and the procedures adopted for execution time and memory consumption analysis. Section 4 presents the experimental results obtained. Finally, Sect. 5 concludes the article and outlines directions for future research.

2 Preliminaries

This section introduces the preliminary concepts underlying the proposed analysis, covering algebraic foundations, Secret Sharing, selected cryptographic primitives, and polynomial interpolation methods.

2.1 Algebraic Theory

This section introduces the algebraic foundations that support the cryptographic schemes and polynomial interpolation techniques employed in this work.

Finite Field. A finite field, also known as a Galois Field (GF), is an algebraic structure composed of a finite number of elements. In this set, the operations of addition and multiplication, as well as their inverses, subtraction and division, with the exception of division by zero, are well-defined and satisfy fundamental properties such as associativity, commutativity, and distributivity. The order of a finite field is always a power of a prime number, that is, p^n, where p is a prime number and n is a positive integer. For more details such as theorems and axioms, see [8].

Root of Unity. The roots of unity are elements of a modular ring that satisfy a property analogous to the classical case: a root of unity of order n is an element r such that $r^n \equiv 1 (\mod M)$, where M is the modulus chosen for the NTT. These roots are fundamental to the definition of the NTT (Number Theoretic Transform), which can be seen as a discrete analogue of the Fourier Transform, but operating in modular arithmetic. This approach allows for efficient transformations on integers, without rounding errors.

Definition 1. *Let $\mathbb{F}_p$ be a finite field with prime modulus p. An element $\omega \in \mathbb{F}_p$ is called an nth root of unity if it satisfies*

$$\omega^n \equiv 1 \pmod{p}. \tag{1}$$

Moreover, if $\omega^k \not\equiv 1 \pmod{p}$ for all $0 < k < n$, then ω is called a primitive nth root of unity. In this case, ω generates the multiplicative cyclic subgroup

$$\{1, \omega, \omega^2, \ldots, \omega^{n-1}\} \subset \mathbb{F}_p^*, \tag{2}$$

whose powers consist of exactly n distinct roots of unity.

According to [9], in the context of the FFT over finite fields, such roots of unity allow the construction of the Discrete Fourier Transform (DFT) matrix, enabling an efficient evaluation of polynomials at multiple points and resulting in a reduction of computational complexity from $O(n^2)$ to $O(n \log n)$.

2.2 Secret Sharing

Secret sharing is a cryptographic scheme that allows a secret, such as a cryptographic key, a password, or any other sensitive data, to be divided into multiple parts, denoting shares. These parts are distributed among different participants. The main characteristic of this scheme is that the original secret can only be reconstructed when a minimum number of participants, known as the threshold, cooperate by combining their actions. If the number of participants is less than this limit, the secret remains unrecoverable, and no relevant information can be compromised by individual actions or any subset that does not meet the established limits.

The Shamir Secret Sharing Scheme. The Shamir Secret Sharing scheme (SSS), proposed by Shamir [15], is based on polynomial interpolation over finite fields. The steps of the scheme are described below.

I **Share generation.** The dealer $\mathcal{D}$ selects a large prime number p, which defines the finite field, and, given the secret s, constructs a polynomial of degree $k-1$ of the form

$$f(x) = a_0 + a_1 x + a_2 x^2 + \cdots + a_{k-1} x^{k-1} \pmod{p}, \tag{3}$$

where $a_0 = s$, p is a prime satisfying $p > n$, with n denoting the total number of participants and k the reconstruction threshold. The coefficients $a_1, a_2, \ldots, a_{k-1} \in \mathbb{Z}_p$ are chosen uniformly at random.

II **Distribution.** The dealer $\mathcal{D}$ computes the n shares of the scheme, $\text{share}_i = (i, f(i))$, and sends them to each participant $\mathcal{P}_i$ through a secure channel, for $i = 1$ to n.

III **Reconstruction.** For the reconstruction phase, it is assumed[1] that the participants $\{\mathcal{P}_i\}_{i=1}^{k}$ send their *shares* $\{s_i\}_{i=1}^{k}$. Using Lagrange interpolation over these shares, the secret is recovered as follows:

$$s = f(0) = \sum_{i=1}^{k} y_i \lambda_i \tag{4}$$

with

$$\lambda_i = \prod_{j=1, j\neq i}^{k} \frac{x_j}{x_j - x_i} \tag{5}$$

where the term y_i represents the share value of participant $\mathcal{P}_i$.

Pedersen's Verifiable Secret Sharing. Pedersen's Verifiable Secret Sharing (VSS) scheme is a cryptographic method that, in addition to dividing the secret into distributed parts among the participants, allows for verification of the validity and consistency of the shares without revealing the secret itself. It uses commitments based on discrete logarithms to ensure that the parts are consistent and that the distributor is honest, offering robustness against fraud or incorrect parts, as stated by [10]. The steps of Pedersen's Verifiable Secret Sharing scheme are described below.

I **Commitment and share generation.** The dealer $\mathcal{D}$ selects a finite field $\mathbb{Z}_p$, where p is a large prime number, and two generators $g, h \in \mathbb{Z}_p^*$ such that the discrete logarithm of h with respect to base g is unknown. Given the secret s, the dealer constructs two polynomials of degree $k-1$:

$$f(x) = a_0 + a_1 x + \cdots + a_{k-1} x^{k-1} \pmod{p}, \tag{6}$$

$$r(x) = b_0 + b_1 x + \cdots + b_{k-1} x^{k-1} \pmod{p}, \tag{7}$$

where $a_0 = s$, b_0 is chosen uniformly at random, $p > n$, and all coefficients $a_i, b_i \in \mathbb{Z}_p$ are chosen uniformly at random, with the exception of a_0.
The dealer then computes the public commitments to the coefficients:

$$C_j = g^{a_j} h^{b_j} \pmod{p}, \quad j = 0, 1, \ldots, k-1. \tag{8}$$

II **Distribution.** The dealer $\mathcal{D}$ computes the n shares of the secret and the random value, $\text{share}_i = (i, f(i), r(i))$, and sends them to each participant $\mathcal{P}_i$, for $i = 1$ to n, through a secure channel. The commitments $\{C_j\}_{j=0}^{k-1}$ are made public.

[1] Note: the notation emphasizes a change in cardinality rather than in the indices of the participants.

III **Share verification.** Each participant $\mathcal{P}_i$ can verify the validity of their share by checking whether the following equality holds:

$$g^{f(i)}h^{r(i)} \equiv \prod_{j=0}^{k-1} C_j^{i^j} \pmod{p}. \tag{9}$$

If the equality does not hold, the participant detects dishonest behavior by the dealer.

IV **Reconstruction.** In the reconstruction phase, it is assumed that the participants $\{\mathcal{P}_i\}_{i=1}^{k}$ send their valid shares $\{f(i)\}_{i=1}^{k}$. Using Lagrange polynomial interpolation over these values, the secret is recovered as

$$s = a_0 = f(0) = \sum_{i=1}^{k} f(i) \prod_{\substack{j=1 \\ j\neq i}}^{k} \frac{-j}{(i-j)} \pmod{p}. \tag{10}$$

2.3 Cryptographic Primitives

Cryptographic primitives are fundamental elements in the construction of information security protocols and schemes, providing mathematical mechanisms that ensure properties such as confidentiality, integrity, protection, and verifiability. This section presents the main primitives used in this work: Pedersen Commitments, as well as those associated with the Paillier, ElGamal, and RSA public-key cryptosystems. Each relies on different hard mathematical problems, such as discrete logarithms and integer factorization, and exhibits specific properties.

The cryptographic primitives described below are based on the works of [5,12–14] respectively.

Pedersen Commitments

Definition 2. *Let p be a large prime, G a cyclic group of order q (with $q \mid p-1$), and let $g, h \in G$ be two generators such that the discrete logarithm of h with respect to base g is unknown. The Pedersen commitment scheme is defined as*

$$C = g^m h^r \pmod{p},$$

where $m \in \mathbb{Z}_q$ is the message to be committed and $r \in \mathbb{Z}_q$ is a random value.

Paillier Cryptosystem

Definition 3. *Let p and q be two large and distinct prime numbers, $n = pq$, and $\lambda = \operatorname{lcm}(p-1, q-1)$. Consider the multiplicative group $\mathbb{Z}_{n^2}^*$ and an element $g \in \mathbb{Z}_{n^2}^*$ such that $\gcd(L(g^\lambda \bmod n^2), n) = 1$, where the function L is defined by*

$$L(u) = \frac{u-1}{n}.$$

The Paillier cryptosystem is defined as follows:

- ***Public key:*** (n, g);
- ***Private key:*** λ.

The encryption of a message $m \in \mathbb{Z}_n$ *is given by* $c = g^m r^n \pmod{n^2}$, *where* $r \in \mathbb{Z}_n^*$ *is chosen uniformly at random.*

ElGamal Cryptosystem

Definition 4. *Let* G *be a cyclic group of order* q *generated by* g, *where the discrete logarithm problem is computationally hard. Let* $x \in \mathbb{Z}_q$ *be the private key and let* $h = g^x$ *be the public key.*

The ElGamal cryptosystem is defined as follows:

- ***Public key:*** (G, q, g, h);
- ***Private key:*** x.

The encryption of a message $m \in G$ *is given by* $(c_1, c_2) = (g^r, m \cdot h^r)$, *where* $r \in \mathbb{Z}_q$ *is chosen uniformly at random.*

RSA Cryptosystem

Definition 5. *Let* p *and* q *be two large and distinct prime numbers,* $n = pq$, *and* $\varphi(n) = (p-1)(q-1)$. *Choose an integer* e *such that* $\gcd(e, \varphi(n)) = 1$, *and let* d *be the multiplicative inverse of* e *modulo* $\varphi(n)$, *that is,*

$$ed \equiv 1 \pmod{\varphi(n)}.$$

The RSA cryptosystem is defined as follows:

- ***Public key:*** (n, e);
- ***Private key:*** d.

The encryption of a message $m \in \mathbb{Z}_n$ *is given by* $c = m^e \pmod{n}$, *and the decryption by* $m = c^d \pmod{n}$.

2.4 Polynomial Interpolation Methods

In the context of interpolation over a finite field $\mathbb{F}_p$ (where p is a prime number or a power of a prime), the operations of addition, subtraction, multiplication, and division (except by zero) are performed modulo p or modulo an irreducible polynomial over $\mathbb{F}_p$ (in the case of $p = q^k$).

Several methods have been developed to perform polynomial interpolation, each with its own characteristics, as well as computational advantages and improvements. We will analyze four of them, namely, Lagrange, Newton, barycentric Lagrange, and the Fast Fourier Transform (FFT), known in the finite domain as the Number Theoretic Transform (NTT).

For standardization, given a set of $n + 1$ distinct points

$$(x_0, y_0), (x_1, y_1), \ldots, (x_n, y_n), \tag{11}$$

where all the values x_i are distinct from each other, the main interpolation methods mentioned are defined below.

Lagrange. The Lagrange interpolation polynomial is the unique polynomial of degree at most n that satisfies $P(x_i) = y_i$ for all $i = 0, 1, \ldots, n$.

According to [4], the general form of the classical Naive Lagrange polynomial, also known simply as the Lagrange polynomial, $P(x)$ is given by

$$P(x) = \sum_{j=0}^{n} y_j L_j(x). \tag{12}$$

In this expression, $L_j(x)$ are the Lagrange basis polynomials, defined by

$$L_j(x) = \prod_{\substack{k=0 \\ k \neq j}}^{n} \frac{x - x_k}{x_j - x_k}. \tag{13}$$

Each polynomial $L_j(x)$ satisfies the interpolation property $L_j(x_j) = 1$ and $L_j(x_k) = 0$ for all $k \neq j$.

Newton. The Newton interpolation polynomial is the unique polynomial of degree at most n that satisfies $P(x_i) = y_i$ for all $i = 0, 1, \ldots, n$. According to [4], the general form of the Newton polynomial $P_n(x)$ is given by

$$\begin{aligned} P_n(x) =& f[x_0] + f[x_0, x_1](x - x_0) + f[x_0, x_1, x_2](x - x_0)(x - x_1) + \\ & f[x_0, x_1, \ldots, x_n](x - x_0)(x - x_1) \ldots (x - x_{n-1}). \end{aligned} \tag{14}$$

In this expression, $f[x_i, x_{i+1}, \ldots, x_j]$ denotes the divided differences of the function f. The divided differences are computed recursively, as described below:

– **Zeroth-order divided difference:**

$$f[x_i] = y_i. \tag{15}$$

– **Divided difference of order $k \geq 1$:**

$$f[x_i, x_{i+1}, \ldots, x_{i+k}] = \frac{f[x_{i+1}, \ldots, x_{i+k}] - f[x_i, \ldots, x_{i+k-1}]}{x_{i+k} - x_i}. \tag{16}$$

Barycentric. The barycentric interpolation method is a reformulation of Lagrange interpolation, designed to improve numerical stability and computational efficiency. Instead of directly computing the Lagrange basis polynomials, barycentric interpolation uses precomputed weights. As described in [2], the interpolating polynomial $P(x)$ associated with $n + 1$ distinct points can be expressed in barycentric form as

$$P(x) = \frac{\sum_{j=0}^{n} \frac{w_j y_j}{x - x_j}}{\sum_{j=0}^{n} \frac{w_j}{x - x_j}}, \tag{17}$$

where the barycentric weights are defined by

$$w_j = \frac{1}{\prod_{\substack{k=0 \\ k \neq j}}^{n} (x_j - x_k)}. \tag{18}$$

NTT. The Fast Fourier Transform (FFT) is an efficient algorithm for computing the Discrete Fourier Transform (DFT). The FFT makes this computation significantly faster than the direct DFT approach, especially for large-scale data sets. However, when considered over a finite domain, known as the Number Theoretic Transform (NTT), it operates on sequences of numbers embedded in a discrete and finite mathematical system.

According to [9], in a finite field $\mathbb{F}_p$, with p prime, the DFT of a sequence $a = (a_0, a_1, \ldots, a_{n-1})$ is defined as

$$A_k = \sum_{j=0}^{n-1} a_j \, \omega^{jk} \pmod{p}, \quad k = 0, 1, \ldots, n-1, \tag{19}$$

where $\omega \in F_p$ is a primitive nth root of unity in $\mathbb{F}_p$, that is, $\omega^n \equiv 1 \pmod{p}$ and $\omega^k \not\equiv 1 \pmod{p}$ for $0 < k < n$.

The inverse transform, which recovers the original sequence a_j from A_k, is given by

$$a_j = n^{-1} \sum_{k=0}^{n-1} A_k \, \omega^{-jk} \pmod{p}, \tag{20}$$

where n^{-1} denotes the multiplicative inverse of n modulo p.

As in the FFT, the NTT employs a divide-and-conquer strategy, splitting the problem into two smaller NTTs, one associated with even indices and the other with odd indices, and combining the results as follows:

$$A_k = A_k^{(\text{even})} + \omega^k A_k^{(\text{odd})} \pmod{p}, \tag{21}$$

$$A_{k+n/2} = A_k^{(\text{even})} - \omega^k A_k^{(\text{odd})} \pmod{p}, \tag{22}$$

for $k = 0, 1, \ldots, \frac{n}{2} - 1$.

This approach reduces the computational complexity from $O(n^2)$, characteristic of the direct modular DFT, to $O(n \log n)$, making the NTT particularly suitable for applications that require fast and accurate modular operations, such as optimized secret reconstruction in Secret Sharing schemes.

Although the FFT and the NTT share the same algorithmic structure, the NTT imposes additional mathematical constraints that affect its practical application. Implementing the NTT requires extra care to avoid arithmetic overflow and to ensure computational efficiency, especially when large moduli are used. As a result, the NTT is less flexible than the FFT in general-purpose signal processing applications, but extremely valuable in contexts that demand modular arithmetic.

3 Experimental Setup and Simulation

The experiments were conducted on a workstation equipped with an Intel Core i9-12900K (16 cores, 3.2 GHz), 125 GB RAM,Ubuntu 22.04, Python 3.12.3. The reported results correspond to the average of 10 independent repetitions of each experiment, with the goal of mitigating the presence of outliers. Execution time was measured using the built-in `time` module, and memory consumption was monitored using the `tracemalloc` library.

For modular inversion, Fermat's Little Theorem was employed. This simple and uniform approach is executed in logarithmic time with respect to p, as stated in [17]. For one experiment, we defined the finite field with $p = 2013265921$, the secret $= 987654321$, the generators $g = 31$ and $h = 17$, in addition to (Table 1).

Table 1. Values of n_{sss} (total number of participants) and k_{sss} (minimum threshold for reconstruction).

n_{sss}	8	16	32	64	128	256	512	1024	2048	4096	8192	16384
k_{sss}	4	8	16	32	64	128	256	512	1024	2048	4096	8192

3.1 Time Analysis

The first topic to be analyzed is the difference in computational performance depending solely on the interpolation method used for secret reconstruction. To this end, we modified the traditional Shamir Secret Sharing method to allow the use of the methods mentioned in Sect. 2.4. Next, for each interpolation method analyzed previously, we added a function to verify the honesty of the participants based on Pedersen's Verifiable Secret Sharing (VSS). For this, we used the cryptographic primitives described in Sect. 2.3. In all analyses performed, the unit of measurement considered is seconds.

3.2 Memory Analysis

Similarly to the previous analysis, memory usage was first evaluated solely for the interpolation method. Subsequently, the cryptographic primitives were assessed according to the adopted interpolation method, enabling a combined analysis of actual memory usage throughout the entire process, considering both components. Finally, peak memory usage was also analyzed for all previously described scenarios. The key difference is that peak memory refers to the maximum amount allocated at any point during the monitoring period, regardless of later releases, whereas memory usage represents consumption at the moment of measurement. In all analyses, the unit of measurement used was megabytes (MB).

4 Experimental Results

This section presents the experimental results obtained from the evaluation of the considered interpolation methods and cryptographic primitives, focusing on computational performance through execution time and memory consumption analysis. Due to the nature of the results, showing differences on the order of microseconds and very close values, we chose to present all graphs on a logarithmic scale. This allows for a more balanced comparison among the methods and clearer identification of trends that would be difficult to observe in a linear representation.

4.1 Time Analysis

As described in Sect. 3, the first aspect analyzed is the variation in computational performance as a function solely of the interpolation method used in the secret reconstruction. Thus, we present Fig. 1. In it, we can see that, among the methods mentioned, the Newton interpolation method, indicated in orange, showed the worst performance, being considered the slowest. Next, the Lagrange method, indicated in blue, and the Barycentric Lagrange, indicated in purple, appear, while the NTT-based method, indicated in pink, obtained the best computational performance.

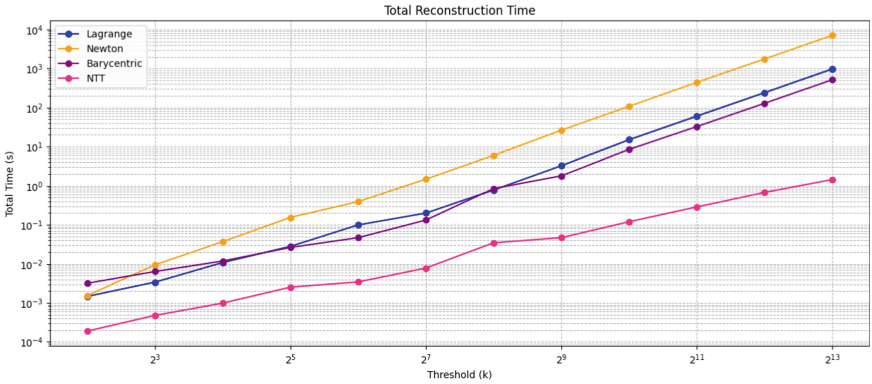

Fig. 1. Time analysis for secret reconstruction depends on the interpolation method used.

Comparing the asymptotic complexity of the methods, the NTT approach showed the best performance, which was expected due to its $O(n \log n)$ complexity, in contrast to the $O(n^2)$ complexity of the other methods. Although Newton's method shares the same quadratic complexity, it involves divided difference computations and recursive polynomial evaluation, increasing the number of required operations and, consequently, processing time. Lagrange's method computes the interpolating polynomial directly, requiring multiplications and divisions for each

point. Barycentric interpolation optimizes Lagrange's formulation by reducing the number of operations during polynomial evaluation, making it slightly more efficient.

It is also worth highlighting the most significant occurs at the equivalence threshold of 2^5, which is when the Lagrange and Barycentric methods intersect. At 2^8, they become equivalent, but at 2^9 the Barycentric method once again stands out as faster than the other two.

Now, in Fig. 2, we present an analysis of the impact of verifying the honesty of participants using Pedersen's VSS with the aforementioned cryptographic primitives, on the interpolation methods. This analysis shows that the behavior of the curves, regardless of the cryptographic primitive used, is very similar. Consequently, we can conclude that the type of primitive used does not significantly influence the temporal performance.

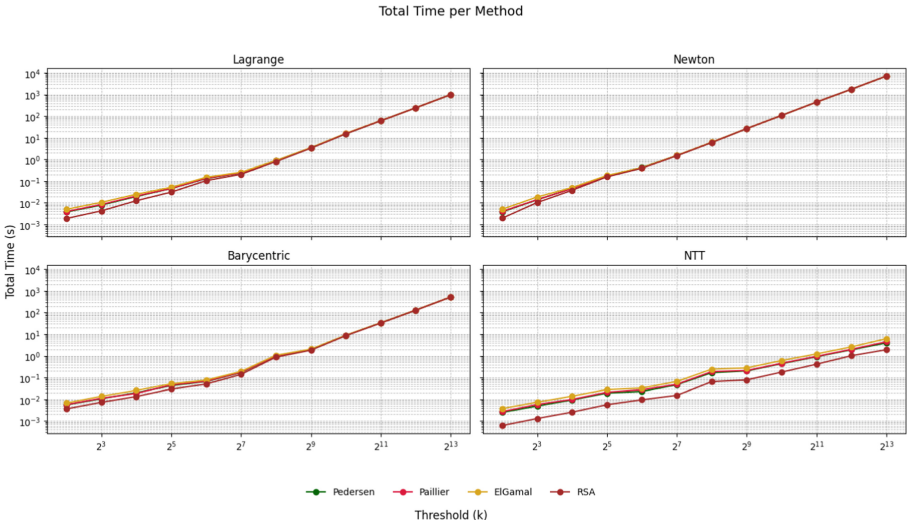

Fig. 2. Time analysis for share verification depends on the interpolation method and the cryptographic primitive used.

4.2 Memory Analysis

As with the previous analysis, the first aspect analyzed is the variation in memory usage as a function solely of the interpolation method used in secret reconstruction. Thus, we present Fig. 3. In it, we can see that, among the methods analyzed, NTT presented the worst performance, being the algorithm that consumes the most memory throughout the entire process. Next, Newton's method stands out, exhibiting oscillatory behavior in the interval from 2^2 to 2^6 interval, stabilizing from 2^7 onwards, when it begins to show practically constant memory

consumption. In turn, the Barycentric and Lagrange methods showed constant behavior, regardless of the threshold value k. Lagrange's method stands out, having obtained the best overall performance, registering the lowest memory consumption among the methods evaluated.

NTT performs complete transformations and complex modular operations, consuming temporary memory during the process. Newton's method constructs a divided difference table, requiring the storage of all intermediate subresults, which significantly increases memory usage. The oscillation observed in the interval from 2^2 to 2^6 can be attributed to the incremental pattern of construction of this table: in smaller sizes, small variations in the number of points cause internal reorganizations and reallocations of data structures, temporarily impacting memory usage. From 2^7 onwards, the structure tends to stabilize, operating with a more uniform allocation pattern and, consequently, with practically constant consumption.

The Barycentric Lagrange is an optimized version of Lagrange interpolation, designed to reduce the cost of evaluating polynomials at multiple points. However, it requires storing the interpolation weights, which increases temporary memory allocation. In contrast, the direct Lagrange method calculates each point independently, without the need to store large intermediate tables, resulting in minimal memory usage and more straightforward calculations.

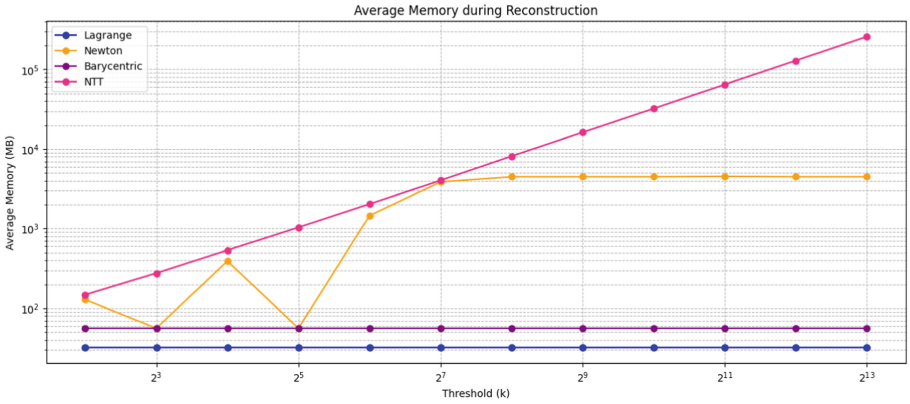

Fig. 3. Memory analysis used to reconstruct the secret depends on the interpolation method used.

Now, analyzing only the cryptographic primitives, regardless of the interpolation method, we observe a quite similar behavior among them. In Fig. 4, it can be observed that Pedersen, Paillier, ElGamal, and RSA present practically overlapping curves in all scenarios, with marginal and stable differences as the threshold k increases.

This pattern indicates that memory consumption is strongly influenced by the structure of the reconstruction method adopted and not by the cryptographic

primitive itself. The choice of interpolation technique has a much more significant impact on memory usage than the cryptographic scheme used for verification or protection of shares. In summary, we can present the memory usage of the cryptographic primitives as shown in Table 2 and Fig. 5 below, with the maximum difference between the primitives being 100 MB (Fig. 4).

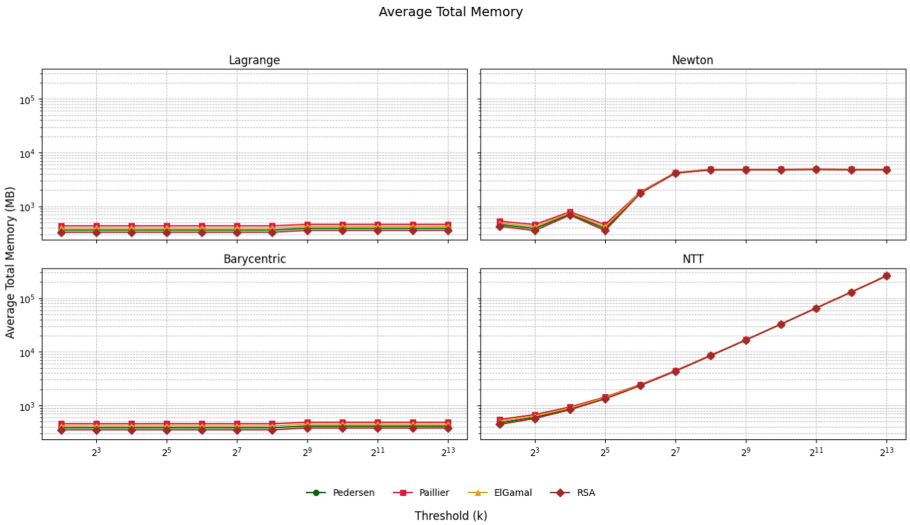

Fig. 4. Memory analysis as a function of the cryptographic primitive used.

Using Fig. 5 as a basis and comparing the primitives among themselves, RSA consistently exhibits the lowest memory usage. This is expected, as its internal operations essentially involve modular exponentiation with few auxiliary objects. However, it is worth noting that this is also a result of the parameters used in the simulation. The public exponent equivalent to $e = 65537$ was used, which is commonly used in RSA cryptosystems because it provides a balance between security and computational efficiency, but small enough to allow efficient cryptography, as stated by [6]. Furthermore, the values of p and q have a size of 16 bits, and the modulus of N has a size of 32 bits. These parameter sizes are intentionally insecure and were selected exclusively to evaluate computational performance rather than cryptographic security. According to [1], these values are considered insecure, but adequate for performance tests, such as evaluating execution time and memory usage, where security is not the primary focus.

Pedersen Commitments come next, with a slightly higher cost. This increase is explained by the need to store commitments, generators, and auxiliary values for verification. The ElGamal algorithm consumes more memory than the Pedersen and RSA algorithms, reflecting the fact that it operates with pairs of group elements and additional random values, which duplicates part of the internal structures during encryption and verification. Paillier's algorithm, on the other

Table 2. Summary of average memory usage (in MB) for different encryption primitives.

k	Pedersen	Paillier	ElGamal	RSA
2^2	328	400	360	296
2^3	328	400	360	296
2^4	328	400	360	296
2^5	328	400	360	296
2^6	328	400	360	296
2^7	328	400	360	296
2^8	356	428	388	324
2^9	356	428	388	324
2^{10}	356	428	388	324
2^{11}	356	428	388	324
2^{12}	356	428	388	324
2^{13}	356	428	388	324

hand, is the most memory-intensive in all scenarios. This is directly related to the use of large integers, operations on larger moduli (n^2), and the maintenance of multiple intermediate values during homomorphic encryption and decryption, which significantly increases the pressure on memory.

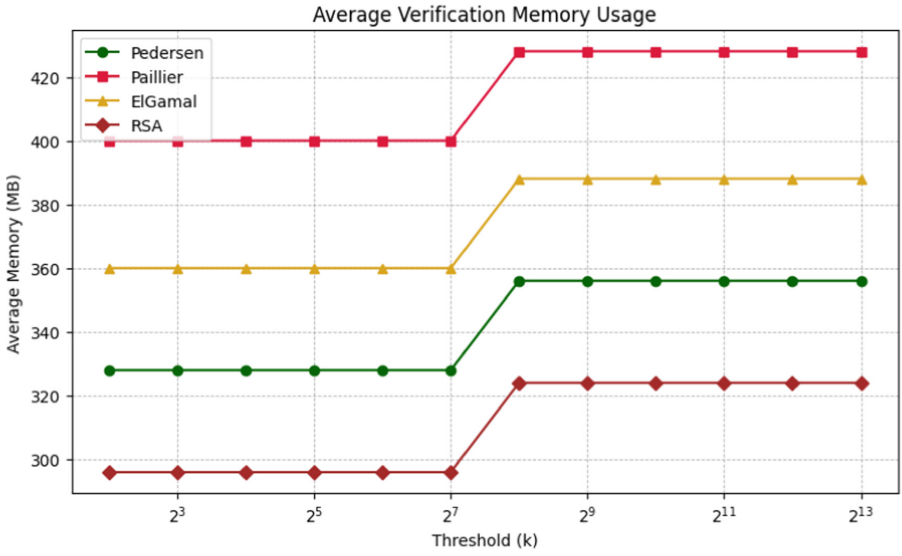

Fig. 5. Average memory usage (in MB) for different encryption primitives.

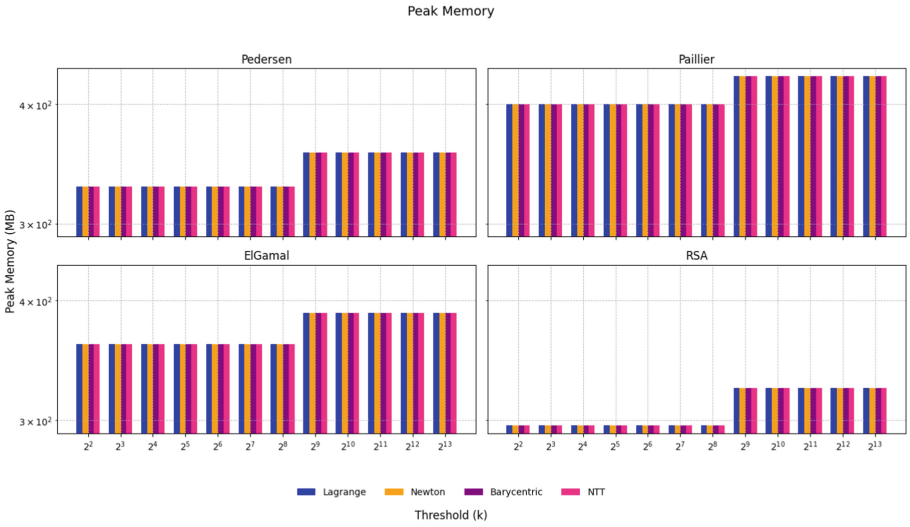

Fig. 6. Analysis of the peak memory required for secret reconstruction as a function of the interpolation method and the cryptographic primitive.

We also analyzed the memory peak, shown in Fig. 5, to reveal the most critical moment of execution. The memory peak is what determines the execution prediction on devices with limited memory, such as smart meters, for example. In all cryptographic primitives, the memory peak was constant, with jumps between limits equivalent to 2^8 and 2^9. This indicates that the peak does not gradually increase with the number of participants, but is controlled by internal allocation limits, typical of vectors (Fig. 6).

The influence of primitive cryptography, on the other hand, is quite clear in the absolute value of the peak. RSA consistently shows the lowest memory peaks, reflecting its simpler structure. Pedersen and ElGamal occupy an intermediate range, while Paillier shows the highest memory peaks for all values of k.

5 Conclusion

First, this work analyzes the performance of the Lagrange interpolation method in Shamir's Secret Sharing compared to other approaches. Although the security of the scheme stems from the polynomial construction over a finite field, the cost of reconstructing the secret depends directly on the interpolation method adopted. By comparing them, it is possible to dissociate the intrinsic cost of the cryptographic scheme from that introduced by the reconstruction algorithm, allowing for a more precise and reproducible evaluation.

Similar logic applies to the verification mechanism in Pedersen's VSS. Considering cryptographic primitives alternative to Pedersen's commitments allows for evaluating how different security assumptions impact performance and resource

consumption. This makes it possible to analyze the trade-offs between security, efficiency, and implementability, isolating the effect of the commitment primitive and selecting the most suitable option without compromising the guarantees of the VSS.

The study, therefore, separates two layers that are usually combined: interpolation cost and verification cost. In interpolation, the NTT-based method presents the best time. In contrast, NTT requires higher memory consumption. This difference is relevant in practice: NTT is more suitable for scenarios requiring low latency and frequent reconstructions, such as aggregation in smart meters, while Lagrange is preferable in memory-constrained environments, such as embedded devices and IoT sensors.

Regarding primitives, the RSA-based scheme consistently presents the shortest execution times for all values of k and interpolation methods, followed by Pedersen, ElGamal, and Paillier, the latter with the highest computational cost. From 2^8 onwards, a discrete and uniform increase in memory is observed across all primitives, suggesting the allocation of auxiliary structures. In general, the impact of k on memory is reduced when compared to the influence of the chosen primitive.

Despite the measurable differences between the primitives, the dominant factor for time and memory is the interpolation method. Thus, the inclusion of verification steps is technically justified, as its cost does not significantly alter the overall performance profile. Its adoption enhances the reliability and robustness of the system, mitigating malicious behavior or failures without significant penalties, strengthening the scheme for applications that require explicit guarantees of correctness and security.

Future work includes extending the analysis to larger cryptographic parameters, implementing the algorithms in compiled languages, and evaluating their performance in distributed environments and real-world applications such as secure multiparty computation and federated learning systems.

Acknowledgments. We thank the National Laboratory for Scientific Computing (LNCC) and CAPES for their institutional support, as well as the team of collaborators for their valuable contributions that enriched the development of this work.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Amalarethinam, D.: Enhanced RSA algorithm with varying key sizes for data security in cloud. 2017 World Congress on Computing and Communication Technologies (WCCCT) pp. 172–175 (2017). https://doi.org/10.1109/wccct.2016.50
2. Berrut, J.P., Trefethen, L.N.: Barycentric Lagrange interpolation. SIAM Rev. **46**(3), 501–517 (2004). https://doi.org/10.1137/S0036144502417715
3. Bogdanov, D.: Foundations and properties of Shamir's Secret Sharing Scheme – Research Seminar in Cryptography. University of Tartu, Institute of Computer Science **1** (2007)

4. Burden, R.L., Faires, J.D.: Numerical Analysis, 10th edn. Cengage Learning, Boston, MA, USA (2021)
5. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Trans. Inf. Theory **31**(4), 469–472 (1985)
6. Feng, Y., Nitaj, A., Pan, Y.: Small public exponent brings more: Improved partial key exposure attacks against RSA. IACR Cryptol. ePrint Arch **2024**, 1329 (2024). https://doi.org/10.62056/ahjbhey6b
7. Ferreira, T.R.N., Borges, F.: Performance evaluation of interpolation techniques in shamir secret sharing scheme. IEEE Access (2026). https://doi.org/10.1109/ACCESS.2026.3679331
8. Garling, D.J.H.: Finite fields. In: Galois Theory and Its Algebraic Background, Cambridge University Press (2021)
9. Lidl, R., Niederreiter, H.: Introduction to Finite Fields and Their Applications, 2nd edn. Cambridge University Press, Cambridge (1997)
10. Mahmoud, Q.A.: Polynomial differential-based information-theoretically secure verifiable secret sharing. International Journal of Information Technology and Computer Science **6**, 18–23 (2014). https://doi.org/10.5815/ijitcs.2014.12.03
11. National Institute of Standards and Technology: Nist threshold cryptography project. https://csrc.nist.gov/Projects/threshold-cryptography (2025), accessed: Jul. 21, 2025
12. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: International conference on the theory and applications of cryptographic techniques, pp. 223–238. Springer (1999)
13. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Annual international cryptology conference, pp. 129–140. Springer (1991)
14. Rivest, R.L., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. Commun. ACM **21**(2), 120–126 (1978)
15. Shamir, A.: How to share a secret. Commun. ACM **22**(11), 612–613 (1979). https://doi.org/10.1145/359168.359176
16. Stadler, M.: Publicly verifiable secret sharing pp. 190–199 (1996). https://doi.org/10.1007/3-540-68339-9_17
17. Xu, S., Gu, H., Wang, L., Guo, Z., Liu, J., Lu, X., Gu, D.: Efficient and constant time modular inversions over prime fields. 2017 13th International Conference on Computational Intelligence and Security (CIS) pp. 524–528 (2017). https://doi.org/10.1109/cis.2017.00122

Computational and Applied Statistics for Data & AI (CAS 2026)

Predicting Periodontal Disease Using Interpretable Machine Learning Models

Joana M. A. Lopes[1], Ana C. Braga[1(✉)], and Marta M. M. Relvas[2]

[1] ALGORITMI Research Centre, LASI, University of Minho, Braga, Portugal
acb@dps.uminho.pt
[2] UNIPRO IUCS CESPU, Gandra, Portugal
marta.relvas@iucs.cespu.pt

Abstract. As well known that periodontal disease is a chronic inflammatory condition and constitutes a major cause of tooth loss globally. An early identification of individuals at risk remains challenging due to the complex interplay among clinical, behavioural, systemic, and biological factors. The present study evaluates supervised Machine Learning models for predicting periodontal disease by integrating multiple data types, with a particular focus on clinical interpretability.

A dataset with 189 participants containing variables as clinical measures of periodontal health, oral hygiene routines, smoking status, diabetes status, and salivary inflammatory marker concentrations was analysed. Several classification algorithms were assessed, including Naive Bayes, Decision Trees, Random Forests, Support Vector Machines, K-Nearest Neighbours, and Logistic Regression. Model performance was evaluated using accuracy, sensitivity, specificity, F1-score, and area under the receiver operating characteristic (ROC) curve.

From this study, Logistic Regression exhibited robust predictive performance and interpretability, supporting its potential utility in clinical decision-making. Some key predictors were find significant as periodontal indicators, behavioural risk factors, and inflammatory biomarkers, particularly IL-1β, IL-6, and IL-8. These results indicate that interpretable Machine Learning approaches may enhance the identification of periodontal risk and inform personalised prevention strategies.

Keywords: Periodontal Disease · Machine Learning · Risk Factors · Inflammatory Biomarkers · Clinical Decision Support

1 Introduction

Periodontal disease (PD) is a chronic inflammatory condition of the tooth-supporting structures and remains an important public health concern. Although advances in preventive care have been made, it is still a leading cause of tooth loss, affecting quality of life, and could be linked to systemic medical problems. So, early diagnosis and risk stratification are important, and traditional diagnostics often focus only on isolated factors and may not reflect the disease's complexity.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 119–132, 2027.
https://doi.org/10.1007/978-3-032-30518-3_8

Recent advances in data science and machine learning (ML) allow for integration among diverse clinical and biological data, affording new opportunities for healthcare prediction. In periodontology, ML is being investigated to improve diagnosis and risk-factor identification, but its broad clinical use is limited by challenges related to interpretability and transparency.

The main objective of our study is to identify machine learning models for predicting periodontal disease that demonstrate both robust predictive accuracy and practical clinical utility. By combining clinical parameters, behavioural factors, systemic conditions, and salivary inflammatory biomarkers, we expect to improve risk assessment and support more effective clinical decision-making.

This manuscript is organized into seven main sections. Section 1 introduces concepts relating to periodontal disease and highlights the importance of interpretable machine learning models for clinical risk stratification. Section 2 provides a brief review of the literature on aetiological factors, salivary biomarkers and supervised algorithms relevant to periodontology. Sections 3 and 4 describe the procedure, including the dataset of 189 individuals, data pre-processing and comparisons between six algorithms. Section 5 presents the model results and the contribution of features, demonstrating the superior performance of Logistic Regression. Finally, Sects. 6 and 7 discuss the clinical implications of the selected model, as well as the limitations and directions for future research.

2 Review of the Literature

As described in [28], periodontal disease is a chronic, multifactorial inflammatory condition that affects the tissues supporting the teeth: the gingiva, the periodontal ligament, the root cementum and the alveolar bone. Initiated by the accumulation of bacterial biofilm on tooth surfaces, this process triggers an immune-inflammatory response in the host. Its progression and severity are influenced by risk factors, including genetic susceptibility, smoking, systemic conditions such as diabetes, and local factors such as dental anatomy and defective restorations [28]. Thus, PD is best seen as a host-biofilm interaction, where the disease trajectory depends on the balance between microbial challenge and host response [9,21]. Rather than a single dominant pathogen, PD involves a dysbiotic shift in biofilm composition. The oral microbiome moves from mainly Gram-positive aerobic species to Gram-negative anaerobes, thereby altering community behaviour and sustaining chronic inflammation [14]. This microbiome represents a dynamic ecosystem in which hundreds of species interact and adapt to changes in nutrient levels, mechanical stress, and pH. Understanding this complexity is key to creating effective prevention and treatment strategies for periodontal diseases [12].

Oral healthcare services and public awareness of oral hygiene have improved in Europe over the past two decades. However, the prevalence of periodontitis has remained stable, which is concerning, as it is a leading cause of tooth loss and a significant public health issue [6]. According to the World Health Organisation, severe periodontal diseases affect about 19% of adults, more than one

billion cases worldwide. Dental conditions, including periodontitis, result in substantial direct treatment costs, estimated at 298 billion US dollars annually, or 4.6% of global health expenditure, alongside considerable indirect costs due to productivity loss [6,19]. Some authors report that in several European countries, oral healthcare can represent almost a third of total healthcare expenses. For example, in Germany, 30–40% of tooth extractions are still due to periodontal disease [6]. This ongoing trend contributes to additional economic burdens, as post-operative discomfort and pain following periodontal surgery negatively affect quality of life and productivity, a situation further exacerbated by lost working hours in countries such as Canada [6].

Economic evaluations show that preventing periodontitis by effectively managing gingivitis is cost-effective. Programs such as supervised toothbrushing, water fluoridation, and the provision of dental supplies in early childhood settings have yielded significant savings for health systems [6]. These authors noted that initiatives like Mouth Care Matters in the United Kingdom aim to increase healthcare professionals' awareness of the link between oral hygiene and general health, with projections indicating substantial direct and broader social benefits for each unit of currency spent [6]. Modelling studies tracking transitions between periodontal health states reveal that individuals with gingivitis have a significant chance of progressing to periodontitis during each follow-up cycle, and some progress even with treatment. These data highlight the importance of early, accurate diagnosis and draw attention to how inconsistent clinical definitions complicate timely disease identification and classification [6].

Recent literature has examined the role of inflammatory mediators in periodontal disease pathogenesis and monitoring, with particular focus on salivary biomarkers such as interleukins IL-1β, IL-6, and IL-8. These biomarkers reflect local immune-inflammatory activity and could aid in risk stratification and personalised care [2,3,7,8,15,17,22]. The biological complexity of periodontitis, characterised by heterogeneity and overlapping risk profiles, has sparked interest in computational methods that integrate clinical, behavioural, sociodemographic, and molecular data [4,24,27]. Within this framework, ML methods—such as Naive Bayes, K-Nearest Neighbours, Decision Trees, Random Forests, Support Vector Machines, and Logistic Regression—are increasingly used in periodontal and biomedical datasets. These approaches seek to improve disease prediction and severity characterisation, exploit nonlinear relationships, and help identify key predictors for clinical support [1,11,18,20,27]. Several studies find that models like Logistic Regression provide strong discriminative performance and interpretability, supporting easier clinical integration through risk scoring and early detection [10,25].

Further supporting these trends, recent studies have emphasised the potential of salivary inflammatory mediators as biomarkers for assessing periodontal disease severity and progression [3,7,16]. Incorporating these biomarkers into machine learning frameworks has demonstrated improved predictive performance. However, challenges persist regarding data preprocessing, feature selection, and clinical translation [1,4,10,11,24,27].

3 Materials and Methods

3.1 Dataset

This study analyzed clinical data from 189 individuals to determine the presence or absence of PD. Data collection followed process no. 22/2022 of the Health Ethics Committee of the Centro Hospitalar do Tâmega e Sousa and were recorded in an Excel database.

The database includes adult participants aged 19–72 years who underwent periodontal clinical examinations, as well as information from questionnaires on oral hygiene habits and tobacco use. Variables also included measurements of inflammatory mediators and cytokines. Saliva samples were collected according to approved ethical protocols, and cytokine levels were quantified using commercial kits.

For participants without diabetes mellitus (DM), salivary cytokine levels were measured in pg/mL using the *LEGENDplex^{TM}Human Macrophage/Microglia Panel (13-plex) Kit (Cat. 740503, BioLegend, San Diego, CA, USA)*. This flow cytometry-based kit allows simultaneous detection of 13 human cytokines and chemokines relevant to immune responses. The assay was performed according to the manufacturer's protocol using fluorescence-coded microbeads specific to each analyte. Readings were acquired by flow cytometry and analyzed with the *LEGENDplex^{TM}Data Analysis Software Suite*. The detection limit for the cytokines included in this study was 2.44 pg/mL [5].

In participants with DM, salivary cytokine concentrations were measured in pg/mL using the LEGENDplex^{TM}Human Essential Immune Response Panel (13-plex) Kit (ref. 740930, BioLegend, CA, USA), which quantifies 13 analytes simultaneously. The assay employed Luminex®xMAP technology, which utilises magnetised microbeads with distinct fluorescent codes to enable highly sensitive multiplex analysis. The experimental protocol adhered to the manufacturer's instructions, including incubation with capture and detection antibodies and measurement on a Luminex®-compatible system. Detection limits were 2.44 pg/mL for most cytokines, except IL-1β (2.69 pg/mL), MCP-1 (2.2 pg/mL), IL-17A (3.66 pg/mL), IL-6 (2.2 pg/mL), IL-8 (2.69 pg/mL), and TGF-β (5.37 pg/mL) [5].

3.2 Data Preprocessing

Data preprocessing comprised exploratory analysis, handling of missing values, normalisation of continuous variables, and logarithmic transformation of skewed biomarker distributions. Values below the detection limit for inflammatory mediators were treated using established imputation strategies. Fig. 1 illustrates the preprocessing flow applied to the dataset, from the original dataset containing 31 variables to the construction of the final dataset with 21 variables, properly processed and prepared for the application of ML algorithms. This workflow ensured the consistency and quality of the data used in predictive modeling.

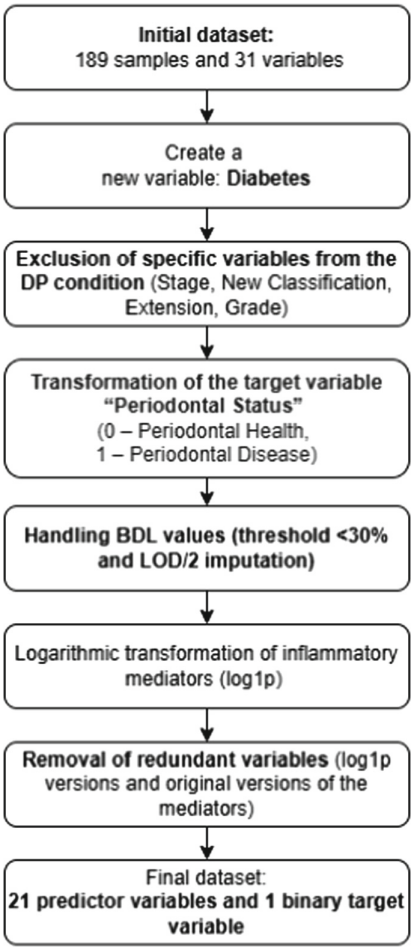

Fig. 1. Data preprocessing workflow.

3.3 Dataset Preparation, Model Selection and Clinical Translation Strategy

The dataset was split into training (80%) and testing (20%) subsets using stratified sampling, which preserved the original class proportions of "Periodontal Health" (38.3% training, 36.8% testing) and "Periodontal Disease" (61.7% training, 63.2% testing). This approach ensured class representativeness across both subsets, minimizing bias in model evaluation. The features were scaled using Standard Scaler to ensure algorithms sensitive to feature magnitudes, such as K-Nearest Neighbors (KNN) and Support Vector Machines (SVM), performed optimally.

Six supervised ML classification algorithms were selected and trained:

1. Naive Bayes (NB)
2. KNN
3. Decision Tree (DT)
4. Random Forest (RF)
5. SVM
6. Logistic Regression (LR)

Model performance was assessed using standard evaluation metrics, including Accuracy, Precision, Recall, F1-Score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC).

To facilitate clinical applicability, a structured methodology for model selection was used. Six binary classification algorithms were initially trained to distinguish between periodontal health and periodontal disease. The top three models were then selected based on their performance. Feature importance analysis was conducted in two steps: first, to identify common predictors across the top models, and second, to determine algorithm-specific top features unique to each model. Three training scenarios were then applied to the selected models: (i) all 21 preprocessed variables, (ii) only the identified common features, and (iii) only the top features for each algorithm. This methodology aims to help dental practitioners by streamlining data collection while preserving predictive performance.

4 Results

Data analysis was carried out in Python 3.12.7 using the following libraries: `NumPy` for mathematical operations and array manipulation, `Pandas` for data processing and analysis, `Matplotlib` for graphical visualisation, and `Scikit-learn` for implementing machine learning models and evaluation metrics. The code for the practical section is available on a GitHub repository - (project in Portuguese available at: https://github.com/joanalopes7/Dissertacao-de-Mestrado-Doenca-Periodontal-e-Machine-Learning.git)

The exploratory data analysis reveals right-skewed distributions in key clinical variables, including `CALtotal`, with severe outliers at 7.42 mm and 9.73 mm. These outliers are consistent with Stage IV periodontitis. Inflammatory mediators also showed skewness but were retained due to their clinical relevance. Mediators with more than 30% of values below the detection limit (BDL) were excluded. The remaining mediators (IP-10, IL-1, MCP-1, IL-6, IL-8) were imputed using LOD/2 and log-transformed (`log1p`) to address skewness.

We use Mann-Whitney U tests to evaluate differences between two groups. There are statistically significant differences revealed between the Periodontal Health and Disease groups for all mediators ($p < 0.05$). IL-1β, IL-6, and IL-8 had the lowest p-values. After imputation, correlations preserved expected (Fig. 2) biological relationships, such as the IL-1β/IL-8 correlation ($r = 0.68$), supporting data quality. These results show the discriminative potential of the

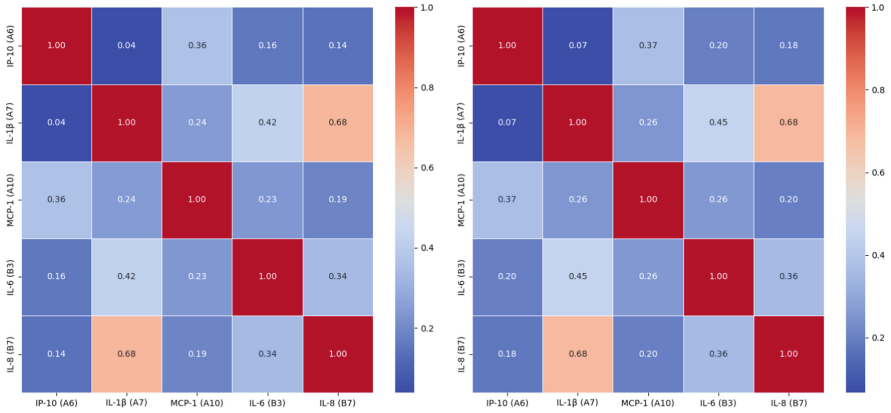

Fig. 2. *Heatmaps* showing the correlation between inflammatory mediators before (left) and after (right) the imputation of BDL values.

selected mediators as robust predictors for periodontal condition classification models.

Subsequent assessment indicated that tree-based models (DT and RF), SVM, and LR achieved high accuracy and area under the ROC curve, indicating strong

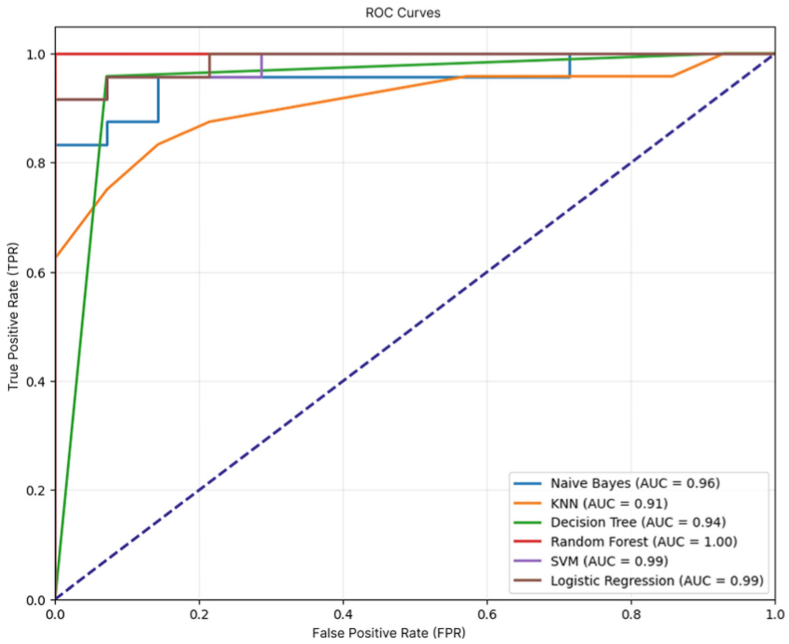

Fig. 3. ROC curves of the six classification models evaluated in this study.

performance in identifying periodontal disease. The Naive Bayes model had high precision, meaning its positive predictions were usually correct, but lower recall, so it missed some cases. The KNN model performed slightly worse, possibly because it is affected by the number of neighbours used and by the data itself.

These results are illustrated by ROC curves in Fig. 3.

Considering the observed performance and ROC curves, the models chosen for the subsequent stage of the methodology adopted in this study were Random Forest, Support Vector Machine, and Logistic Regression. Among these three evaluated models (Table 1), Logistic Regression demonstrated a strong balance between predictive performance (Accuracy=0.974, F1-score=0.979, AUC=1.000) and interpretability. Although ensemble models achieved comparable or slightly higher accuracy, Logistic Regression provided transparent coefficient estimates that facilitated clinical interpretation (Table 2).

Feature importance analysis (Table 3) revealed that periodontal clinical indicators, such as probing depth and plaque index, were among the strongest predictors.

Table 1. Model performance across different feature sets.

Model	Feature Set	Accuracy	Precision	Recall	F1-score	AUC
Random Forest	All Features	0.947	1.0	0.917	0.957	1.000
Random Forest	Shared Features	0.947	1.0	0.917	0.957	0.997
Random Forest	Top Specific Features	0.947	1.0	0.917	0.957	1.000
Logistic Regression	All Features	0.947	1.0	0.917	0.957	0.988
Logistic Regression	Shared Features	0.974	1.0	0.958	0.979	0.988
Logistic Regression	**Top Specific Features**	**0.974**	**1.0**	**0.958**	**0.979**	**1.000**
Support Vector Machine	All Features	0.947	1.0	0.917	0.957	0.985
Support Vector Machine	Shared Features	0.947	1.0	0.917	0.957	0.991
Support Vector Machine	Top Specific Features	0.974	1.0	0.958	0.979	1.000

Table 2. Coefficients, Odds Ratios, and clinical interpretation of the Logistic Regression model with top specific features.

Variable	Coefficient (β)	OR = e^{β}	Clinical Interpretation
BOP_total_percentage	7.4197	1670.7	Each 1% increase in BOP multiplies the odds by ∼1670x
PPD_total	1.4394	4.22	Each additional mm in probing depth increases the odds by 322%
PI_total_percentage	0.8359	2.31	Each 1% increase in PI approximately doubles the odds
log_IL-8 (B7)	0.1595	1.17	Each unit increase in IL-8 increases the odds by 17%
Cigarettes_per_day	0.1028	1.11	Each additional cigarette increases the odds by 11%
log_IL-1β (A7)	-0.4564	0.63	Each unit increase in IL-1β decreases the odds by 37%
No. teeth	-0.9929	0.37	Each additional tooth decreases the odds by 63%
log_IL-6 (B3)	0.2323	1.26	Each unit increase in IL-6 increases the odds by 26%
Brushing_frequency	0.5141	1.67	Each additional brushing increases the odds by 67%
log_MCP-1 (A10)	0.1876	1.21	Each unit increase in MCP-1 increases the odds by 21%

Table 3. Description of variables of selected model - Logistic Regression with Top Specific Features.

Variable	Description	Unit
BOP_total_percent.	Total percentage of bleeding on probing, indicating active gingival inflammation.	%
PPD_total	Total probing pocket depth, reflecting the severity of periodontal tissue destruction.	mm
PI_total_percentage	Total plaque index percentage, associated with oral hygiene control.	%
log_IL-8 (B7)	Log-transformed interleukin-8 level, an inflammatory biomarker related to the local immune response.	pg/mL
Cigarettes_per_day	Number of cigarettes smoked per day, a recognized risk factor for periodontitis.	cigarettes/day
log_IL-1β (A7)	Log-transformed interleukin-1β level, a pro-inflammatory marker strongly associated with tissue destruction.	pg/mL
No._teeth	Total number of teeth present, reflecting the patient's oral health history.	count
log_IL-6 (B3)	Log-transformed interleukin-6 level, an inflammatory biomarker associated with chronic processes.	pg/mL
Brushing_frequency	Daily toothbrushing frequency, a behavioral variable related to oral hygiene.	times/day
log_MCP-1 (A10)	Log-transformed monocyte chemoattractant protein-1 (MCP-1) level, involved in inflammatory cell recruitment.	pg/mL

Smoking status and brushing frequency strongly predicted disease. IL-1β, IL-6, and IL-8 also played key roles.

By combining traditional clinical factors and specific laboratory measurements, our model draws on diverse sources of information, thereby enhancing its predictive ability and providing unique insights into risk mechanisms not available through standard approaches.

5 Discussion

As mentioned before, the main objective of this study was to explore the potential of machine learning models to predict PD using clinical, sociodemographic, behavioural and molecular data. Throughout the entire workflow developed—from data collection and pre-processing to algorithm training and model evaluation—it was possible to identify the key variables associated with the risk of developing PD and to assess the discriminatory power of different

predictive approaches. The discussion that follows seeks to relate the results obtained to those in the existing literature.

Our results showed that total Bleeding on Probing (BOP) was among the strongest predictors of PD. This finding is consistent with previous works that identifies BOP as a fundamental clinical indicator of gingival inflammation and disease progression [6,13,18]. Within the group of diseases that includes gingivitis and periodontitis—both of which fall under the general concept of periodontal disease—bleeding is known to be an unequivocal clinical sign: in gingivitis, this may manifest as bleeding whilst brushing the teeth, whereas in periodontitis it may occur spontaneously for no apparent reason. In this sense, bleeding on probing (BOP) reflects active inflammation and constitutes an objectively reproducible measure that is highly informative for risk assessment.

Statistical analyses conducted in this work show the value of inflammatory mediators in PD diagnosis. All salivary inflammatory markers used to train the model showed statistically significant differences between the 'Periodontal Health' and 'Periodontal Disease' groups. IL-1β, IL-6, and IL-8 have consistently emerged as relevant biomarkers in the three main models, as indicated in previous studies which associate these mediators with periodontal inflammation and tissue degradation [23]. Consequently, clinicians can regularly use saliva-based diagnostics for the early detection, monitoring, prognosis and management of PD, through the detection of host-derived biomarkers. It should always be kept in mind that the diagnostic validity of these tests depends on their sensitivity and specificity [7,23]. Our findings reinforce the importance of certain inflammatory mediators as clinical indicators and potential therapeutic targets in PD.

Furthermore, the analysis of the significance of the variables revealed a wide range of influential factors, with particular emphasis on classic periodontal parameters described in the literature, such as the bacterial plaque index, probing depth and clinical attachment loss, as well as systemic and behavioural factors, such as diabetes and a history of smoking. This variety of significant predictors reinforces the multifactorial nature of PD and the need to consider multiple risk factors that interact with each other to obtain a robust assessment of the risk and severity of the disease [21,26]. Local, systemic and behavioural factors do not operate in an isolated manner; instead, they interact to shape the host's response to the oral biofilm and, ultimately, the clinical manifestation of the disease.

All the machine learning models tested in this study performed well, with an accuracy of over 84% in all cases, and several algorithms—including Random Forest, Decision Tree, Support Vector Machine, and Logistic Regression—achieved an accuracy of over 94%. This high performance reinforces the use of computational approaches for the early identification and prediction of PD and suggests that these models can assist clinicians with risk assessment and clinical decision-making. Although various performance measures (precision, accuracy, recall and F1-score) were calculated, the main metric used to compare the classifiers was the area under the ROC curve. The AUC was preferred because it captures overall discriminatory power independent of the choice of decision threshold and is

less vulnerable to potential class imbalance than precision. In clinical applications, the AUC is particularly suitable because it reflects the trade-off between sensitivity and specificity, which is important for minimising both false negatives and false positives in the diagnosis of PD.

With regard to the chosen model, based on the results presented, the final model selected for clinical application was a logistic regression classifier trained using a subset of the best-performing features. This choice was based not only on its near-perfect performance (AUC = 1.0, with high accuracy, precision and F-score), but also on its interpretability and relevance to clinical practice. Compared to more complex methods, such as Random Forest or Support Vector Machines, Logistic Regression provides coefficients that can be easily interpreted as odds ratios (ORs), allowing clinicians to understand the relative contribution of each variable to the risk of disease. This form of clarity is a major advantage in medical contexts, where model explanations are needed to support clinical judgement and build confidence. The analysis of the coefficients of the Logistic Regression model, together with its principal characteristics, clarified which variables contributed most strongly to the risk of PD. Specifically, the total percentage of BOP, total probing depth and total plaque index had very high ORs (approximately 1670, 4.22 and 2.31, respectively), indicating that even small increases in these clinical parameters substantially increase the likelihood of PD. On the LR model, probability estimates close to 1 indicate a high risk of periodontal disease, whilst those close to 0 indicate a low risk. It should also be pointed out that the coefficients represent adjusted effects, showing the combined influence of all variables in the model. Thus, variables with low multicollinearity, such as smoking or toothbrushing frequency, can be interpreted more directly. In contrast, inflammatory markers or correlated clinical variables should be interpreted within the context of the full multivariate model.

The LR model shows a positive coefficient for daily brushing frequency. This does not mean that more frequent brushing increases the risk of periodontal disease; rather, it reflects the multivariate adjustment. One possible explanation is that patients with poorer periodontal health (higher values for bleeding on probing, probing depth or plaque index) may brush their teeth more frequently in response to symptoms or medical advice, creating a positive association in the data. Clinically, tooth brushing remains protective—the positive coefficient indicates a statistical association after adjustment for other variables. This model should be used as a tool for jointly predicting PD risk, not for evaluating the effects of the coefficients separately.

The final proposed model remains robust across various decision thresholds. Performance analysis revealed high sensitivity and specificity, enabling the reliable detection of cases of periodontal disease without a significant increase in the false positive rate. Its high discriminatory power shows that the predicted probability distributions clearly separate the ‘Periodontal Health’ and ‘Periodontal Disease’ groups, reducing ambiguity in borderline cases.

6 Limitations and Future Work

Several limitations should be addressed. The sample size was relatively small. Thus, the models were not tested on an independent dataset. Subsequent research should evaluate these models using larger, more varied populations. Additionally, employing longitudinal data that tracks the same subjects over time may improve the prediction of disease progression.

7 Conclusion

In conclusion, we can say that this study demonstrates that interpretable machine learning models are capable of accurately predicting periodontal disease using clinical, behavioural and inflammatory data. The logistic regression model achieved high predictive accuracy whilst maintaining clinical practicality, demonstrating its suitability as a decision-support tool in periodontal care. This concise model, which includes ten variables (BOP_total_percentage, PPD_total, PI_total_percentage, log_IL-8 (B7), Cigarettes_per_day, log_IL-1β (A7), No. teeth, log_IL-6 (B3), Brushing_frequency, log_MCP-1 (A10)) confirmed established risk factors and highlighted the importance of salivary cytokines as non-invasive and quantifiable biomarkers of disease activity. These results support the practical utility of logistic regression for personalised risk assessment and management in periodontal clinical practice.

Acknowledgments. This work has been supported by FCT - Fundação para a Ciência e Tecnologia within the R&D Units Project Scope: UID/00319/2023.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alqahtani, H.M., Koroukian, S.M., Stange, K., Schiltz, N.K., Bissada, N.F.: Identifying factors associated with periodontal disease using machine learning. Journal of International Society of Preventive and Community Dentistry **12**, 612–620 (11 2022). https://doi.org/10.4103/jispcd.JISPCD_188_22
2. Ashwood, P., Krakowiak, P., Hertz-Picciotto, I., Hansen, R., Pessah, I., de Water, J.V.: Elevated plasma cytokines in autism spectrum disorders provide evidence of immune dysfunction and are associated with impaired behavioral outcome. Brain Behav. Immun. **25**(1), 40–45 (2011). https://doi.org/10.1016/j.bbi.2010.08.003
3. BinShabaib, M., AlHarthi, S.S., Akram, Z., Khan, J., Rahman, I., Romanos, G.E., Javed, F.: Clinical periodontal status and gingival crevicular fluid cytokine profile among cigarette-smokers, electronic-cigarette users and never-smokers. Arch. Oral Biol. **102**, 212–217 (2019). https://doi.org/10.1016/j.archoralbio.2019.05.001
4. Binson, V.A., Thomas, S., Subramoniam, M., Arun, J., Naveen, S., Madhu, S.: A review of machine learning algorithms for biomedical applications. Ann. Biomed. Eng. **52**(5), 1159–1183 (2024). https://doi.org/10.1007/s10439-024-03459-3

5. BioLegend: Legendplex^{TM} mull-analyte flow assay kit for accurate quantification of multiple human th (t helper cell) cytokines from cell culture supernatant, serum, plasma and other biological samples legendplex^{TM} mull-analyte flow assay kit. Tech. rep. (2025), https://sl1nk.com/ohloU
6. Bishop, C.: Time to take gum disease seriously: The societal and economic impact of periodontitis. Economist Intelligence Unit report, London (2021), commissioned by the European Federation of Periodontology
7. Cafiero, C., Spagnuolo, G., Marenzi, G., Martuscelli, R., Colamaio, M., Leuci, S.: Predictive periodontitis: The most promising salivary biomarkers for early diagnosis of periodontitis (4 2021). https://doi.org/10.3390/jcm10071488
8. Chang, M.C., Lin, S.I., Pan, Y.H., Lin, L.D., Wang, Y.L., Yeung, S.Y., Chang, H.H., Jeng, J.H.: Il-1-induced icam-1 and il-8 expression/secretion of dental pulp cells is differentially regulated by irak and p38. J. Formos. Med. Assoc. **118**(8), 1247–1254 (2019). https://doi.org/10.1016/j.jfma.2018.11.015
9. Chapple, I.L.C., Avila-Ortiz, G., Nibali, L., Graziani, F., Sanz, M., Tonetti, M., Giannobile, W.V., et al.: Periodontal health and gingival diseases and conditions on an intact and a reduced periodontium: Consensus report of workgroup 1 of the 2017 world workshop on the classification of periodontal and peri-implant diseases and conditions. In: Journal of Clinical Periodontology. vol. 45, pp. S68–S77 (06 2018https://doi.org/10.1111/jcpe.12940
10. Ciobanu-Caraus, O., Aicher, A., Kernbach, J.M., Regli, L., Serra, C., Staartjes, V.E.: A critical moment in machine learning in medicine: On reproducible and interpretable learning. Acta Neurochir. **166**(1), 65–78 (2024). https://doi.org/10.1007/s00701-024-05892-8
11. Deng, Y., et al.: Artificial intelligence and machine learning for periodontal diagnosis and prognosis. (Complete with the exact journal details you are using) (2023)
12. of Dental, N.I., Research, C.: Probing periodontal disease. https://www.nidcr.nih.gov/news-events/nidcr-news/2024/probing-periodontal-disease (03 2024), accessed 21 Oct. 2024
13. European Federation of Periodontology: Gum diseases explained. https://www.efp.org/for-patients/gum-diseases/gum-diseases-explained/ (2021), acedido em: 19 out. 2024
14. Hasan, A., Palmer, R.M.: A clinical guide to periodontology: Pathology of periodontal disease. British Dental Journal **216**(8), 457–461 (04 2014https://doi.org/10.1038/sj.bdj.2014.299
15. Li, W., Xia, Y., Yang, J., Sanyal, A.J., Shah, V.H., Chalasani, N.P., Yu, Q.: Disrupted balance between pro-inflammatory lipid mediators and anti-inflammatory specialized pro-resolving mediators is linked to hyperinflammation in patients with alcoholic hepatitis. Frontiers in Immunology **15**, 1377236 (2024). https://doi.org/10.3389/fimmu.2024.1377236
16. Majeed, A.M., et al.: Salivary cytokines as biomarkers for periodontal disease severity: A systematic review. (Fill in exact journal details as per your original reference) (2022)
17. Manzardo, A.M., Poje, A.B., Penick, E.C., Butler, M.G.: Multiplex immunoassay of plasma cytokine levels in men with alcoholism and the relationship to psychiatric assessments. Int. J. Mol. Sci. **17**(4), 472 (2016). https://doi.org/10.3390/ijms17040472
18. Marotz, C., et al.: Machine learning for oral microbiome and periodontal disease prediction. (Complete with journal, volume, pages, doi) (2022)
19. Organization, W.H.: Oral health. https://www.who.int/news-room/fact-sheets/detail/oral-health (03 2023), accessed 10 Nov. 2024

20. Özden, F., et al.: Application of machine learning methods in periodontal disease classification. (Complete with journal details according to your original source) (2015)
21. Papapanou, P.N., Sanz, M., Buduneli, N., Dietrich, T., Feres, M., Fine, D.H., Flemmig, T.F., Garcia, R., Giannobile, W.V., Graziani, F., Greenwell, H., Herrera, D., Kao, R.T., Kebschull, M., Kinane, D.F., Kirkwood, K.L., Kocher, T., Kornman, K.S., Kumar, P.S., Loos, B.G., Machtei, E., Meng, H., Mombelli, A., Needleman, I., Offenbacher, S., Seymour, G.J., Teles, R., Tonetti, M.S.: Periodontitis: Consensus report of workgroup 2 of the 2017 world workshop on the classification of periodontal and peri-implant diseases and conditions. In: Journal of Clinical Periodontology. vol. 45, pp. S162–S170 (06 2018).https://doi.org/10.1111/jcpe.12946
22. Polinski, K.J., Bemis, E.A., Yang, F., Crume, T., Demoruelle, M.K., Feser, M., Seifert, J., O'Dell, J.R., Mikuls, T.R., Weisman, M.H., Gregersen, P.K., Keating, R.M., Buckner, J., Reisdorph, N., Deane, K.D., Clare-Salzler, M., Holers, V.M., Norris, J.M.: Association of lipid mediators with development of future incident inflammatory arthritis in an anti–citrullinated protein antibody–positive population. Arthritis and Rheumatology **73**(6), 955–962 (2021). https://doi.org/10.1002/art.41631
23. Relvas, M., Mendes-Frias, A., Gonçalves, M., Salazar, F., López-Jarana, P., Silvestre, R., da Costa, A.V.: Salivary il-1, il-6, and il-10 are key biomarkers of periodontitis severity. International Journal of Molecular Sciences **25** (8 2024https://doi.org/10.3390/ijms25158401
24. Sarker, I.H.: Machine learning algorithms, real-world applications and research directions. SN Computer Science **2**(3), 160 (2021). https://doi.org/10.1007/s42979-021-00592-x
25. Schober, P., Vetter, T.R.: Logistic regression in medical research. Anesthesia & Analgesia **132**(2), 365–366 (2021). https://doi.org/10.1213/ANE.0000000000005247
26. Steffens, J.P., Marcantonio, R.A.C.: Classificação das doenças e condições periodontais e peri-implantares 2018: guia prático e pontos-chave. Revista de Odontologia da UNESP **47**, 189–197 (8 2018).https://doi.org/10.1590/1807-2577.04704
27. Uddin, S., Khan, A., Hossain, M.E., Moni, M.A.: Comparing different supervised machine learning algorithms for disease prediction. BMC Med. Inform. Decis. Mak. **19**(1), 281 (2019). https://doi.org/10.1186/s12911-019-1004-8
28. Zhang, M., Zhang, X., Wang, K., Tang, X., Sun, J., Zhang, Y.: An update on periodontal inflammation and bone loss. Frontiers in Immunology **15**, 1406798 (2024https://doi.org/10.3389/fimmu.2024.1406798

Modeling Error Propagation Quantization In Neural Networks

Jasur Kiyamov[1,2(✉)], Aleksandr Dik[1], and Egor Savkov[1]

[1] St. Petersburg University, Saint Petersburg, Russia
z.kiyamov@spbu.ru, a.dik@spbu.ru
[2] Samarkand Branch of Tashkent University of Information Technologies, Samarkand, Uzbekistan

Abstract. This paper tackles the issue of quantization error propagation analysis in computational graphs of neural networks. To estimate the contribution of different operators to the resulting error, we suggest an analytical model based on Jacobians. Error accumulation along the operator chain is examined and an experimental comparison between the analytical and empirical errors is carried out. It is demonstrated that the analytical model accurately predicts the direction and magnitude of the error; however, substantial error accumulation is seen when the model passes through nonlinear and MLP blocks. We are able to localize important model components and optimize quantization schemes thanks to the obtained results.

Keywords: Error Propagation · Linear Approximation · Taylor Expansion · Quantization Error · Post-Training Quantization

1 Introduction

In computer vision, natural language processing, and other domains, modern neural networks are extensively employed. However, computational resources restrict their use in real-world systems, which makes quantization techniques important [1,2]. Error results from quantization's reduction of parameter and activation representation precision. The mechanisms by which quantization error spreads throughout the network are still not well understood, despite its extensive use [3]. Nonlinearities distort the structure, linear layers amplify the error as it spreads through the product of Jacobians, and normalization partially stabilizes but does not stop the error from building up.

2 Statement of the Problem

A neural network is considered as a composition of sequentially applied operators [4–6]:

$$X_{k+1} = F_k(X_k), \quad k = 0, \ldots, n-1$$

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 133–143, 2027.
https://doi.org/10.1007/978-3-032-30518-3_9

where $X_k \in \mathbb{R}^{d_k}$ is the input at the k-th step, F_k is the computational graph operator (linear transformation, normalization, nonlinearity, etc.). Quantization is represented as adding a perturbation to the inputs and parameters:

$$X_k^q = X_k + \Delta X_k, \quad W_k^q = W_k + \Delta W_k$$

where ΔX_k and ΔW_k are the quantization errors that arise due to discretization. The output of the quantized model is defined as:

$$X_{k+1}^q = F_k(X_k^q, W_k^q)$$

The error occurring at the (k)-th step is defined as:

$$\Delta X_{k+1} = X_{k+1} - X_{k+1}^q = F_k(X_k, W_k) - F_k(X_k^q, W_k^q)$$

Thus, the error depends on both the input disturbance and the parameter errors. The error propagates along the chain of operators, forming a sequence:

$$\Delta X_0 \to \Delta X_1 \to \cdots \to \Delta X_n$$

Of primary interest is the analysis of dynamics:

- accumulation of errors
- changes in its structure
- amplification or attenuation at different stages

The following tasks are solved within the framework of this work [7–9]:

1. Error estimation at each stage of calculations

$$\Delta X_k \quad \text{for all } k$$

2. Comparison of analytical and empirical error models
   - analytical: through Jacobians and linear approximation
   - empirical: direct calculation of the difference in outputs
3. Analysis of error propagation along a chain of operators
   - nature of accumulation
   - presence of critical areas
4. Identifying the most sensitive operators
   - linear transformations (MatMul)
   - nonlinearities (ReLU, SiLU, Softmax)
   - normalization (LayerNorm, RMSNorm)
5. Assessing the influence of model structure on quantization resistance It is required to build a model that:
   - describes the error propagation

$$\Delta X_{k+1} = \Phi_k(\Delta X_k, \Delta W_k)$$

– provides an exact approximation:

$$\Delta X_k^{analytical} \approx \Delta X_k^{empirical}$$

– allows you to quantify the contribution of each operator to the final error:

$$|\Delta X_n|$$

As a result of solving the problem, it is necessary to obtain an analytical model of error propagation

3 Linear Approximation (1st Order)

The first-order *linear* approximation is a fundamental tool for analyzing how small perturbations in the input propagate through a function. Given a differentiable function F(x), its behavior around a point x can be approximated using the first-order Taylor expansion:

$$F(x + \Delta x) \approx F(x) + J_F(x), \Delta x$$

where $J_F(x)$ is the Jacobian matrix of F at point x. This approximation assumes that higher-order terms are negligible, which is valid when the perturbation Δx is sufficiently small [10,11]. In the context of neural networks, the linear approximation provides an efficient way to estimate the effect of input noise, quantization errors, or parameter perturbations on the output. It allows modeling error propagation layer by layer, where each transformation contributes according to its local Jacobian. However, this approach ignores nonlinear interactions between perturbations, which may lead to significant inaccuracies in highly nonlinear operations such as normalization layers and activation functions.

Let F: $\mathbb{R}^n \to \mathbb{R}^m$ Taylor expansion:

$$F(X + \Delta X) = F(X) + J_F(X)\Delta X + O(\|\Delta X\|^2)$$

where

$$J_F(X) = \frac{\partial F}{\partial X}$$

Error:

$$\Delta Y = F(X + \Delta X) - F(X)$$

Linear model:

$$\Delta Y \approx J_F(X)\Delta X$$

Norm assessment:

$$\|\Delta Y\| \leq \|J_F(X)\| \cdot \|\Delta X\|$$

The Jacobian determines the sensitivity of the operator

3.1 Second-Order Approximation

The second-order approximation extends the linear model by incorporating curvature information of the function. For a twice-differentiable function $F(x)$, the second-order Taylor expansion around point x is given by [12,13]:

$$F(x + \Delta x) \approx F(x) + J_F(x), \Delta x + \frac{1}{2}, \Delta x^T H_F(x), \Delta x$$

where $J_F(x)$ is the Jacobian and $H_F(x)$ is the Hessian matrix of second derivatives. The additional quadratic term captures interactions between perturbations and reflects the local curvature of the function. In neural networks, this approximation provides a more accurate model of error propagation, especially for nonlinear components such as normalization layers (e.g., LayerNorm, RMSNorm) and activation functions (e.g., SiLU, Softmax). Unlike the first-order approximation, which assumes purely linear behavior, the second-order model accounts for how errors interact and amplify due to nonlinearity [14–16].

To improve accuracy:

$$F(X + \Delta X) = F(X) + J_F(X)\Delta X + \frac{1}{2}\Delta X^T H_F(X)\Delta X$$

where

$$H_F(X) = \frac{\partial^2 F}{\partial X^2}$$

Then:

$$\Delta Y \approx J_F(X)\Delta X + \frac{1}{2}\Delta X^T H_F(X)\Delta X$$

important for:

- softmax
- SiLU
- normalization

Propagation of error along the chain For the model:

$$X_{k+1} = F_k(X_k)$$

the error is spreading:

$$\Delta X_{k+1} \approx J_{F_k}(X_k)\Delta X_k$$

Iteratively:

$$\Delta X_n \approx \left(\prod_{k=0}^{n-1} J_{F_k}(X_k)\right) \Delta$$

Growth Rating:

$$\|\Delta X_n\| \leq \prod_{k=0}^{n-1} \|J_{F_k}\| \cdot \|\Delta X_0\|$$

possible:

– attenuation (stable)
– exponential growth (unstable)

Taking into account weights (parametric sensitivity) for the operator F(X,W):

$$\Delta Y \approx J_X \Delta X + J_W \Delta W$$

where:

$$J_X = \frac{\partial F}{\partial X}, \quad J_W = \frac{\partial F}{\partial W}$$

For the linear layer:

$$\Delta Y \approx W \Delta X + \Delta W X$$

Quantization model

$$X_q = X + \delta X, \quad W_q = W + \delta W$$

where:

$$\delta X \sim \mathcal{Q}(X) - X$$

It is often assumed:

$$\mathbb{E}[\delta X] = 0, \quad Var(\delta X) = \sigma^2$$

or uniform noise:

$$\delta X \sim \mathcal{U}(-\epsilon, \epsilon)$$

Error covariance model If errors are independent:

$$\mathrm{Cov}(\Delta Y) = J_X \Sigma_X J_X^T + J_W \Sigma_W J_W^T$$

For the linear layer:

$$\mathrm{Cov}(\Delta Y) \approx W \Sigma_{\Delta X} W^T + \Sigma_{\Delta W} R_X$$

where:

$$R_X = \mathbb{E}[X X^T]$$

the error depends on:

– scale structures
– login statistics

4 Description of the Experiment

This experiment aims to analyze how quantization errors propagate through neural network operators and to evaluate the accuracy of first-order and second-order approximations in modeling these errors [17,18]. The study considers perturbations introduced into both inputs and model parameters and compares the estimated error with the true observed deviation (Table 1).

Table 1. Jacobians of Common Operators

Operator	Function $F(x)$	Jacobian J_F
Matrix Multiplication/Linear Layer	$y = Wx + b$	$J = W$
Addition	$y = x + r$	$J = I$
Subtraction	$y = x - r$	$J = I$
Element-wise Multiplication	$y = x \odot m$	$J = \text{diag}(m)$
Element-wise Division	$y = x \odot \frac{1}{d}$	$J = \text{diag}\left(\frac{1}{d}\right)$
Transpose	$y = Px$	$J = P$
Reshape	$y = \text{reshape}(x)$	$J = P$
Softmax	$y = \text{softmax}(x)$	$J = \text{diag}(y) - yy^T$
ReLU	$y = \max(0, x)$	$J = \text{diag}(\mathbf{1}_{x>0})$
SiLU/Swish	$y = x\sigma(x)$	$J = \sigma(x) + x\sigma(x)(1 - \sigma(x))$

1. The input tensor is specified (Fig. 1):

$$X_0$$

2. An error is added to the input and/or parameters:

$$X_k^q = X_k + \Delta X_k, \quad W_k^q = W_k + \Delta W_k$$

where $\Delta X_k, \Delta W_k$ are quantization errors.

For each operator F_k the following are calculated:

$$X_{k+1} = F_k(X_k, W_k)$$
$$X_{k+1}^q = F_k(X_k^q, W_k^q)$$

Calculating the error at each step

$$\Delta X_{k+1} = X_{k+1} - X_{k+1}^q$$

Calculation of metrics

– absolute error:

$$|\Delta X_k|_2$$

– relative error:

$$\frac{|\Delta X_k|_2}{|X_k|_2}$$

– maximum error:

$$\max |\Delta X_k|$$

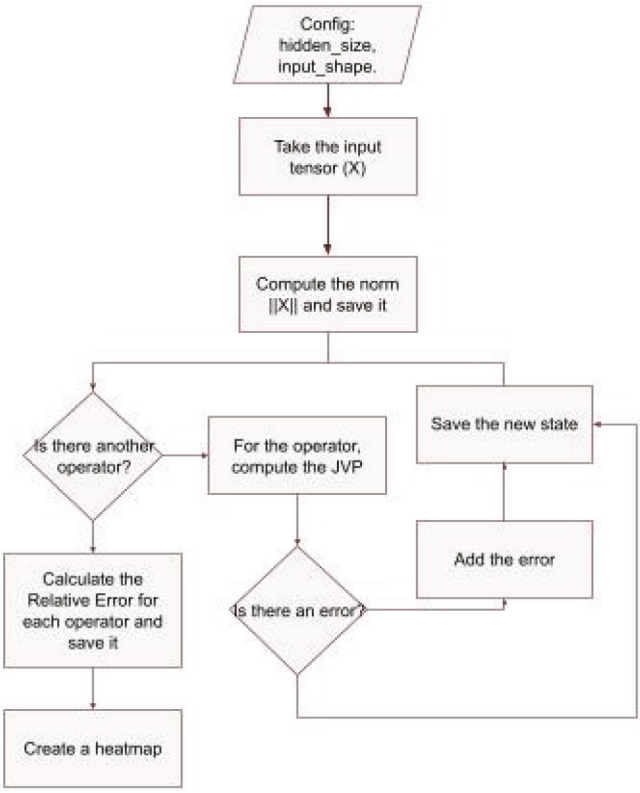

Fig. 1. Error Propagation Modeling.

5 Model Accuracy

In this section, we experimentally validate the suggested error propagation model on a series of operators, such as the Softmax function, a linear transformation (MatMul), nonlinearity (ReLU), and normalization (LayerNorm).

The findings shown in Fig. 2 illustrate how the relative output error varies with the degree of input disturbance and the operator's location along the computational chain. The error is found to increase as it moves through the model, but the type of operator greatly influences the type of increase. The linear transformation (MatMul) contributes most to the increase in error. The input error is scaled by the weight matrix, which explains this:

$$\Delta Y \approx W \Delta X$$

The primary source of error accumulation in the model is linear layers since the error increases proportionately to the weight norm. An extra error transformation related to zeroing out negative components is introduced by the ReLU operator. In contrast to the linear layer, this modifies the error's structure but

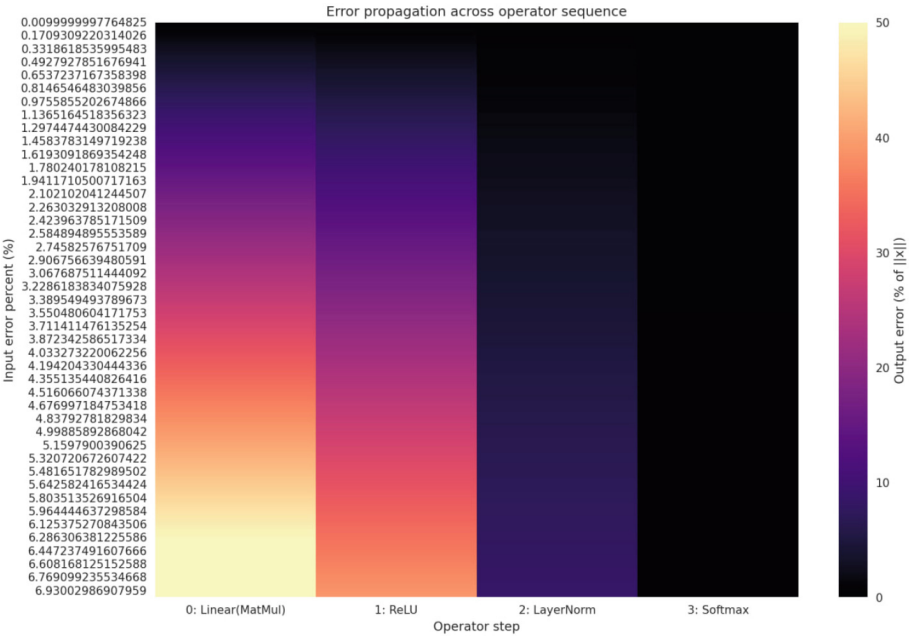

Fig. 2. Accuracy verification for the chain.

does not greatly increase its magnitude. Therefore, nonlinearity primarily affects the error's distribution rather than its scale.

The error is stabilised by LayerNorm, which limits further growth and lessens its spread. The normalisation of input activations is the cause of this:

$$y = \frac{x - \mu}{\sigma}$$

Normalisation reduces sensitivity to input disturbances and equalises the scale of errors. The Softmax function shows the least amount of error growth contribution. Despite being nonlinear, it simply redistributes the error among the output vector's components rather than amplifying it. This demonstrates how resilient Softmax is to minor input perturbations. The following conclusions can be made based on the experimental results:

- The linear transformation (MatMul) is the primary cause of the error's growth;
- Nonlinear functions alter the error's structure but do not significantly amplify it;
- Normalisation prevents errors from spreading;
- Error accumulation does not occur from Softmax.

Normalisation and nonlinear functions serve as stabilisers, but linear transformations play a major role in determining error propagation in the operator

chain. This enables us to pinpoint important model elements and utilise the outcomes to improve quantisation strategies. The obtained values are narrow and concentrated, according to an analysis of the error distribution. The quantisation error exhibits predictable behaviour since the majority of observations are concentrated within a narrow range. Rare but critically large deviations do not occur during the computational process, as indicated by the lack of significant outliers. This indicates that even with quantisation, the model does not display unstable operating modes. The fact that the error distribution is stable and well-controlled attests to the accuracy of the analytical model employed and its suitability for evaluating the effects of quantisation. The scale of the output signal and the quantisation error's magnitude did not significantly correlate, according to the analysis. The relative error does not significantly change when the output vector $|\alpha Y|$'s norm is changed. This suggests that the operator's structural characteristics and the input data's statistics, rather than the signals' absolute magnitude, are what ultimately determine the error. Specifically, the following are important:

- roperties of weight matrices (norm, spectrum)
- type of nonlinearity
- distribution of input activations

The quantization error is structurally determined and is not a simple consequence of the scaling of the output values.

6 Influence of Second-Order Terms

The analysis shows that the accuracy of the error approximation depends significantly on the magnitude of the perturbation. For small quantization errors, the first-order linear model:

$$\Delta Y \approx J_F(X)\Delta X$$

provides high accuracy and adequately describes the error propagation. However, as the magnitude of the disturbance increases, the contribution of the second-order terms becomes significant (Table 2):

$$\Delta Y \approx J_F(X)\Delta X + \frac{1}{2}\Delta X^T H_F(X)\Delta X$$

Ignoring the second order in this case leads to a noticeable discrepancy between the analytical and empirical error. The greatest second-order influence is observed for operators with pronounced nonlinearity and high curvature, such as:

- Softmax, due to exponential dependence
- SiLU (Swish), due to smooth nonlinearity
- LayerNorm, due to dependence on input statistics

Table 2. Modeling Second-Order Errors Across Different Operators and Noise Levels

Operator	Error (%)	True	1st Order (%)	2nd Order (%)	Improvement (%)
LayerNorm					
LayerNorm	0.01	0.000789	0.124596	0.046161	0.078435
LayerNorm	0.10	0.007893	0.061336	0.004509	0.056827
LayerNorm	1.00	0.078921	0.508784	0.004902	0.503882
LayerNorm	2.00	0.157804	1.016194	0.019512	0.996682
LayerNorm	5.00	0.394055	2.540733	0.121827	2.418906
RMSNorm					
RMSNorm	0.01	0.000794	0.057812	0.040447	0.017365
RMSNorm	0.10	0.007941	0.050564	0.003842	0.046722
RMSNorm	1.00	0.079401	0.510653	0.004956	0.505696
RMSNorm	2.00	0.158762	1.021473	0.019730	1.001744
RMSNorm	5.00	0.396443	2.555264	0.123285	2.431979
SiLU					
SiLU	0.01	0.000522	0.069817	0.043226	0.026591
SiLU	0.10	0.005219	0.051649	0.004097	0.047551
SiLU	1.00	0.052193	0.511570	0.002566	0.509004
SiLU	2.00	0.104397	1.022947	0.010072	1.012874
SiLU	5.00	0.261079	2.555468	0.062831	2.492637
Softmax					
Softmax	0.01	1.8×10^{-5}	0.097104	0.057421	0.039683
Softmax	0.10	0.00018	0.066108	0.005894	0.060214
Softmax	1.00	0.001801	0.647290	0.003427	0.643863
Softmax	2.00	0.003603	1.294104	0.013413	1.280691
Softmax	5.00	0.009017	3.231060	0.083720	3.147340

Taking the second order into account is necessary for correct modeling of errors under conditions of significant disturbances.

The study demonstrates that an analytical model based on linear approximation provides highly accurate quantization error prediction. Empirical results confirm that both the magnitude and direction of the error are in good agreement with the analytical estimates. Analysis of the propagation of the error along the chain of operators demonstrates its cumulative nature : as it passes through the model, the error is successively transformed and increases, and the growth occurs unevenly. It was found that the main contribution to the degradation of accuracy is made by MLP blocks, where linear transformations with large weights lead to a significant increase in error. Nonlinear operators (such as ReLU, SiLU, and Softmax) not only change the structure of the error, but in some cases can amplify it due to their high sensitivity to input disturbances.

It is shown that despite the change in scale, the direction of the error remains the same, which is confirmed by the high value of cosine similarity between the analytical and empirical errors.

7 Conclusion

This paper proposes an effective analytical model for studying the propagation of quantization error in neural networks at the level of individual operators and their composition. The model is based on linear approximation using Jacobians and is supplemented with statistical error analysis. It is shown that first-order linear approximation provides high accuracy for small perturbations and agrees well with empirical results. Furthermore, it is established that the quantization error is structured and not random noise, but is determined by the properties of the operators and the distribution of the input data. An analysis of error propagation along the operator chain revealed key nodes where the greatest error amplification occurs. These nodes include attention blocks and MLPs, which play a decisive role in the degradation of model accuracy after quantization.

References

1. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016)
2. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
3. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. Nature **323**(6088), 533–536 (1986)
4. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: ICCV (2015)
5. Hendrycks, D., Gimpel, K.: Gaussian Error Linear Units (GELUs). arXiv:1606.08415 (2016)
6. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv:1607.06450 (2016)
7. Zhang, B., Sennrich, R.: Root mean square layer normalization. In: NeurIPS (2019)
8. Vaswani, A., et al.: Attention is all you need. In: NeurIPS (2017)
9. Higham, N.J.: Accuracy and stability of numerical algorithms. In: SIAM (2002)
10. Banner, R., Nahshan, Y., Soudry, D.: Post training 4-bit quantization of convolutional networks for rapid-deployment. In: NeurIPS (2018)
11. Jacob, B., et al.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: CVPR (2018)
12. Nagel, M., et al.: Data-free quantization through weight equalization and bias correction. In: ICCV (2019)
13. Frantar, E., et al.: GPTQ: accurate post-training quantization for generative pre-trained transformers. arXiv:2210.17323 (2022)
14. Xiao, G., et al.: SmoothQuant: accurate and efficient post-training quantization for large language models. In: ICML (2023)
15. Dettmers, T., et al.: LLM.int8(): 8-bit matrix multiplication for transformers at scale.' arXiv:2208.07339 (2022)
16. Krishnamoorthi, R.: Quantizing deep convolutional networks for efficient inference: a whitepaper. arXiv:1806.08342 (2018)
17. Martens, J.: New insights and perspectives on the natural gradient method. In: JMLR (2020)
18. Papyan, V.: The prevalence of neural collapse during the terminal phase of deep learning training. In: PNAS (2020)

Check for updates

Email Scanner – A Machine Learning Approach to Staffing/Non-Staffing Automatic Classification

Sandra M. Aleixo[1,2(✉)], Cátia Dias[1], Artur Ferreira[1,3], Mariana Lopes[1], Inês Prazeres[4], and Ana Silveira[4]

[1] Department of Mathematics and Department of Computer Science and Engineering, ISEL – Instituto Superior de Engenharia de Lisboa, IPL – Instituto Politécnico de Lisboa, Rua Conselheiro Emídio Navarro, 1, 1959-007 Lisbon, Portugal
{sandra.aleixo,catia.dias,artur.ferreira}@isel.pt, A49804@alunos.isel.pt
[2] CEAUL – Centro de Estatística e Aplicações da Universidade de Lisboa, Faculdade de Ciências da Universidade de Lisboa, Bloco C6, Piso 4, 1749-016 Lisbon, Portugal
[3] IDMEC, IST – Instituto Superior Técnico, Avenida Rovisco Pais, 1049-001 Lisbon, Portugal
[4] Independent Researcher, Lisbon, Portugal

Abstract. A company operates on a consulting basis, and one of the key tasks of the staffing teams is to manage project requests from project managers. This process has three stages: receiving, registering, and managing requests. Reception and registration rely on information sent by the company's Project Managers via email, describing client needs for specific roles. This communication is not standardized and often lacks project details. There is the need to process a vast set of incoming emails, classifying them as staffing if their contents is related to project-request emails, and as non-staffing otherwise. Currently, this classification is made through an empirical rule-based approach with unsatisfactory results. In this paper, we propose a solution to improve the tool for automatic data collection, performing email classification as staffing or non-staffing. Our approach enhances the previous system by improving classification performance using Machine Learning (ML) models while expanding the solution's functional scope. Since data vectorization is required to use ML techniques, several methods for this purpose are evaluated, selecting the most suitable one. Along with the best vectorization, some classifiers are trained and evaluated using appropriate metrics. For the best performing classifiers, their parameters are fine tuned, and the resulting metrics are analyzed, showing that ML models significantly improve the tool's classification performance. This work falls within the scope of industrial mathematics, as it solves a real-world problem through the joint work and collaboration between academia and industry.

Keywords: Artificial Intelligence · Big Data · Classification · Email Analysis · Industrial Mathematics · Machine Learning · Text Classification · Text Vectorization

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 144–162, 2027.
https://doi.org/10.1007/978-3-032-30518-3_10

1 Introduction

A leading multinational company has activity in the fields of management, consulting, information technology, and outsourcing services. The company operates globally in the development and implementation of innovative solutions combining strategy, technology, and digital transformation, supporting organizations across a wide range of industries. The company is on Fortune Global 500 ranking, being present worldwide, employing thousands of people. It partners with many of the world's largest organizations, delivering solutions in artificial intelligence, cloud computing, automation, data analytics, and cybersecurity. It is committed to innovation, sustainability, and inclusion.

In their daily operation, the company needs to deal with many staffing requests. These are sent by email by project managers and describe client needs in terms of specific roles, technical profiles, seniority levels, geographic location, and expected allocation duration. Traditionally, this process involved a substantial manual component, namely reading emails individually, identifying relevant information, and recording it in Excel files with the request date and the requester identifier. The high volume of incoming messages, combined with the variability of email content, makes this procedure time-consuming, prone to human error, and inefficient. This process is essential for planning and allocating human resources to ongoing technology projects.

In the literature, one may find some approaches to the email screening problem, with ML techniques. The EmFore solution [23] is an online learning system developed by Microsoft for automatically learning email folder interpretable symbolic classification rules. The system learns incrementally from user actions, updating rules in real time as new emails arrive. It introduces mechanisms for rule consistency, confidence-based suppression, and fast adaptation. Overall, it enables scalable and transparent email organization for modern email clients. In [9], we have an approach to classify customer support emails, from a telecom company, with word embeddings using the Support Vector Machines (SVM), Naïve Bayes, and Long-Short Term Memory (LSTM) networks. LSTM is found to achieve the best performance. Recently, an email sorting system based on DistilBERT, a lightweight transformer model was proposed [17]. It classifies emails into multiple categories and extracts key information such as deadlines and urgency. The system integrates with Gmail and prioritizes messages using context-aware Natural Language Processing (NLP) techniques.

The use of instruction-following Large Language Models (LLM) for enterprise email classification is approached by [22]. It addresses challenges such as class imbalance, domain-specific language, and limited labeled data. The proposed solution combines instruction-based classification with semantic enrichment techniques. The experimental results show that LLM provide robust contextual understanding across diverse email structures. They also improve handling of previously unseen technical terms and evolving vocabulary.

In [20], the authors address zero-shot spam email classification using pretrained LLM. It evaluates different models on the SpamAssassin dataset, with raw email content and with LLM-generated summaries. Results show strong

performance, with GPT-4 reaching 95% acceptable results without fine-tuning, at the expense of high computational and operational costs.

A comparison between Machine Learning (ML), Deep Learning (DL), and LLM for email classification is reported in [25]. Traditional ML models such as SVM, NB, SVM, and Random Forest (RF) [4] are found efficient requiring fewer resources. The DL models achieve higher accuracy by capturing complex patterns in email data. LLM models further improve performance by understanding context and semantics. The study concludes that hybrid approaches combining ML, DL, and LLM offer the best overall performance. For a recent survey on this topic, please see [1].

The previous email staffing requests classification method used by the company showed insufficient performance for real-world use, with a high rate of incorrect classifications and limited robustness to the linguistic, structural, and semantic diversity of the emails received daily. This limitation compromised the overall reliability of the tool, since the correct identification of staffing emails is a fundamental step in process automation. Incorrect message classification frequently required manual intervention, undermining the initial objective of reducing human effort and increasing operational efficiency. It became evident that the adopted approach to email classification needed to be reformulated by developing a more robust and effective classifier capable of reliably and consistently identifying emails related to staffing requests from those that fall outside this scope.

In this paper, we propose an approach for the unit responsible for the intake, registration, and management of staffing requests, within the company's technology department. A tool based on Artificial Intelligence (AI) [21] and Machine Learning (ML) techniques [2] techniques was developed to partially automate this process. The solution uses AI, namely Large Language Models (LLM) [6] to extract information from staffing emails, which was then stored into a database. This solution enables access to the mailbox, message reading, and extraction of some relevant data, which are then recorded in a database, representing a first step toward information centralization and historical request management. The proposed solution improves the existing one for the binary classification of emails into staffing and non-staffing categories.

The remainder of the paper is organized as follows. Section 2 summarizes the main methodological concepts outlining the materials and methods, considered in our work. Section 3 presents the detailed case study and the experimental evaluation with standard metrics. Section 4 concludes the paper, with the main findings, practical implications, and directions for future work.

2 Methodology

This section presents the methodology adopted for developing the automatic email classification system. It describes the main technical choices considered throughout the work, from text vector representations to the ML classification models employed, as well as the validation and optimization strategies adopted,

and concludes with the evaluation metrics to assess model performance. This structure provides a systematic rationale for the methodological choices that underpin the experimental analysis presented in the following section.

2.1 Vectorization

Text vectorization methods represent textual documents with numerical-valued vectors, suitable for the use of ML algorithms. For automatic text classification and information retrieval, the choice of an adequate vectorization method is a fundamental step, as it determines how textual information is presented to ML models. In the following, we briefly review some of these common representations [15,16].

The Bag-of-Words (BoW) model is one of the simplest techniques. A document is represented as a collection of words or terms, ignoring the order in which they occur while preserving their frequency. A vocabulary containing all distinct terms found in the corpus is then built. Each document is represented by a numerical vector holding the frequency of occurrence of each vocabulary term, yielding a fixed-dimensional vector representation. The conceptual simplicity of Bag-of-Words, combined with ease of implementation and computational efficiency, makes it suitable for applications with large volumes of text. Thus, it is widely used for document classification, clustering, and information retrieval. However, this simplicity entails important shortcomings. By ignoring word order and syntactic structure, the model loses relevant semantic information, being unable to capture relationships between words or distinguish contextual meanings. Moreover, synonyms are treated as distinct terms and lexical ambiguity is not addressed, limiting model performance for some tasks.

Consider a collection $\mathcal{D} = \{d_1, d_2, \ldots, d_n\}$ of n documents and a vocabulary $V = \{t_1, t_2, \ldots, t_p\}$ with p distinct terms. Let $f(t, d)$ denote the number of occurrences of term t in document d. Each document d is represented by a vector $\mathbf{x}_d = (x_{d,1}, \ldots, x_{d,p}) \in \mathbb{R}^p$, where each component $x_{d,i}$ corresponds to the frequency of term t_i in document d, $x_{d,i} = f(t_i, d)$.

The Term Frequency–Inverse Document Frequency (TF-IDF) is one of the most widely used methods for vector representations of text in automatic classification tasks. It is simple, computationally efficient, and interpretable, and is frequently used as a baseline in text classification systems. However, like the BoW model, it ignores word order and context and does not capture semantic relationships between terms. It is a term-weighting scheme whose objective is to quantify the importance of each term in a document while accounting for both its local relevance (within the document) and its global relevance (across the document collection). The underlying idea of TF–IDF is that a term is informative if it occurs frequently in a given document but does not appear ubiquitously across the corpus. Thus, very common words tend to receive lower weights, whereas more specific and discriminative terms become more relevant in the vector representation.

Let $f(t, d)$ denote the number of occurrences of term t in document d, and let $df(t)$ denote the document frequency of term t, i.e., the number of documents

in $\mathcal{D}$ that contain t at least once. The final weight of term t in document d (TF-IDF) is given by the product of the Term Frequency (TF) and Inverse Document Frequency (IDF) components and is defined as $\text{tfidf}(t, d) = \text{tf}(t, d) \cdot \text{idf}(t)$. The TF component measures the importance of the term within the document. A common definition is the relative frequency $\text{tf}(t, d) = f(t, d)/p$. The IDF component penalizes terms that appear in many documents in the corpus, since such terms tend to have lower discriminative power, and is commonly defined as $\text{idf}(t) = \log(n/df(t))$. Thus, each document is represented by a TF-IDF vector $\mathbf{x}_d = (x_{d,1}, \ldots, x_{d,p}) \in \mathbb{R}^p$, where $x_{d,i} = \text{tfidf}(t_i, d)$. This representation is typically sparse, since only a small fraction of the vocabulary terms occurs in each document; the same usually happens with the BoW representation.

To reduce the influence of document length and enable more consistent comparisons between documents, it is common practice to normalize the vector (typically using the Euclidean norm). One direct application of this vector representation is the computation of similarity between documents, for which cosine similarity is the most common measure. In detail, it evaluates the cosine of the angle between the document vectors, quantifying their similarity in terms of content. Regardless of vector magnitude, cosine similarity enables comparison of documents with different lengths, assigning higher values to semantically closer documents and values near zero to weakly related documents.

Word embedding methods provide a recent approach to text vector representations, developed to overcome the main limitations of the previous models. Unlike them, embeddings aim to capture semantic and contextual information by representing each word with a vector. The central idea is that words with similar meanings tend to occur in similar contexts. By analyzing large volumes of text, these models learn vector representations in which geometric proximity reflects semantic and syntactic relationships between words. For example, synonymous or semantically related words are represented by vectors that are close in the embedding space. Despite their greater expressive power, word embeddings still have some limitations in text classification. Simple aggregation of word vectors may lead to loss of structural information and of the relative importance of terms. They are also more complex to train, requiring greater computational resources and, in some cases, large volumes of data to produce high-quality representations.

Each word (term) t is represented by a vector $\mathbf{v}_t \in \mathbb{R}^q$, with $q \ll p$, whose components capture semantic relationships between words. Once word vector representations are obtained, a document representation d can be constructed by aggregating the vectors of the words that compose it. A simple approach is to compute the mean of the word vectors in d, $\mathbf{x}_d = \frac{1}{p} \sum_{i=1}^{p} f(t_i, d)\mathbf{v}_{t_i}$. More details about text vectorization can be found in [15,16].

2.2 Machine-Learning Classifiers

In the context of this work, the problem under study falls within automatic binary text classification being similar to classic tasks such as email SPAM filtering. The objective is to automatically assign a category to each request based

on the content of the message. To this end, Natural Language Processing (NLP) models represent the text in a structured manner and to apply ML techniques. These techniques can be either supervised or unsupervised, depending on the availability of labeled data [2].

In this paper, a supervised learning approach is adopted, since a set of previously human labeled requests is available. The model is trained on these examples, to enable the automatic classification of new requests.

In summary, the main stages of the process adopted in this work for the automatic classification of requests are presented in Fig. 1. Based on this process, different supervised classification models (Näive Bayes, KNN, Logistic Regression, Decision Trees, Support Vector Machines, and Ensemble Methods), learned from several vectorization techniques, were considered in order to compare their performance and select the most effective combination.

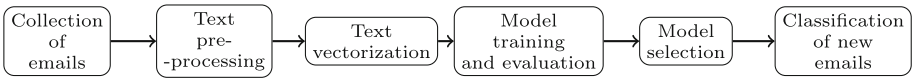

Fig. 1. Stages of the process for the automatic classification of requests.

The system thus follows a model-based approach, in which a classifier is trained on historical data and the resulting model is used to produce predictions for new incoming messages. This approach is appropriate for the problem under study, given that labeled data are available and the objective is the automatic classification of new requests [3,10,15,26].

Naive Bayes (NB) is a probabilistic classifier based on Bayes' theorem [26], and it produces predictions by computing simple and conditional probabilities. This classifier assumes that all predictor variables are conditionally independent of one another given the class label. This hypothesis implies that each variable contributes independently and with equal importance to the final outcome. Bayesian classification aims to estimate the probability that an instance is assigned to a class L (e.g., the category of a text), given a set of features. A decision between two possible classes can be made by computing the ratio between the conditional probabilities associated with each class.

The term *naive* stems from the strong independence assumption between features. Variants of this model include Bernoulli Näive Bayes, Multinomial Näive Bayes, and Complement Näive Bayes. Bernoulli Näive Bayes is particularly suitable for datasets with binary features. In this model, each attribute indicates only the presence or absence of a given feature, taking values in $\{0, 1\}$. Multinomial Näive Bayes is particularly suitable for datasets whose features represent counts or frequencies. In this model, each attribute takes non-negative integer values, corresponding, for example, to the number of occurrences of a word in a document. Complement Näive Bayes is a variant of Multinomial Näive Bayes in which the conditional probabilities of the attributes are estimated from the empirical distribution of the remaining classes, yielding more robust weights, especially

in settings with class imbalance and high dimensionality. Overall, despite their simplicity, Näive Bayes classifiers are often highly efficient, achieving satisfactory levels of accuracy and performance across a variety of classification tasks, particularly in text processing contexts, where they may even outperform more complex classifiers.

The K-Nearest Neighbours (KNN) classifier [26] assumes that similarity between instances is directly related to their proximity. Similar instances (labeled with the class) tend to be close in the feature space, in terms of some distance metric, which is typically computed using the Euclidean norm. The parameter K in the KNN algorithm represents the number of nearest neighbours considered in the decision process. Choosing this parameter is critical, as it defines the region of influence used to classify or estimate a new point. The optimal value of K can be determined empirically or through standard methods such as k-fold cross-validation. The procedure then proceeds with distance computation, identification of the nearest neighbours, and voting for the final class. One of the main advantages of KNN is its simplicity. There is no explicit training phase and no need to tune complex parameters. However, as the number of instances or features increases, the algorithm becomes computationally slower, which makes it less suitable for scenarios requiring fast labeling responses.

Logistic Regression (LR) [10] models are widely used in text classification tasks due to its simplicity, computational efficiency, and strong performance with high-dimensional vector representations such as Bag-of-Words or TF-IDF. This classifier assumes that the decision boundary can be obtained via a linear combination of features, learning a weight vector that maximizes class separation and directly modeling the probability that an instance belongs to a class through the logistic function.

Decision Trees (DT) are classification models based on a tree-like structure [18,19]. Classification consists of traversing the tree from the root to a leaf, following the branches defined by the conditions evaluated at each node. In this way, each root-to-leaf path corresponds to a decision rule used to classify the data. The leaf nodes hold the class labels. For the DT learning process, a function is used to compute the cost associated with each branch, and the branch with the lowest cost is selected. This classifier is sensitive to small variations in the training data and their presence may yield completely different trees. Moreover, its greedy nature does not always guarantee a globally optimal tree and, in the presence of dominant classes, biased trees may emerge.

Support Vector Machines (SVM) [7,24] are among the most popular and well-established classifiers for ML applications. They are versatile models capable of performing linear and nonlinear classification. The main objective of SVM is to identify, in the multidimensional feature space, a hyperplane that correctly separates the data while maximizing the margin, i.e., the distance between the hyperplane and the closest points from each class, called support vectors. Maximizing the margin enables more confident classification of new data, reducing the probability of classification errors. When the data are not linearly separable, SVM can use kernel functions to project the data into a higher-dimensional

space where linear separation becomes possible. Due to its robustness and strong generalization ability, it is frequently used in high-dimensional problems such as text classification [11].

Ensemble methods [27] combine multiple base models to produce a more robust classifier with improved generalization ability. In text classification tasks, approaches such as Bagging, Random Forest, Extremely Randomized Trees, AdaBoost, and Gradient Boosting build several classifiers (typically DT) and aggregate their predictions via voting or weighted combination [4]. These methods can improve performance as compared to individual models and are particularly useful when the goal is to capture complex patterns in textual data.

Bagging [5] reduces variance by averaging models trained on random subsets of the data. No differentiated weights are assigned to individual classifiers, which contribute equally to the final decision. Random Forest and Extremely Randomized Trees rely on an ensemble of trees. Each tree produces a prediction, and the final class is obtained by aggregating individual predictions, typically via majority voting. The main difference between these methods lies in the degree of randomness introduced during tree construction. While Random Forest randomly selects a subset of features and searches for the best split at each node, Extra Trees introduces additional randomness by selecting split points at random for each input feature. Boosting methods build models sequentially, giving greater weight to misclassified examples. AdaBoost (Adaptive Boosting) [8] iteratively combines weak classifiers by adjusting the weights of misclassified observations and training based on errors attained by previous models. In Gradient Boosting, each new model is trained to correct the errors made by the ensemble built so far, explicitly optimizing a loss function via gradient descent and fitting each new model to the residuals of the previous model.

2.3 Evaluation Metrics for Classifier Performance

Evaluating the performance of a ML model is a fundamental step. Evaluation metrics quantify model effectiveness and enable objective comparison of different approaches. Depending on the nature of the problem, selecting an appropriate metric is crucial for model selection [15,26]. In the following, we describe the key evaluation metrics for classifier performance assessment.

The confusion matrix is a fundamental evaluation tool for classifier performance. It compares true classes with predicted classes and forms the basis for computing several evaluation metrics, such as Accuracy, Precision, Recall, and F_1-Score. In binary classification, the confusion matrix has dimensions 2×2 and consists of four fundamental elements: True Positives (TP) - positive cases correctly classified; True Negatives (TN) - negative cases correctly classified; False Positives (FP) - negative cases incorrectly classified as positive, and False Negatives (FN) - positive cases incorrectly classified as negative. Accuracy measures the proportion of correct predictions made by the model among all predictions,

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}.$$

Although it provides a quick overall view of model performance, Accuracy can be misleading on imbalanced datasets, as it may convey a false impression of good performance when the misclassification rate for minority classes is high.

Precision assesses the proportion of positive predictions made by the model that are actually correct,

$$\text{Precision} = \frac{TP}{TP + FP}.$$

This metric is particularly relevant in scenarios where the cost associated with false positives is high. Precision helps ensure that, when the model predicts a positive outcome, that prediction has a high probability of being correct.

Recall, also known as Sensitivity, measures the proportion of actual positive cases that are correctly identified by the model,

$$\text{Recall} = \frac{TP}{TP + FN}.$$

This metric is particularly important in scenarios where failing to detect a positive case (a false negative) has a higher cost than a false positive.

The F_1-Score corresponds to the harmonic mean of Precision and Recall,

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}},$$

and is particularly useful when both must be balanced into a single measure. A high F_1-Score indicates that the model performs well simultaneously in terms of both Precision and Recall.

The F_β-Score is a generalization of the F_1-Score that allows adjusting the relative importance of Precision and Recall through a parameter β,

$$F_\beta = (1 + \beta^2) \cdot \frac{\text{Precision} \cdot \text{Recall}}{\beta^2 \cdot \text{Precision} + \text{Recall}},$$

thereby allowing one metric to be prioritized depending on the objectives of the problem under analysis. In particular, when $\beta > 1$ greater emphasis is placed on Recall, resulting in a stronger penalty for false negatives. Conversely, when $0 < \beta < 1$, greater emphasis is placed on Precision, resulting in a stronger penalty for false positives.

The Receiver Operating Characteristic (ROC) curve is a graphical representation used in binary classification assessment problems that shows the relationship between the True Positive Rate (TPR), which coincides with Recall, and the False Positive Rate (FPR) given by $FPR = FP/(FP + TN)$, for different threshold values. This curve allows analysis of the trade-off between the ability to detect positive instances and the number of false alarms. The Area Under the Curve (AUC) metric corresponds to the area under the ROC curve and quantifies the model's overall ability to distinguish between classes. AUC takes values in the range $[0, 1]$, with higher values indicating better performance. More details about these evaluation metrics can be found in [26].

2.4 Classifier Assessment and Fine Tuning

To ensure a model's ability to generalize to unseen data, validation and optimization strategies are adopted to robustly estimate classifier performance and to select suitable hyperparameter configurations. A hyperparameter is defined as a setting chosen before or during training validation that controls how the model learns or its ability, and is not directly estimated from the data as it happens for model parameters [3,26].

Cross-Validation (CV) is a statistical technique widely used in ML to assess the generalization ability of a predictive model. Its main objective is to estimate model performance on unseen data with different random samples from the input datasets (data folds) mitigating overfitting issues associated with evaluations performed on a single training/test split. In classification, overfitting occurs when a classifier achieves very high performance on the training set, but performance drops sharply on the validation/test set, indicating that the model memorised training-specific patterns/noise and was not able to generalize to new examples.

The most common CV method is k-fold cross-validation, in which the dataset is divided into k subsets or folds. In each iteration, the model is trained using $k-1$ folds and evaluated on the remaining fold, which acts as the test set. This process is repeated k times, ensuring that each subset is used exactly once for testing. Final model performance is obtained by averaging the results across all iterations, providing a more stable and robust estimate of performance. In imbalanced-data settings, i.e., with much more instances from one class than the other, stratified k-fold cross-validation is commonly used to preserve class proportions within each fold.

One of the main advantages of cross-validation is its ability to reduce the variance of the model performance estimate by evaluating classifier behaviour across multiple data partitions. However, it increases computational cost, since the model is trained multiple times, which can be limiting for large datasets or complex models. It is frequently used together with hyperparameter tuning procedures such as Grid Search, enabling systematic evaluation of different parameter configurations based on robust validation.

Grid Search is a systematic approach to hyperparameter optimization in which predefined combinations of values for a set of hyperparameters are exhaustively evaluated. The objective is to identify the configuration that maximizes (or minimizes) a selected performance metric, yielding a model with improved generalization ability. Despite its simplicity and interpretability, Grid Search can become computationally expensive, since the total number of trained models grows with the product of the number of values considered for each hyperparameter and the number of folds.

Stochastic Gradient Descent (SGD) is an iterative optimization method widely used in supervised learning models, particularly in settings with large datasets or high dimensionality. Unlike standard Gradient Descent, which computes the gradient of the cost function using the entire training set, SGD updates model parameters progressively using individual examples or small batches, resulting in a significant reduction in computational cost. The Log Loss function,

combined with SGD, evaluates whether the probabilities predicted by the model match the true classes, penalizing incorrect predictions with high probability.

3 Case Study: Staffing/Non-Staffing Classification

The process of classifying emails as Staffing/Non-Staffing is essential for planning and resource management at the company. A staffing request is defined as a formal request for allocating human resources to a given project or set of projects. These requests are typically made by project managers or team leads, who identify the need to reinforce resources to meet operational requirements and deadlines agreed with the client. Each staffing request specifies a set of requirements associated with the desired profile, is usually communicated via email, and exhibits high variability in structure and level of detail, since it does not follow a rigidly standardized format or taxonomy. This heterogeneity makes manual process management even more complex, particularly in high-volume settings, reinforcing the importance of automated mechanisms that can correctly identify, classify, and handle requests.

3.1 Data Description

The data consists of real corporate emails from an internal mailbox, covering communications recorded since 2023. Emails are written in multiple languages, namely Portuguese, English, Spanish, and, to a lesser extent, French, reflecting the linguistic diversity of the organizational context. To increase the dataset size and improve model generalization, additional emails had to be generated using AI techniques, based on existing emails.

The database initially comprised 1132 real emails and was used only to evaluate the existing empirical classifier. It was subsequently augmented to 1672 real emails and used for vectorization experiments and in an initial classification phase, that produced some interesting performance results. To analyze classifier behaviour on a larger and more balanced dataset, company-approved AI-generated emails were incorporated, yielding a total of 3263 emails. Emails were manually labeled into two classes, staffing and non-staffing, and the labeling was validated by the company. The database includes the following attributes/variables: a unique email identifier (id), the email subject (subject), the email body (body), an indicator identifying whether the email is part of a reply chain (email_chain), and the corresponding label (classification). The subject and body variables constitute the primary sources of textual information used by the classification models.

3.2 Data Pre-processing and Preparation

An initial pre-processing phase was performed on the textual data to reduce noise, standardize text representation, and improve the quality of extracted features. Pre-processing included normalization of the text in the subject and body

fields, encompassing conversion of all characters to lowercase, removal of accents, and elimination of redundant whitespace. This reduces vocabulary size and prevents superficial writing variations from being treated as distinct terms by the models. The subject and body fields were concatenated, resulting in a single textual field per email.

3.3 Pre-existing Classification Rule

The existing tool under development, used to automate email classification, incorporated a model based on a set of empirical rules and keywords, without the use of ML techniques. The decision logic relied exclusively on the presence and frequency of predefined keywords, applying deterministic rules to email textual content. In detail, an email was classified by an empirical rule, as *staffing* if the text did not contain the word "férias" (the Portuguese word for holidays) and contained at least a minimum number of keywords.

The performance of this empirical classifier was evaluated on an initial database consisting of 1132 real emails. The obtained predictions were compared with the true labels, yielding TP=140, TN=522, FP=428 and FN=42, hence Accuracy=0.5848. Precision, Recall, and F_1-Score values are reported in Table 1.

Table 1. Performance of the initial rule-based approach.

	Staff	No-Staff
Precision	0.2465	0.9255
Recall	0.7692	0.5495
F_1-Score	0.3733	0.6896

These results indicate limited overall performance. Although the classifier achieves relatively high Recall for the staffing class, its Precision for that class is very low, indicating that while it identifies most staffing emails, it incorrectly classifies a large number of non-staffing emails as staffing. For the non-staffing class, the opposite behaviour is observed: high Precision but moderate Recall, resulting in several emails that in fact belong to this class being missed. Overall, the classifier shows a clear tendency to overclassify emails as staffing.

This classifier depends entirely on a manually defined keyword list. Small linguistic variations, differences in request phrasing, or the introduction of new terms not covered by the rules frequently lead to incorrect classifications. This rigidity makes system maintenance difficult and compromises its ability to generalize to new data. Therefore, adopting a ML-based approach becomes necessary, as it can exploit statistical and semantic patterns present in the texts.

3.4 Vectorization Experiments

To enable the use of ML classifiers, vector representations are required. Accordingly, the methods presented in Sect. 2.1 were evaluated to assess the impact of text representation on the performance of some of the classifiers presented in Sect. 2.2. The selected classifiers were Logistic Regression (LogReg), SVM (LinearSVC), and Random Forest (RandomForest), as they are widely used in text classification tasks and perform well in high-dimensional spaces. The BoW and TF-IDF vectorization methods were tested with and without stopwords (BoW-SW, BoW, TFIDF-SW, and TFIDF). Stopwords are highly frequent words in a language (e.g., articles, prepositions, and conjunctions) that carry limited semantic content and tend to introduce noise while contributing little relevant information for class discrimination. Word embeddings (WEMB) were tested only without stopwords, since pre-trained embeddings already provide vectors that capture context, and the presence of stopwords may be useful because it helps the model capture grammar and relationships between words. For each vectorization method, stratified five-fold cross-validation was applied, allowing identification of the best mean F_1-Score for the staffing class, as reported in Table 2.

Table 2. Mean F_1-Score for the different classifier/vectorization combinations. The best result is in boldface.

	BoW-SW	BoW	TFIDF-SW	TFIDF	WEMB
LogReg	**0.8572**	**0.8530**	0.8321	0.8257	0.7619
LinearSVC	0.8526	0.8398	**0.8390**	**0.8522**	**0.7770**
RandomForest	0.8278	0.8176	0.8101	0.8168	0.7451

For each vectorization method, the classifier with the highest mean F_1-Score was selected, and Grid Search was applied to optimize its hyperparameters. The selected classifier was then trained on the dataset with 1672 emails using a test split consisting of 75% of the emails, obtaining the metrics reported in Table 3.

Table 3. Metrics for the best classifier/vectorization combination. The best result is in boldface.

	Accuracy	Precision		Recall		F_1-Score	
		Staff	No-Staff	Staff	No-Staff	Staff	No-Staff
LogReg + BoW-SW	0.8373	0.8063	0.8634	0.8324	0.8412	0.8191	0.8522
LogReg + BoW	0.8373	0.8063	0.8634	0.8324	0.8412	0.8191	0.8522
LinearSVC + TFIDF-SW	0.8421	0.7960	**0.8848**	**0.8649**	0.8240	0.8290	0.8533
LinearSVC + TFIDF	**0.8541**	**0.8196**	0.8839	0.8595	**0.8498**	**0.8391**	**0.8665**
WEMB	0.7990	0.7306	0.8744	**0.8649**	0.7468	0.7921	0.8056

In this setting, for BoW the stopword removal does not affect the metrics. For Accuracy and F_1-Score, the highest value is obtained with TF-IDF vectorization without stopwords. However, in the context of this work, evaluation should be primarily driven by Recall for the staffing class, since the system's main objective is to minimize missed staffing requests, i.e., to reduce the number of false negatives. Both TF-IDF with stopwords and word embeddings achieve the highest value for this metric. However, since TF-IDF with stopwords yields higher values for the remaining metrics, it was selected as the vectorization method to be used for automatic classification.

3.5 Classifier Selection

After selecting TF-IDF vectorization with stopwords, several classifiers were trained and evaluated on the final dataset, which comprised 3263 emails: 1737 staffing emails and 1526 non-staffing emails, corresponding to 53.23% and 46.77% of the total, respectively. Thus, the dataset is reasonably balanced and suitable for subsequent stages. The dataset was split into training and test subsets using a stratified train–test split strategy, ensuring preservation of class proportions in both subsets. The adopted split used 75% of the data for training and 25% for testing, and it was performed before any numerical transformation of the text in order to avoid data leakage, i.e., to prevent information from the test set from influencing the training process. If vectorization were performed prior to splitting, the statistics used would be computed on the full dataset, including test emails, giving the classifier indirect access to information that should remain unknown during training.

In parallel, an analysis of the vocabulary most characteristic of emails labeled as staffing was conducted based on mean TF-IDF weights computed exclusively for this class. This analysis showed that the highest-weight terms are strongly associated with the context of resource allocation and project management.

The classifiers presented in Sect. 2.2 were considered, as well as a linear classifier optimized with SGD using Log Loss. Evaluation was performed using five-fold cross-validation, preserving the class distribution in each fold. The following metrics were considered for the staffing class: Accuracy, Precision, Recall, F_1-Score, F_2-Score, and AUC, as reported in Table 4.

For selecting the classifier to be used in the tool, the F_2-Score metric was defined as the priority, since it emphasizes Recall for the staffing class. For the company, it is more important that a staffing email is not classified as non-staffing rather than a non-staffing email is classified as staffing (no false negatives). The results indicate that the SVM classifier achieves the best overall performance, reflecting a strong ability to detect staffing emails, combined with a balanced trade-off with Precision. However, the SGD classifier with Log Loss achieves a very similar value, with slightly higher Accuracy and AUC than SVM.

Based on this analysis, the hyperparameters of both classifiers were optimized using Grid Search to obtain the final evaluation. After applying Grid Search, the F_2-Score values were 0.9315 for SVM and 0.9244 for SGD with Log Loss. The improvement achieved by SVM is larger, increasing the gap between the

Table 4. Mean metrics for the classifiers. The best result is in boldface.

Classifier	Accuracy	Precision	Recall	F_1-Score	F_2-Score	AUC
BernoulliNB	0.8789	0.8759	0.9035	0.8894	0.8978	0.9504
MultinomialNB	0.8896	0.8932	0.9021	0.8976	0.9002	0.9644
KNN	0.8799	0.9084	0.8589	0.8827	0.8681	0.9493
LogReg	0.8730	0.8628	0.9115	0.8863	0.9011	0.9476
DecisionTree	0.7824	0.7668	0.8653	0.8127	0.8434	0.8126
SVM	0.9104	0.9080	0.9280	0.9177	**0.9238**	0.9707
Bagging	0.8792	0.8875	0.8870	0.8870	0.8869	0.9538
RandomForest	0.8774	0.8675	0.9136	0.8899	0.9039	0.9622
ExtRandTrees	0.8874	0.8801	0.9172	0.8981	0.9094	0.9667
AdaBoost	0.8327	0.8074	0.9143	0.8573	0.8905	0.9241
GradientBoosting	0.8490	0.8184	**0.9338**	0.8722	0.9081	0.9434
SGDLogLoss	**0.9116**	**0.9105**	0.9273	**0.9187**	**0.9238**	**0.9744**

two models. An increase in AUC to 0.9740 is also observed, indicating excellent separability between the staffing and non-staffing classes. The ROC curve shape, depicted in Fig. 2, positioned mostly above the diagonal, further supports the conclusion that the model performs substantially better than chance.

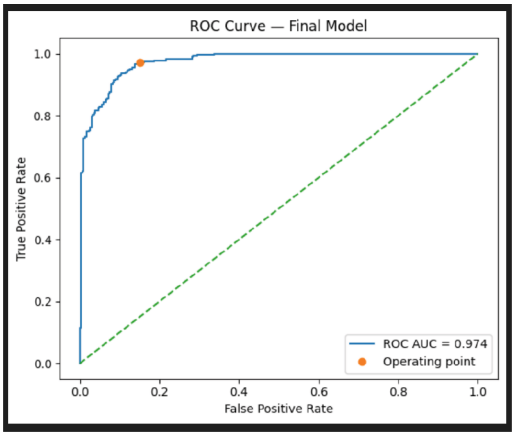

Fig. 2. ROC Curve and AUC.

In addition, the classifier achieves an Accuracy of 0.9142, indicating strong overall classification performance. For the staffing class, Precision and Recall are 0.8802 and 0.9713, respectively, showing effective identification of most true staffing requests, in line with the objective of prioritizing Recall for this class. For

the non-staffing class, Precision and Recall are 0.9628 and 0.8492, respectively. These results confirm the suitability of the classifier for the problem under study, demonstrating consistent discrimination between the two classes.

After selecting the classifier, the decision threshold was further adjusted to prioritize minimizing false negatives for the staffing class. The threshold was selected by maximizing the F_2-Score on a validation set (a subset of the training set) to avoid overfitting. On the test set, a Recall of 0.9770 was obtained for the staffing class. Despite a slight reduction in Precision to 0.8416, the F_2-Score increased to 0.9465, confirming that threshold tuning aligned the classifier with our objectives. Using the adjusted threshold, the classifier achieves an Accuracy of 0.9465, indicating strong overall classification performance. For the staffing class, Precision and Recall are 0.8416 and 0.9770, respectively, and for the non-staffing class, Precision and Recall are 0.9679 and 0.7902, respectively.

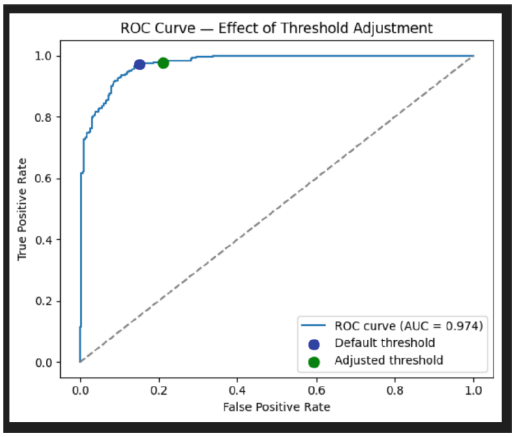

Fig. 3. ROC curve with default and adjusted thresholds.

The ROC curve of the selected classifier in Fig. 3 shows the two operating points corresponding to the default and the adjusted threshold The adjusted threshold shifts the operating point along the curve, increasing TPR at the expense of a slight increase in FPR. This behaviour confirms that threshold optimization allows prioritizing the detection of staffing emails, aligning the model with our objectives without compromising its overall discriminative capability.

4 Conclusions

The company faces a daily challenge within their technology department responsible for managing staffing requests. These requests, sent by project managers via email, include details such as technical profiles, seniority levels, location, and allocation duration. The process is largely manual (emails are read individually

and relevant information is recorded in Excel files). The company faces a big data problem (in both volume and variability) since the current approach is time-consuming, inefficient, and prone to human error. The first attempt of the company to screen out these emails was to devise an empirical classification rule. However, this lead to unsatisfactory results, since it provides a large number of false negatives (true staffing emails wrongly classified as non-staffing).

Our approach aims to provide a new solution for the email staffing/non-staffing classification, grounded on ML techniques. The goal was to improve the classification results for staffing and non-staffing emails, integrated into an AI-based system to support the management of project requests within Staffing teams. The previously developed empirical, rule-based solution was evaluated, leading to the conclusion that improvements could be implemented at the classification level, by devising a different approach from a labeled dataset.

Several vectorization methods were evaluated. TF-IDF with stopwords proved to be the most effective and was subsequently used to train and evaluate different ML classifiers. For the two best-performing models, SVM and SGD with Log Loss, hyperparameter optimization was performed using Grid Search to select the classifier to be integrated into the system. The proposed solution significantly improves the previous results attained by the empirical rule approach.

Overall, this entire study and the experimental evaluation demonstrate that an approach based on Natural Language Processing and Machine Learning techniques enables more robust and consistent email classification, reducing manual intervention and increasing process efficiency. The developed system therefore contributes to more structured information management, facilitating the logging, tracking, and historical analysis of project requests.

As future work, we aim to explore dimensionality reduction with feature selection techniques. The reduction of data dimensionality will identify the relevant, irrelevant, and redundant features of our dataset. With lower dimensionality of the data, we expect to further improve the classification accuracy and further reduce the false negatives, achieving the maximum Recall.

Acknowledgments. This work is funded by national funds through FCT Fundação para a Ciência e a Tecnologia, I.P., under CEAUL Research Unit, UID/00006/2025, DOI: https://doi.org/10.54499/UID/00006/2025, and by the European Union Next GenerationEU through the project UID/PRR/00006/2025, DOI: https://doi.org/10.54 499/UID/PRR/00006/2025 and ISEL. The authors acknowledge also FCT for its financial support via LAETA (project https://doi.org/10.54499/UID/50022/2025).

Disclosure of Interests. The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this article. All analyses were conducted independently, and no external funding sources affected the study design, data collection, or interpretation of results. During the preparation of this manuscript/study, the authors used ChatGPT-5 (OpenAI) for the purposes of improving the clarity and readability of the manuscript. The

authors have reviewed and edited the output and take full responsibility for the content of this publication.

References

1. Alshaikh, M., Alrajeh, Y., Alamri, S., Melhem, S., Abu-Khadrah, A.: Supervised methods of machine learning for email classification: a literature survey. Syst. Sci. Control Eng. **13**(1), 2474450 (2025)
2. Alpaydin, E.: Introduction to Machine Learning, 2nd edn. MIT Press, Cambridge (2010)
3. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer, Hiedelberg (2009)
4. Breiman, L.: Random forests. Mach. Learn. **45**(1), 5–32 (2001)
5. Breiman, L.: Bagging predictors. Mach. Learn. **24**(2), 123–140 (1996)
6. Brown, T., et al., : Language models are few-shot learners. In: Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901 (2020)
7. Cortes, C., Vapnik, V.: Support-vector networks. Mach. Learn. **20**(3), 273–297 (1995)
8. Freund, Y., Schapire, R.: Experiments with a new boosting algorithm. In: Proceedings of the Int. Conference on Machine Learning (ICML), pp. 148–156 (1996)
9. García, V., Sánchez, J.S., Mollineda, R.A.: E-mail classification with machine learning and word embeddings for improved customer support. Neural Comput. Appl. **31**, 3891–3903 (2019)
10. Hastie, T., Tibshirani, R., Friedman, J.H.: The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer, Heidelberg (2017)
11. Hsu, C., Lin, C.: A comparison of methods for multi-class support vector machines. IEEE Trans. Neural Netw. **13**(2), 415–425 (2002)
12. Jebara, T.: Machine Learning: Discriminative and Generative. Springer, Heidelberg (2004)
13. Liu, H., Motoda, H.: Feature Selection for Knowledge Discovery and Data Mining. Kluwer Academic Publishers (1998)
14. Mackay, D.: Information Theory, Inference, and Learning Algorithms. Cambridge University Press, Cambridge (2003)
15. Manning, C.D., Raghavan, P., Schütze, H.: Introduction to Information Retrieval. Cambridge University Press, Cambridge (2009)
16. Patil, R., Boit, S., Gudivada, V., Nandigam, J.: A survey of text representation and embedding techniques in NLP. IEEE Access **11**, 36120–36146 (2023)
17. Priya, S.Y., Reddy, B.I.S., Malleswar, R.U., Kumar, B.V., Jeevana Priya, C.: Smart email sorting: a machine learning approach for enhanced productivity. Int. J. Res. Appl. Sci. Eng. Technol. (IJRASET) (2026). https://doi.org/10.22214/ijraset.2026.78656
18. Quinlan, J.: Induction of decision trees. Mach. Learn. **1**(1), 81–106 (1986)
19. Rokach, L., Maimon, O.: Top-down induction of decision trees classifiers - a survey. IEEE Trans. Syst. Man Cybern. Part C Appl. Rev. **35**(4), 476–487 (2005)
20. Rojas-Galeano, S.: Zero-shot spam email classification using pre-trained large language models. arXiv preprint arXiv:2405.15936 (2024)
21. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach, 3rd edn. Pearson Education, England (2016)

22. Sarıyıldız, A.Ç., Durukan-Odabaşı, Ş: Enterprise e-mail classification using instruction-following large language models. Appl. Sci. **16**(5), 2173 (2026)
23. Singh, M., Cambronero, J., Gulwani, S., Le, V., Verbruggen, G.: EmFore: online learning of email folder classification rules. In: Proceedings of Int. Conference on Information and Knowledge Management (CIKM 2023), pp. 2280–2290 (2023)
24. Vapnik, V.: Statistical Learning Theory. Wiley, New York (1998)
25. Wang, X.: Emails classification: comparing machine learning, deep learning and large language models. In: Proceedings of the International Conference on Data Analysis and Machine Learning (DAML 2024), pp. 451–458 (2024)
26. Witten, I.H., Frank, E., Hall, M.A., Pal, C.J., Foulds, J.: Data mining : practical machine learning tools and techniques. 5th Edition, Morgan Kaufmann (2025)
27. Zhang, C., Ma, Y.: Ensemble Machine Learning: Methods and Applications. Springer, Heidelberg (2012)

Forecast Difficulty Detection via Latent Geometry in N-BEATS

Lauryane Santos Siqueira[1(✉)], Leonardo Nogueira Matos[1], Paulo Novais[2], and Pedro José Costa de Oliveira[2]

[1] Department of Computer Science, Graduate Program in Computer Science (PROCC), Federal University of Sergipe (UFS), Sergipe, Brazil
{lauryanesiqueira,lnmatos}@academico.ufs.br
[2] Department of Informatics, Algoritmi Research Centre, University of Minho, Braga, Portugal
pjon@di.uminho.pt, pedro.jose.oliveira@algoritmi.uminho.pt

Abstract. Identifying difficult forecasts and assessing model reliability remain important challenges in time series forecasting. This work investigates whether geometric properties of the latent space associated with the interpretable coefficients of the N-BEATS model can provide informative signals about predictive difficulty. For this purpose, a metric based on the aggregated magnitude of latent activations is proposed and compared with probabilistic density models (*Gaussian Mixture Models* and *Kernel Density Estimation*) and baselines computed directly from the input signal, such as amplitude and variance. Experiments conducted on the M4 dataset showed that input-space baselines achieved the highest global AUC and Spearman correlation values, indicating a strong relationship between simple statistical properties of the original signal and forecasting difficulty. However, among the latent-space approaches, the proposed geometric metric consistently outperformed the probabilistic models, reaching AUC values up to 0.765 and Spearman correlation close to 0.59 in weekly series. In addition, t-SNE analyses indicated partial organization of samples according to difficulty levels in the latent space. The results suggest that, although part of the discriminative signal is associated with the original temporal signal itself, latent representations learned by N-BEATS also capture relevant structural information for characterizing forecasting difficulty.

Keywords: Time Series Forecasting · Interpretable Neural Networks · Latent Space Analysis · Uncertainty Estimation · N-BEATS · Predictive Difficulty

1 Introduction

Forecasting is widely used to support decisions related to resource planning and organizational strategies across different time horizons. In this context, time

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 163–180, 2027.
https://doi.org/10.1007/978-3-032-30518-3_11

series forecasting plays a fundamental role in several applied domains, including economics, energy, healthcare, and urban planning [9].

In recent years, deep learning architectures such as Recurrent Neural Networks (RNNs), Long Short-Term Memory networks (LSTMs), and Gated Recurrent Units (GRUs) have gained significant popularity in time series forecasting across various applied domains, demonstrating strong performance, particularly in scenarios characterized by large volumes of data. However, variability in the performance of these architectures, associated with factors such as random weight initialization, raises concerns regarding their consistency and reliability [28].

Most of these architectures, however, operate as "black-box" models, in which the relationship between input variables and predictions remains opaque, requiring the use of post hoc interpretability methods, such as SHAP and LIME, to assist in understanding the model's decision-making process [6].

In contrast to this trend, the N-BEATS model (Neural Basis Expansion Analysis for Time Series Forecasting) proposes a deep architecture capable of producing interpretable forecasts by decomposing the predicted signal into basis components, enabling the extraction of explainable factors analogous to traditional decomposition approaches into trend and seasonality [15].

Despite the structural interpretability introduced by models such as N-BEATS, understanding the reliability of forecasts in deep architectures remains a relevant challenge, as explainability methods often face limitations related to the trade-off between transparency, performance, and computational cost [12].

Predictive uncertainty estimation has been extensively investigated in the literature, since point forecasts are often insufficient to represent the variability and randomness inherent in modeled systems. Probabilistic approaches aim to address this limitation by providing additional information that supports more robust analysis and decision-making in complex scenarios [11].

However, many uncertainty quantification approaches in deep neural networks require modifications to model structures or inference strategies based on multiple executions, such as Bayesian methods, ensembles, and test-time augmentation techniques. These characteristics may imply increased computational and operational complexity, which can hinder their adoption in certain applied contexts [7].

In this sense, investigating whether internal representations learned by the model provide signals about forecast reliability emerges as a promising direction. In N-BEATS, forecasts are constructed through basis function expansions parameterized by interpretable coefficients, which can be viewed as a latent representation of temporal behavior. Deep learning studies suggest that such representations integrate correlated information and offer compact system descriptions, supporting both performance and interpretability [15]. Furthermore, latent representation learning has been associated with improved predictive capability across different settings [19].

In this context, this work investigates whether the geometry of the coefficient space generated by the N-BEATS model can be explored as a proxy measure of

forecasting difficulty. In particular, we analyze whether the magnitude of latent activations is associated with prediction error and can be used to identify difficult forecasts. Additionally, we compare this geometric approach with probabilistic latent-space density estimation methods, including *Gaussian Mixture Models* and *Kernel Density Estimation*, as well as with baselines computed directly from the input signal, such as amplitude and variance.

The main contribution of this work is the investigation of whether forecasting difficulty signals emerge both from simple statistical properties of the temporal signal and from the geometry of latent representations learned by N-BEATS. To this end, we propose a geometric metric based on the aggregated magnitude of latent activations and compare it with latent-space probabilistic models and input-space baselines.

This paper is organized as follows. Section 2 presents related work on time series forecasting, interpretability, and latent-space uncertainty. Section 3 describes the proposed methodology, including the N-BEATS architecture, the construction of latent representations, and the strategies used to model forecasting difficulty. Section 4 presents the experimental setup, including dataset, evaluated models, and metrics. Section 5 presents and discusses the empirical results. Finally, Sect. 6 concludes the paper and outlines directions for future research.

2 Related Work

Recent advances in time series forecasting have been driven by deep learning models, as well as by the growing interest in interpretability and uncertainty quantification. These studies indicate that internal representations learned by neural networks influence both predictive performance and the understanding of model behavior, motivating the investigation of structural and probabilistic properties of latent spaces.

2.1 Deep Learning for Time Series Forecasting

Deep learning models have become effective approaches for time series forecasting due to their ability to capture nonlinear relationships and automatically learn complex representations, as discussed by [4]. Different architectures, such as recurrent, convolutional, and attention-based networks, allow the exploration of temporal dependencies across multiple forecasting horizons.

Recent studies, such as [23], investigate more sophisticated architectures, including decomposition-based and multi-scale structures, reinforcing the role of deep learning as a dominant forecasting paradigm. However, as highlighted by [10], the increasing complexity of these models also amplifies challenges related to interpretability and the analysis of internal representations. Surveys such as [14] point out that latent representation learning plays a central role in capturing relevant temporal variation factors, although questions regarding generalization and transparency remain open.

2.2 Interpretability in Time Series Models

Interpretability in neural forecasting models has received increasing attention due to the opaque nature of these architectures. In this context, [8] demonstrate that interpretable recurrent models can learn variable-specific representations and identify their temporal relevance. Hybrid approaches combining neural networks and statistical decomposition also favor the interpretation of latent patterns, as discussed by [13], while intrinsically interpretable architectures aim to reduce the "black-box" character by incorporating structural mechanisms directly into the model [27].

The advancement of Explainable Artificial Intelligence (XAI) has expanded the repertoire of techniques aimed at interpreting forecasts, including global and local explanations and feature attribution methods, as summarized by [12]. From a methodological perspective, *ante-hoc* and *post-hoc* approaches can be distinguished, which in time series forecasting include attention mechanisms, saliency maps, and representative-pattern-based methods, still subject to challenges related to stability and long-range dependency interpretation [18].

2.3 Uncertainty Estimation in Forecasting

Uncertainty quantification in neural networks has been investigated under different paradigms. Surveys such as [7] organize these approaches into Bayesian methods, ensembles, test-time augmentation strategies, and deterministic solutions that estimate confidence from a single network.

Concerns regarding reliability are not recent: early studies already pointed out limitations related to noise sensitivity and out-of-sample generalization in neural networks applied to time series forecasting [5]. More recently, it has been argued that explanations produced by complex models should also be accompanied by explicit uncertainty measures in order to improve the practical usefulness of machine-learning-based decisions [21].

2.4 Latent-Space Uncertainty

Recent research lines investigate uncertainty estimation directly in the space of learned representations. The DUQ method proposed by [24] uses latent proximity measures to identify unreliable predictions or out-of-distribution samples. Complementarily, [16] show that the geometric organization of these representations reflects data coverage properties.

In generative models, latent vectors tend to organize complex variation factors into semantically coherent directions, favoring interpolation and probabilistic modeling [3]. Furthermore, analyses based on differential geometry indicate that latent spaces may assume the form of nonlinear manifolds associated with different levels of uncertainty and representational distortion [2].

3 Methodology

This section presents the proposed methodology to investigate the relationship between latent representations of the N-BEATS model and forecasting difficulty. First, the forecasting problem formulation and the adopted architecture are described. Next, the construction of the latent representation, the proposed geometric metric, and the comparison models are presented. Finally, the forecasting difficulty criteria and the evaluation strategy are defined. :contentReference[oaicite:0]index=0

3.1 Forecasting Problem Formulation

Time series forecasting consists of estimating future values of a variable from historical observations sequentially organized over time. It is assumed that patterns present in the historical data contain relevant information for forecasting, although they are subject to random variability and unobserved factors [9].

Consider a univariate time series $y_t{}_{t=1}^{T}$. Given an input window of length L, the objective is to learn a forecasting function

$$f : \mathbb{R}^L \to \mathbb{R}^H, \tag{1}$$

where H represents the forecasting horizon. Thus, the model produces point estimates $\hat{\mathbf{y}}_{t+1:t+H}$ from past observations $\mathbf{y}_{t-L+1:t}$.

3.2 N-BEATS Architecture

The N-BEATS (*Neural Basis Expansion Analysis for Time Series Forecasting*) model consists of a fully connected deep architecture designed for univariate time series forecasting tasks. The network is organized into stacks composed of multiple residual blocks, which operate sequentially over the input signal, allowing the progressive modeling of different temporal components [15].

Consider an input vector $x \in \mathbb{R}^L$ corresponding to the historical observation window used as model input, and a forecasting horizon H. Each block b receives as input a residual vector $x_b \in \mathbb{R}^L$ and produces two sets of expansion coefficients: θ_b^b, associated with the reconstruction component (*backcast*), and θ_b^f, associated with the forecasting component (*forecast*). These coefficients are obtained through a stack of fully connected layers with nonlinear activations.

From these coefficients, the block generates its two outputs through specific basis functions, given by

$$\hat{x}_b = g_b^b(\theta_b^b), \qquad \hat{y}_b = g_b^f(\theta_b^f), \tag{2}$$

where $\hat{x}_b \in \mathbb{R}^L$ represents the block backcast and $\hat{y}_b \in \mathbb{R}^H$ represents its partial contribution to the future forecast.

The residual flow across blocks is iteratively updated as

$$x_{b+1} = x_b - \hat{x}_b, \tag{3}$$

so that each subsequent block operates on the portion of the historical signal not yet explained by previous blocks. In parallel, the forecast contributions are aggregated to compose the final model prediction:

$$\hat{y} = \sum_{b=1}^{B} \hat{y}_b, \tag{4}$$

where B denotes the total number of blocks in the architecture.

Figure 1 illustrates this organization at different levels: the basic block, the residual structure within each stack, and the hierarchical aggregation of partial forecasts until the final model output is obtained.

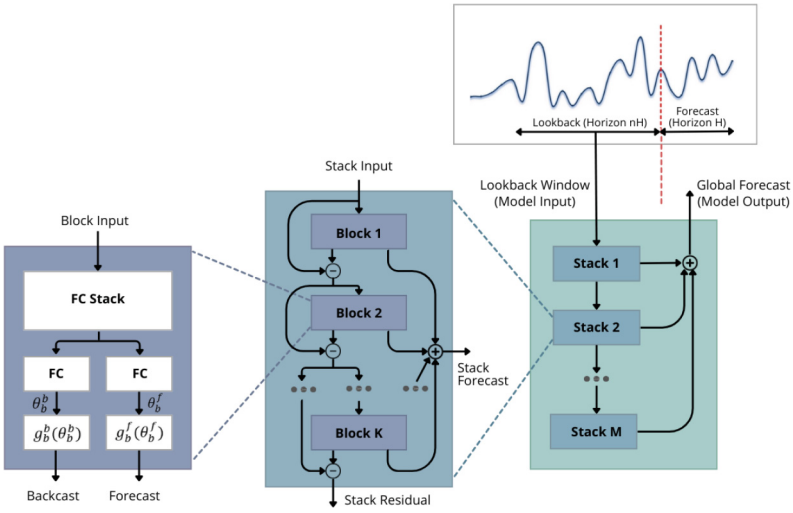

Fig. 1. General architecture of the N-BEATS model. Each block uses a stack of fully connected layers to produce expansion coefficients θ^b and θ^f, which are subsequently transformed by the basis functions $g^b(\cdot)$ and $g^f(\cdot)$ into backcast and forecast components. The final prediction is obtained through the aggregation of contributions across residual stacks. Adapted from [15].

3.3 Latent Representation in the Coefficient Space

A central characteristic of N-BEATS is the construction of forecasts through an expansion in basis functions parameterized by coefficients learned by the network. In each block, the forecast and backcast outputs are obtained from linear projections of latent coefficient vectors.

Formally, let $\theta_b^{(f)} \in \mathbb{R}^{d_f}$ and $\theta_b^{(b)} \in \mathbb{R}^{d_b}$ be the coefficient vectors associated with the forecast and reconstruction components of block b, respectively. The block outputs can be expressed as

$$f_b = G_f \theta_b^{(f)}, \quad b_b = G_b \theta_b^{(b)}, \tag{5}$$

where $G_f \in \mathbb{R}^{H \times d_f}$ and $G_b \in \mathbb{R}^{L \times d_b}$ represent basis-function matrices.

The coefficient vectors produced by the blocks can be interpreted as a structured latent representation of the temporal behavior modeled by the network. Considering the forecast coefficients of all blocks, the aggregated representation is defined as

$$\Theta = \left[\theta_1^{(f)}, \theta_2^{(f)}, \ldots, \theta_B^{(f)}\right], \tag{6}$$

and its vectorized form as

$$z = \mathrm{vec}(\Theta), \tag{7}$$

where $z \in \mathbb{R}^d$ encodes the temporal patterns used to generate the forecast. This latent representation forms the basis for constructing the forecasting difficulty metrics proposed in this work.

3.4 Geometric Metric

After model training, the forecast coefficients of each interpretable block are extracted. To handle differences in vector length and the presence of padding, the vectors are truncated at the last valid coefficient, ensuring alignment with the evaluation variables.

For each block b, the Euclidean norm of the coefficients is computed as

$$s_b = \|\theta_b^f\|_2, \tag{8}$$

defining the latent activation vector

$$\mathbf{s} = [s_1, s_2, \ldots, s_B]. \tag{9}$$

From this vector, the global scalar metric is defined as

$$S = \|\mathbf{s}\|_2, \tag{10}$$

denoted as *radial_norm*, interpreted as an indicator of the internal activation intensity mobilized by the model to generate each forecast. This metric is evaluated as a proxy for forecasting difficulty under the hypothesis that greater latent activation intensity is associated with higher representational effort by the model.

Figure 2 presents the adopted methodological pipeline.

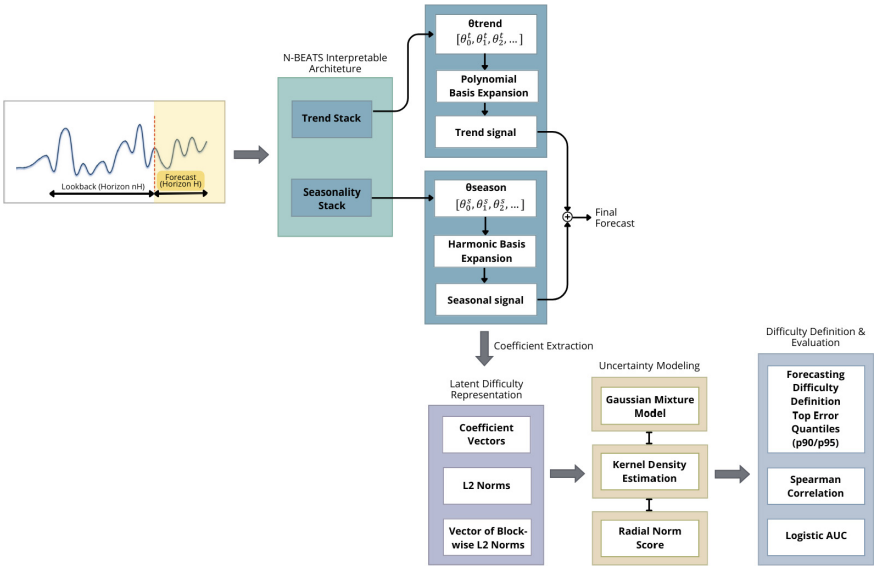

Fig. 2. Overview of the proposed methodological pipeline. The interpretable N-BEATS architecture generates latent coefficients associated with trend and seasonality blocks, which are transformed into geometric representations used for uncertainty modeling and forecasting difficulty definition.

3.5 Latent-Space Uncertainty Modeling

As a comparison to the proposed geometric metric, classical probabilistic models applied to the latent space were considered.

Gaussian mixture models are widely used in signal and time series modeling due to their ability to represent complex and multimodal distributions through combinations of simple components. Approaches using GMM to model trajectories in reconstructed phase spaces have demonstrated good performance in temporal classification [17]. Discriminative extensions based on large margins have also been proposed to improve class separability when modeling latent distributions [20]. In addition, applications in different domains indicate that GMM provides a robust strategy for representing structural variations and supporting likelihood-based decisions [22].

Nonparametric density estimation methods, such as *Kernel Density Estimation* (KDE), offer greater flexibility by allowing unknown distributions to be modeled without assuming specific functional forms. In time series analysis, recursive variants enable adaptive updating of estimates over time [1]. KDE has also been employed in tasks such as anomaly detection and learned representation modeling [29]. Its application in temporal classification problems reinforces the potential of this method for characterizing probabilistic structures in sequential data [25].

In this work, these approaches were used to model the distribution of latent representations extracted from the forecasting model. It was assumed that such a distribution could be approximated by a finite Gaussian mixture in the case of GMM, allowing the capture of possible multimodal structures. Alternatively, KDE was employed to estimate the local density of latent activations in a non-parametric manner, with the respective scores defined from the observations' negative log-likelihood (NLL).

In contrast to these probabilistic approaches, the proposed geometric metric is based exclusively on the magnitude of latent activations, enabling the investigation of whether the aggregated intensity of the internal representation contains informative signals regarding forecasting difficulty, independently of explicit assumptions about data distribution.

3.6 Input-Space Baselines

To verify whether the discriminative signal associated with forecasting difficulty originates from the latent representation or from properties of the original signal, baselines computed directly from the historical input window used by N-BEATS were added.

More specifically, two simple statistical descriptors of the input sequence X were considered: signal variance and signal amplitude (difference between maximum and minimum values). These baselines allow the comparison between information derived from the original input space and that extracted from the latent representation learned by the model.

It is important to highlight that these metrics are computed from the same historical window responsible for generating the latent coefficients θ, enabling a direct comparison between input-space and latent-space methods.

3.7 Forecasting Difficulty Definition

Forecasting difficulty was defined from the continuous error associated with each temporal instance. Binary scenarios based on percentile thresholds were also considered, particularly $P90$ and $P95$.

3.8 Evaluation Strategy

The relationship between the scores and forecasting error was analyzed using the Spearman correlation coefficient. Additionally, discriminative capability was evaluated using the area under the ROC curve (AUC), considering as positive class the observations belonging to the upper error quantiles.

4 Experimental Configuration

This section presents the experimental setup used to evaluate the relationship between latent representations of the N-BEATS model and forecasting difficulty

in time series. Different seasonal frequencies, model configurations, and latent-space density modeling strategies were considered. The evaluation was based on correlation and discriminative performance metrics, allowing the analysis of the potential of geometric signals derived from interpretable coefficients as indicators of predictive difficulty.

4.1 Dataset

The experiments were conducted using time series from the M4 dataset, widely adopted as a benchmark for forecasting tasks. Monthly, weekly, and yearly categories were considered, since they exhibit distinct seasonal patterns and forecasting horizons.

The interpretable version of N-BEATS (*NBEATS-I*) was adopted, investigating the impact of different configurations defined by the combinations *repeat*=0, *lookback* $\in \{2, 4, 6\}$, and *loss* $\in \{\text{MAPE}, \text{MASE}\}$.

4.2 Observation Extraction and Selection

During training, the interpretable coefficients produced by the blocks were stored for each pair (time series, forecasting instant). Initially, approximately 15.36 million observations were obtained for monthly series, 512 thousand for weekly series, and 15.36 million for yearly series.

To focus on representations learned after training stabilization, only the last 50 iterations were considered, resulting in approximately $51,200$ observations per category. Repeated pairs were then removed by keeping only the last occurrence, reducing each set to approximately $45,000$ unique observations.

4.3 Latent Representation

Each observation is associated with six interpretable blocks of the model, including three trend blocks and three seasonality blocks. For each block b, the Euclidean norm of the forecast coefficients was computed as

$$s_b = \|\theta_b^f\|_2, \tag{11}$$

defining the latent activation vector

$$\mathbf{s} = [s_1, s_2, \ldots, s_6], \tag{12}$$

which summarizes the intensity of activations produced by the architecture.

4.4 Preprocessing and Data Splitting

For each experiment and series type, the data were randomly divided into fitting (70%) and evaluation (30%) sets, without replacement and using a fixed random seed.

Before density modeling, the latent vectors were transformed using the *Yeo-Johnson* transformation, a generalization of the Box-Cox transformation that allows handling both positive and negative values and is commonly used for variance stabilization and approximation to normality [26], followed by standardization using parameters estimated only from the fitting set.

4.5 Density Modeling and Scores

Three types of latent-space scores were considered: a parametric Gaussian mixture model with diagonal covariance and automatic component selection via BIC (*gmm_bestK*); a nonparametric *Kernel Density Estimation* model with Gaussian kernel (*kde*); and a deterministic geometric metric based on latent activation magnitude (*radial_norm*).

For *gmm_bestK*, the number of components was selected within the range $k \in \{2, \ldots, 9\}$, and the score was defined as the observations' negative log-likelihood (NLL) on the evaluation set. For KDE, a fixed bandwidth equal to 1.0 was adopted, also using NLL as a surprise measure.

4.6 Difficulty Definition and Evaluation

Forecasting difficulty was defined from the final forecasting error. Both a continuous measure (*difficulty_continuous*) and a binary definition based on the 90th percentile of the error distribution (*difficulty_high_p90*) were considered.

The relationship between latent scores and forecasting error was evaluated using Spearman correlation, while discriminative capability in identifying difficult forecasts was measured through the area under the ROC curve (AUC).

5 Results and Discussion

This section presents and discusses the results obtained from the analysis of the relationship between N-BEATS latent representations and forecasting difficulty, considering geometric analyses, correlation with continuous forecasting error, and performance in difficult forecast detection tasks. :contentReference[oaicite:0]index=0

5.1 Latent Space Visualization

Figure 3 presents two-dimensional projections of the latent space obtained through t-SNE applied to the *block_l2_vector* vectors, constructed from the norms of the N-BEATS latent coefficients. Before applying t-SNE, a preliminary dimensionality reduction step using PCA was performed.

Each panel represents a specific experimental configuration and temporal frequency. Points were colored according to forecasting difficulty level, defined from the continuous forecasting error: *Low* ($\leq P50$), *Medium* ($P50$–$P90$), and

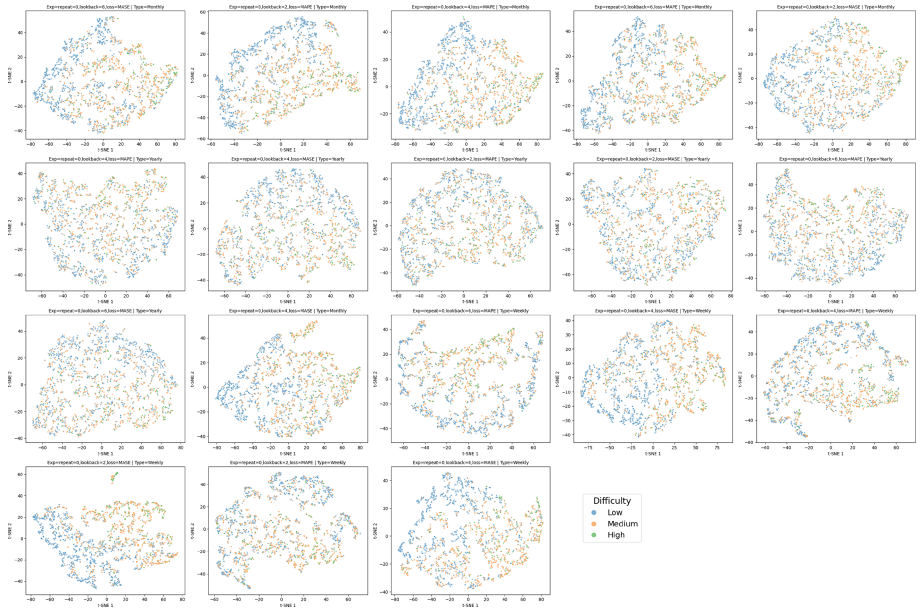

Fig. 3. t-SNE projections of the latent space for different experimental configurations and temporal frequencies. Points represent samples projected from the *block_l2_vector* vectors after preliminary dimensionality reduction using PCA. Colors indicate forecasting difficulty levels defined from the continuous forecasting error: *Low* ($\leq P50$), *Medium* ($P50$–$P90$), and *High* ($> P90$).

High ($> P90$). To facilitate visualization, stratified sampling with up to 2000 points per panel was used.

Overall, monthly and weekly series present more organized regions and greater separation between low- and high-difficulty samples. In contrast, yearly series exhibit greater overlap between groups, consistent with the lower AUC and correlation values previously observed.

Although t-SNE does not strictly preserve global distances, the obtained projections provide qualitative evidence that the latent representations learned by N-BEATS partially organize difficult and non-difficult samples into different latent-space regions.

5.2 Score Distribution by Difficulty Level

Figure 4 presents the distribution of the evaluated scores considering the binary formulation of forecasting difficulty, in which observations belonging to the upper error percentile are labeled as the positive class.

Overall, a consistent shift in the distributions associated with high-difficulty samples can be observed, reflected by increased median and dispersion of the scores in the positive class. This behavior occurs both in latent-space methods and in baselines computed directly from the input signal.

Among the latent methods, *radial_norm* presents clearer visual separation between difficulty groups when compared to the probabilistic *GMM* and *KDE* models, consistent with the quantitative results previously observed in terms of AUC and Spearman correlation.

Variance and amplitude baselines also exhibit strong separation between groups, especially in monthly series, indicating that simple statistical properties of the original signal are significantly related to forecasting difficulty.

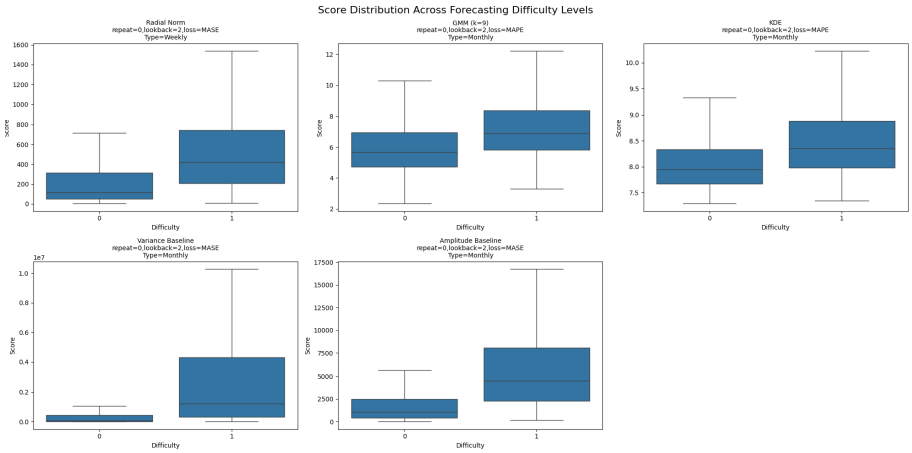

Fig. 4. Distribution of the evaluated scores for different forecasting difficulty levels, including latent-space methods (*radial_norm*, *GMM*, and *KDE*) and baselines computed directly from the input space (variance and amplitude).

5.3 Performance in Difficult Forecast Detection

The results show that baselines computed directly from the input space, especially amplitude and variance, achieved the highest global AUC and Spearman correlation values, mainly in monthly and weekly series. This behavior suggests that part of the discriminative signal associated with forecasting difficulty is already present in the variability and scale of the temporal signal itself.

Among the latent-space methods, the geometric *radial_norm* approach consistently outperformed the probabilistic *KDE* and *GMM* models while maintaining competitive performance across different experimental scenarios.

Table 1 presents a comparison between the best results obtained by the evaluated methods.

Overall, the results indicate that forecasting difficulty depends both on simple statistical properties of the original signal and on structural information organized in the latent space learned by N-BEATS.

Table 1. Comparison between the best results obtained by input-space baselines and latent-space methods.

Space	Method	Series	Spearman	AUC
Input	amplitude	Monthly	**0.662**	**0.818**
Input	variance	Monthly	0.649	0.809
Input	amplitude	Weekly	0.616	0.781
Latent	radial_norm	Weekly	0.588	0.765
Latent	radial_norm	Monthly	0.528	0.745
Latent	radial_norm	Weekly	0.528	0.732
Latent	gmm_bestK	Monthly	0.413	0.699
Latent	kde	Monthly	0.337	0.696
Latent	gmm_bestK	Yearly	0.251	0.686

5.4 ROC Curves of the Evaluated Models

Figure 5 presents the ROC curves of the best results obtained by the evaluated methods. Baselines based on amplitude and variance achieved the highest AUC values, indicating strong discriminative capability of the input signal itself. Among the latent-space methods, *radial_norm* consistently outperformed the probabilistic *GMM* and *KDE* models, suggesting that simple geometric measures derived from latent activations can capture relevant information associated with forecasting difficulty.

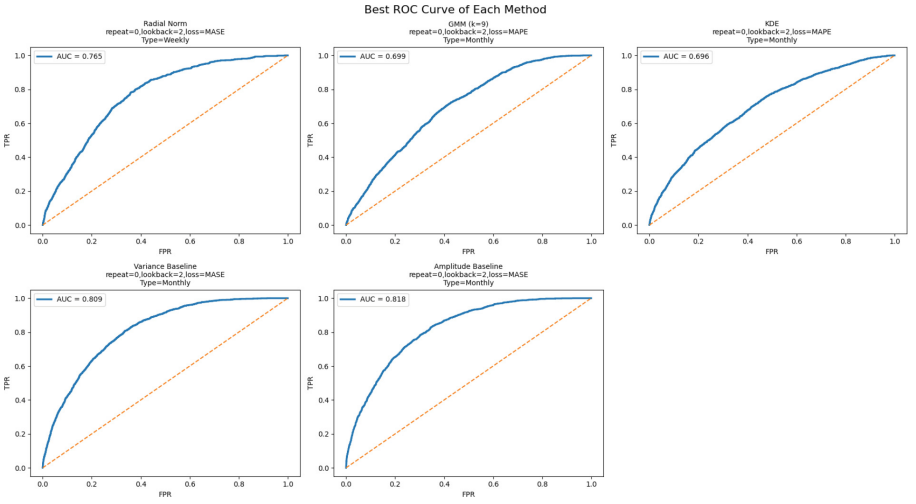

Fig. 5. ROC curves of the best results obtained by the evaluated methods in the task of detecting highly difficult forecasts.

5.5 Correlation Analysis

Figure 6 presents the Spearman correlation coefficients between the different scores and the continuous forecasting error. The highest coefficients were obtained by baselines computed in the input space, especially for monthly and weekly series. Among latent-space methods, *radial_norm* consistently outperformed *GMM* and *KDE*. In contrast, yearly series presented lower correlations in almost all experimental configurations.

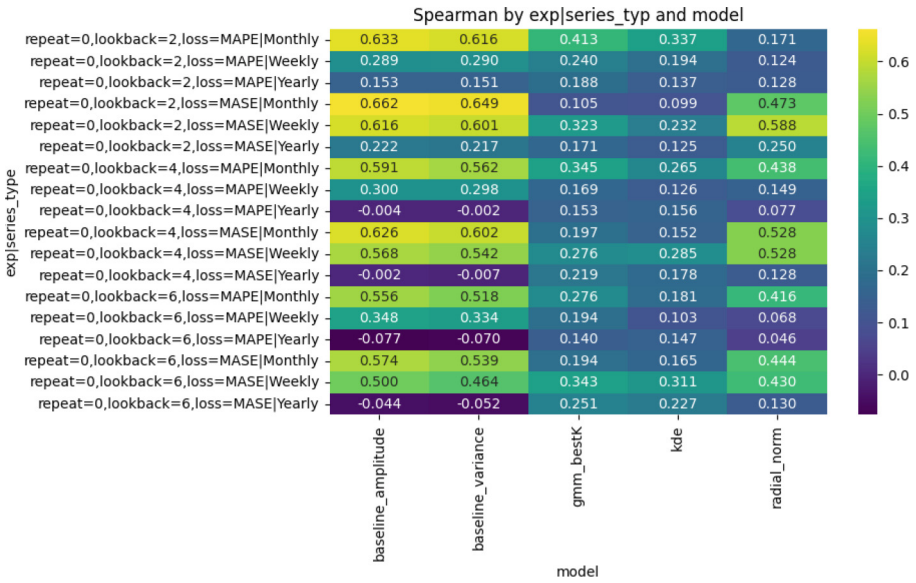

Fig. 6. Spearman correlation coefficients between the evaluated scores and the continuous forecasting error.

6 Conclusion

This work investigated whether latent representations derived from the interpretable coefficients of the N-BEATS model can provide informative signals about forecasting difficulty in time series. For this purpose, a geometric approach based on the magnitude of latent activations was proposed and compared with probabilistic density models and baselines computed directly from the input signal.

The results showed that baselines based on signal amplitude and variance achieved the highest global AUC and Spearman correlation values, indicating that a relevant part of the discriminative signal associated with forecasting difficulty is already present in the variability and scale of the temporal signal itself.

Despite this, latent-space methods remained competitive across different experimental scenarios. In particular, the *radial_norm* metric consistently outperformed the probabilistic *GMM* and *KDE* models, suggesting that simple geometric measures derived from latent activations can capture additional structural information associated with predictive difficulty.

The analyses based on t-SNE projections also indicated partial organization of samples according to difficulty levels in the latent space, especially for monthly and weekly series. Overall, the results suggest that forecasting difficulty depends both on statistical properties of the original signal and on the latent structure learned by the model.

The results also indicate that internal representations based on the interpretable coefficients of N-BEATS can provide relevant information about predictive difficulty, even without structural modifications to the architecture or the use of complex probabilistic strategies.

As future work, we highlight the investigation of the proposed approach in other forecasting architectures and the exploration of more robust geometric metrics for monitoring and calibration of forecasting models.

Acknowledgments. The work of Paulo Novais has been supported by FCT Fundação para a Ciência e Tecnologia within the R&D Unit Project Scope UID/00319/2025 - Centro ALGORITMI (ALGORITMI/UM) https://doi.org/10.54499/UID/00319/2025.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aboubacar, A., El Machkouri, M.: Recursive kernel density estimation for time series. IEEE Trans. Inf. Theory **66**(10), 6378–6388 (2020)
2. Arvanitidis, G., Hansen, L.K., Hauberg, S.: Latent space oddity: on the curvature of deep generative models. arXiv preprint arXiv:1710.11379 (2017)
3. Bojanowski, P., Joulin, A., Lopez-Paz, D., Szlam, A.: Optimizing the latent space of generative networks. arXiv preprint arXiv:1707.05776 (2017)
4. Casolaro, A., Capone, V., Iannuzzo, G., Camastra, F.: Deep learning for time series forecasting: advances and open problems. Information **14**(11), 598 (2023)
5. Connor, J.T., Martin, R.D., Atlas, L.E.: Recurrent neural networks and robust time series prediction. IEEE Trans. Neural Netw. **5**(2), 240–254 (1994)
6. Freeborough, W., van Zyl, T.: Investigating explainability methods in recurrent neural network architectures for financial time series data. Appl. Sci. **12**(3), 1427 (2022)
7. Gawlikowski, J., et al.: A survey of uncertainty in deep neural networks. Artif. Intell. Rev. **56**(Suppl 1), 1513–1589 (2023)
8. Guo, T., Lin, T., Lu, Y.: An interpretable LSTM neural network for autoregressive exogenous model. arXiv preprint arXiv:1804.05251 (2018)
9. Hyndman, R.J., Athanasopoulos, G.: Forecasting: Principles and Practice. OTexts, 3 edn. (2021). https://otexts.com/fpp3/

10. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., Muller, P.-A.: Deep learning for time series classification: a review. Data Min. Knowl. Disc. **33**(4), 917–963 (2019). https://doi.org/10.1007/s10618-019-00619-1
11. Kabir, H.D., Khosravi, A., Hosen, M.A., Nahavandi, S.: Neural network-based uncertainty quantification: a survey of methodologies and applications. IEEE access **6**, 36218–36234 (2018)
12. Kalasampath, K., Spoorthi, K., Sajeev, S., Kuppa, S.S., Ajay, K., Maruthamuthu, A.: A literature review on applications of explainable artificial intelligence (XAI). IEEE Access **13**, 41111–41140 (2025)
13. Liu, S., Wu, K., Jiang, C., Huang, B., Ma, D.: Financial time-series forecasting: towards synergizing performance and interpretability within a hybrid machine learning approach. arXiv preprint arXiv:2401.00534 (2023)
14. Miller, J.A., et al.: A survey of deep learning and foundation models for time series forecasting. arXiv preprint arXiv:2401.13912 (2024)
15. Oreshkin, B.N., Carpov, D., Chapados, N., Bengio, Y.: N-beats: neural basis expansion analysis for interpretable time series forecasting (2020). https://arxiv.org/abs/1905.10437
16. Postels, J., Blum, H., Cadena, C., Siegwart, R., Van Gool, L., Tombari, F.: Quantifying aleatoric and epistemic uncertainty using density estimation in latent space. arXiv preprint arXiv:2012.03082 **1** (2020)
17. Povinelli, R.J., Johnson, M.T., Lindgren, A.C., Ye, J.: Time series classification using gaussian mixture models of reconstructed phase spaces. IEEE Trans. Knowl. Data Eng. **16**(6), 779–783 (2004)
18. Rojat, T., Puget, R., Filliat, D., Del Ser, J., Gelin, R., Díaz-Rodríguez, N.: Explainable artificial intelligence (xai) on timeseries data: a survey. arXiv preprint arXiv:2104.00950 (2021)
19. Rongali, S., et al.: Learning latent space representations to predict patient outcomes: model development and validation. J. Med. Internet Res. **22**(3), e16374 (2020)
20. Sha, F., Saul, L.K.: Large margin gaussian mixture modeling for phonetic classification and recognition. In: 2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings, vol. 1, pp. I–I. IEEE (2006)
21. Slack, D., Hilgard, A., Singh, S., Lakkaraju, H.: Reliable post hoc explanations: modeling uncertainty in explainability. Adv. Neural. Inf. Process. Syst. **34**, 9391–9404 (2021)
22. Słoński, M.: Gaussian mixture model for time series-based structural damage detection. Comput. Assist. Methods Eng. Sci. **19**(4), 331–338 (2012)
23. Song, X., Deng, L., Wang, H., Zhang, Y., He, Y., Cao, W.: Deep learning-based time series forecasting. Artif. Intell. Rev. **58**(1), 23 (2024)
24. Van Amersfoort, J., Smith, L., Teh, Y.W., Gal, Y.: Uncertainty estimation using a single deep deterministic neural network. In: International Conference on Machine Learning, pp. 9690–9700. PMLR (2020)
25. Wei, C., Peng, B., Li, C., Liu, Y., Ye, Z., Zuo, Z.: A two-stage optimized robust kernel density estimation for Bayesian classification with outliers. Int. J. Mach. Learn. Cybern. **16**(12), 9999–10023 (2025)
26. Weisberg, S.: Yeo-johnson power transformations. Department of Applied Statistics, University of Minnesota. Retrieved June 1 (2003), 8 (2001)
27. Yan, Y., Zhu, J., Duda, M., Solarz, E., Sripada, C., Koutra, D.: Groupinn: grouping-based interpretable neural network for classification of limited, noisy brain data. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 772–782 (2019)

28. Yunita, A., et al.: Performance analysis of neural network architectures for time series forecasting: a comparative study of RNN, LSTM, GRU, and hybrid models. MethodsX **15**, 103462 (2025)
29. Zayara, A., Fandeyev, V.: Kernel probability density estimation in solving the problem of classification of the technical condition of complex systems. J. Commun. Technol. Electron. **70**(5), 226–231 (2025)

Virtuality in Teams: A Multilevel Analysis of Its Effects on Innovation and Satisfaction

Isabel Dórdio Dimas[1(✉)], Teresa Rebelo[2], Paulo Renato Lourenço[2], and Marta Alves[3]

[1] CeBER, Faculty of Economics, University of Coimbra, Coimbra, Portugal
idimas@fe.uc.pt
[2] CeBER, Faculty of Psychology and Education Sciences, University of Coimbra, Coimbra, Portugal
{terebelo,prenato}@fpce.uc.pt
[3] NECE-UBI, University of Beira Interior, Covilhã, Portugal
mpalves@ubi.pt

Abstract. This study contributes to clarifying how virtuality impacts team innovation and team member satisfaction through its influence on teamwork management. Teamwork management refers to the degree to which team members structure and plan their work to facilitate the implementation of new practices and monitor the work progress and involves the following dimensions: work planning/organization, resource management and support for innovation. Three hundred and fifty-eight team members from 88 teams were surveyed. A two-wave data collection procedure was implemented with a time lag of about two months. Data was collected from two sources: team members were surveyed about virtuality, teamwork management and team member satisfaction, while team leaders assessed the levels of innovation of their teams. Due to the multilevel nature of data (team member satisfaction is at the individual level and the remaining variables at the group level), hypotheses were tested through Generalized Structural Equation Modelling in Stata. Results revealed that the indirect effect of virtuality on innovation through teamwork management was positive and significant, as well as the indirect effect on team member satisfaction. The findings emphasize that virtuality, by encouraging greater engagement in behaviours such as planning work, organizing resources, and providing mutual support, fosters more positive attitudes toward the team and facilitates the introduction of new ideas and procedures.

Keywords: Virtuality · teamwork management · team effectiveness

1 Introduction

Advances in communication and information technologies have fundamentally transformed how organizations structure work. These developments have enabled companies to adopt a wide range of virtual work arrangements that reduce, and in some cases eliminate, traditional constraints of time and physical location. Employees can now collaborate from different cities, countries, or even continents with relative ease (Gibson

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 181–191, 2027.
https://doi.org/10.1007/978-3-032-30518-3_12

et al., 2023). Although virtual work had already been steadily increasing, the COVID-19 pandemic accelerated this shift dramatically. The sudden and widespread move to remote work during the pandemic represented what many scholars consider one of the most profound job design disruptions of our time (George et al., 2020). Organizations were forced to rethink how work is coordinated, how teams communicate, and how collaboration occurs under conditions of physical separation.

As a result of these changes, contemporary work teams increasingly rely on digital tools, such as videoconferencing platforms, collaborative document systems, and asynchronous communication channels, to coordinate tasks, exchange information, and pursue shared goals (Handke et al., 2026). Virtual interaction has therefore become a defining feature of modern teamwork rather than a supplemental or temporary arrangement. This evolution raises important questions about how team processes unfold when face-to-face interaction is limited or replaced by technology-mediated communication.

In this context, understanding the role of virtuality in shaping team dynamics, effectiveness, and outcomes has become crucial (Purvanova & Kenda, 2022). Team virtuality can be conceptualized as a continuum that reflects the degree to which team members use communication media to coordinate their actions and perform their tasks (Dixon & Panteli, 2010). As organizations increasingly adopt hybrid work models, teamwork has likewise become more hybrid, with many teams operating at different points along the spectrum between "never" and "always" meeting face-to-face (Cousins et al., 2007; Handke et al., 2026).

Research on the effects of team virtuality on team effectiveness has mostly focused on performance and has produced contradictory results (Purvanova & Kenda, 2022). While some studies show that team virtuality contributes to performance by increasing available resources, others point out that team virtuality hinders communication and cooperation, which in turn leads to a decrease in performance levels (Zheng et al., 2024).

In this study, we go beyond the commonly studied dimension of team effectiveness (i.e., team performance) to explore other equally important and less studied criteria (Hackman, 1987). Specifically, we intend to contribute to clarifying how virtuality impacts team innovation (i.e., the intentional introduction and adoption of new ideas and solutions in the group) and team member satisfaction (i.e., an affective response of members toward the group, its characteristics and functioning) through its influence on teamwork management.

Teamwork management refers to the way team members organize themselves to plan their tasks and perform their work with the aim of enhancing overall team effectiveness (Rousseau et al., 2006). According to Rousseau and colleagues, teamwork management includes the following dimensions: work planning/organization, resource management and support for innovation dimensions. Work planning involves coordinating and scheduling tasks to ensure they are completed efficiently and within established deadlines. Resource management concerns how team members use and distribute available resources, strategically allocating them to achieve desired goals and outcomes. Support for innovation reflects the extent to which the team encourages the generation of new ideas and practices, as well as their practical implementation and adoption within the group.

Drawing on conservation of resources theory (COR; Hobfoll et al., 2018), virtuality can be understood as a condition that simultaneously facilitates resource gains and exposes individuals to potential resource losses (Zhang et al., 2024). In terms of gains, virtual work often expands access to information through shared digital platforms, increases transparency of work processes, and provides employees with greater flexibility in managing their tasks and schedules (Handke et al., 2025). These benefits can enhance efficiency and reduce time pressures, thereby enriching individuals' resource pools.

At the same time, virtuality can create conditions that drain resources. Technology-mediated communication may limit spontaneous interaction and reduce the richness of social cues, making it more difficult to build trust, coordinate tasks smoothly, or maintain strong interpersonal relationships (Carter et al., 2019). These challenges can weaken collaboration and increase the cognitive effort needed to interpret messages or resolve misunderstandings. Thus, virtuality introduces a dynamic interplay of resource gains – such as increased information availability – and resource losses – such as diminished collaboration and social connectedness (Zhang et al., 2024).

Individuals working in virtual environments are motivated to engage in behaviours that help them obtain additional resources while safeguarding those they already possess (Hobfoll et al., 2018). Therefore, when team virtuality increases, team members are likely to invest more effort in team management behaviours – that is, behaviours aimed at structuring, coordinating, and supporting collective work processes – to obtain the benefits of working virtually (Rousseau et al., 2006).

These behaviours become especially important in virtual contexts that require deliberate coordination through digital tools, where clarity of roles, timelines, and expectations must be intentionally constructed rather than emerging organically. Research on virtual teamwork emphasizes that effective planning, transparent communication, and shared monitoring help reduce ambiguity and maintain high-quality interaction even when team members are geographically dispersed (Handke et al., 2025). Such behaviours function simultaneously as resource-protective strategies – preventing misunderstandings, delays, or interpersonal strain – and as resource-building strategies, enabling access to information and strengthening team cohesion.

Team management behaviours are well-established predictors of team effectiveness. Rousseau et al. (2006) argue that coordinated planning, mutual support, and efficient resource allocation facilitate not only smoother task execution but also promote a shared understanding of goals and interdependencies. When team members engage in these behaviours consistently, they create a structured and supportive team environment that enhances members' confidence in their collective capabilities. This structured teamwork environment is particularly valuable in virtual teams, where coordination and communication require more intentional effort (Purvanova & Kenda, 2022).

By planning tasks effectively, supporting one another, and distributing resources strategically, teams cultivate positive attitudes among members, reflected in higher individual satisfaction. Such attitudes emerge because well-managed teamwork reduces stressors, increases predictability, and fosters a sense of belonging and mutual reliability, resources that are especially important when physical distance limits informal social interaction (Mahadevan, 2024). Moreover, teamwork management also creates

the psychological and cognitive conditions necessary for innovation (Rousseau et al., 2006) When teams manage their work effectively, they free up cognitive resources that can be redirected toward creative thinking, exploration of new ideas, and collective problem-solving.

Thus, as virtuality increases, teams will tend to invest more in team management behaviours not only to avoid the negative effects of resource losses but also to leverage resource gains, ultimately enhancing both individual satisfaction and team innovation. Therefore, the following hypotheses are proposed:

H1: Team virtuality has an indirect relationship with team members' satisfaction through teamwork management.

H2: Team virtuality has an indirect relationship with team innovation through teamwork management.

Our conceptual model is depicted in Fig. 1.

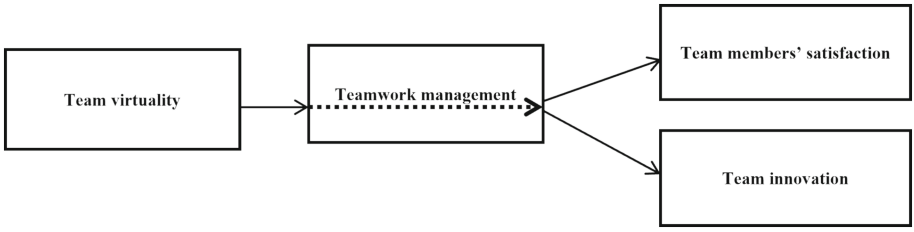

Fig. 1. Hypothesized model under analysis

2 Materials and Methods

2.1 Procedure and Sample

Organizations were identified using both the authors' personal and professional networks and national company databases that are accessible online. Data were collected through surveys, which were mostly completed online (in a few cases, surveys were completed on paper). To mitigate common method variance (Podsakoff et al., 2003), we employed a two-wave data collection approach with a two-month interval between measurements. Team virtuality and teamwork management was collected in the first wave, while team innovation and team member satisfaction was collected in the second wave. This time frame was strategically chosen to balance two objectives: minimizing common method bias while ensuring the interval remained sufficiently brief for the hypothesized antecedents to potentially influence the outcome variable. In addition, data were obtained from two sources: team members were surveyed about team virtuality, teamwork management and individual satisfaction, while team innovation was obtained from the leader.

The sample is composed of 358 members from 88 teams from different companies. Team size varied from 3 to 43 members with an average of approximately six members (SD = 5.37) and teams have been working together for about eight years (SD = 10.50). Teams developed different types of activities, most of them were services teams (46%).

2.2 Measures

To measure team virtuality, we used the indicator proposed by De Jong et al.'s (2008), which is based on the criteria of Baltes et al. (2002) that consider three factors: (a) the frequency of technology use, (b) the degree of synchronisation, and (c) the presence of nonverbal and para verbal signs. Team members are asked to provide the percentage of use (up to 100%) of different modes of communication in team interactions in the last month (including options such as in-person, video conference, conference call, and chat). Then, the equation proposed by the authors is implemented (see De Jong et al., 2008). Finally, the results obtained are subtracted from 100% to give the level of virtuality. Higher values indicate a greater degree of virtuality.

Teamwork management was measured using the scale developed by Rousseau et al., (2006), which is composed of six items that are measured using a 5-point Likert-type scale, from 1 (not true) to 5 (completely true). A sample item is "We organize our work activities".

To measure team member satisfaction, we used the scale developed by Dimas and Lourenço (2015). This scale is composed of seven items that measure members' satisfaction with different aspects related to the team. A sample item is "team functioning". Statements are evaluated on a 7-point Likert-type scale ranging from 1 (totally dissatisfied) to 7 (totally satisfied).

Finally, team innovation was measured through the three-item scale of Batarseh et al. (2017), which, in turn, was based on the two-item scale of Vera and Crossan (2005). The responses are evaluated on a Likert scale ranging from (1) totally disagree to (7) totally agree. A sample item is: "The team frequently introduces new product/service innovations."

Team size and team tenure were considered as control variables in the present study, as previous research has shown that these variables may influence team outcomes (Standifer et al., 2015; West & Andreson, 1996). As some teams were surveyed before the start of the COVID-19 pandemic (2020) and others during the pandemic (2021, 2022), a dummy variable (0 = surveyed before the pandemic; 1 = surveyed during the pandemic) was introduced to account for the possible influence of the pandemic crisis on the outcome variables.

3 Results

3.1 Data Aggregation

Since teamwork management was operationalized at the team level but obtained at the individual level using referent shift consensus (Chan, 1998), within-team homogeneity and between-team heterogeneity need to be assured. Accordingly, we calculated the inter-rater reliability index r_{wg} (James et al., 1993) and the intra-class correlation coefficients ICC (1) and ICC (2) (Bliese, 2000).

The r_{wg} index relates the within-group variance of a single item *j* in a group of K raters (S^2_{jk}) to an expected variance that assumes all ratings were due to random responding

(σ^2_{EU}). Since our scales have multi-items, $r_{wg(j)}$ was implemented which is given by the following formula:

$$r_{wg(j)} = \frac{J \times \left(1 - \frac{S^2_{jk}}{\sigma^2_{EU}}\right)}{1 + (j-1) \times \left(1 - \frac{S^2_{jk}}{\sigma^2_{EU}}\right)}$$

Following Bliese (2000), *ICC (1)* was calculated using Bartko (1976) formula:

$$ICC(1) = \frac{MSB - MSW}{MSB + [(k-1) \times MSW]}$$

MSB is the between-group mean square, *MSW* is the within-group mean square, and K is the group size. This formula is a way of contrasting the between-group and within-group variance components from the ANOVA model. Since group sizes were different among teams, the group size average was used. While *ICC (1)* may be considered as a measure of the reliability associated with a single assessment of the group mean, *ICC (2)* provides an estimate of the reliability of the group means and is given by the following formula (Bliese, 2000):

$$ICC(2) = \frac{MSB - MSW}{MSB}$$

The results showed adequate average r_{wg} values for teamwork management (M = .89, SD = .21), well above the recommended cut-off of point of.70 (James et al., 1993). The results for *ICC* (1) and *ICC* (2) for teamwork management were.30 and.67, respectively, also revealing adequate results. Taken together, the r_{wg}, *ICC* (1) and *ICC* (2) values provide sufficient justification for aggregating the data to the team level.

3.2 Measurement Model Assessment

We conducted a confirmatory factor analysis on Stata (version 16), using the maximum likelihood method of estimation to ensure that the latent variables under study presented adequate convergent and discriminant validity. The three-factor measurement model showed adequate results ($\chi 2$ (101) = 354.06, $p < .001$, TLI = .94, CFI = .95, RMSEA = .08). Internal consistency was assessed via Cronbach's alpha (α), which exceeded.70 for all measures (cf. Table 1).

3.3 Hypothesis Testing

Descriptive statistics and intercorrelations among studied variables are presented in Table 1. Regarding control variables, team tenure was negatively and significantly related to team member satisfaction ($r = -.22$, $p < .01$), to teamwork management ($r = -.26$, $p < .01$), as well as to team innovation ($r = -.17$, $p < .01$). Teams collected during the pandemic presented higher levels of team virtuality ($r = .56$, $p < .01$), teamwork management ($r = .22$, $p < .01$), team member satisfaction ($r = .11$, $p < .05$) and team

innovation ($r = .14$, $p < .01$). Team size was negatively and significantly related to team innovation ($r = -.32$, $p < .01$), but not with teamwork management and team member satisfaction. Therefore, team size was only included as a control variable in the estimation of team innovation, while team tenure. Team virtuality was positively related to teamwork management ($r = .28$, $p < .01$) and to team member satisfaction ($r = .10$, $p < .05$) while the correlation with team innovation was not significant ($r = -.05$, *ns*). Teamwork management was positively related to team innovation ($r = .22$, $p < .01$) and to team member satisfaction ($r = .33$, $p < .01$).

Table 1. Descriptive statistics and intercorrelation matrix.

Variable	1	2	3	4	5	6	7	*M*	*SD*
1. Team tenure	–							8.72	10.76
2. Team size	−.04							8.53	7.72
3. Pandemic situation	.03	.03						–	–
4. Team virtuality	−.15**	.04	.56**					44.45	16.36
5. Teamwork management	−.26**	−.01	.22**	.28**	(.94)			3.95	.50
6. Team innovation	−.17**	−.32**	.14**	−.05	.22**	(.89)		5.04	1.09
7. Team member satisfaction	−.22**	.03	.11*	.10*	.33**	.15**	(.94)	5.60	1.00

Note: N = 358 members/88 teams. *p < .05; **p < .01. *M* = Mean; *SD* = Standard Deviation; Pandemic situation (0 = before the pandemic; 1 = after the beginning of the pandemic). Team member satisfaction and team innovation were collected two months after the remaining variables. Cronbach's alphas are provided between brackets.

Data for this study is hierarchical in nature since we are interested in exploring the effect of two variables that are operationalized at the team level, team virtuality and teamwork management, on one variable from the individual level, team member satisfaction and one variable from the team level, team innovation.

Because individuals within the same group are exposed to shared influences, applying ordinary regression analysis may produce unreliable estimates; this is because the assumption of independent observations would be violated (Bryk & Raudenbush, 1992). Therefore, we need to explore the need to consider the influence of the group on team member satisfaction, by computing the intraclass correlation coefficient (Bliese, 2000). Therefore, the initial step involves decomposing the variance in team member satisfaction (σ^2) into within-group and between-group components. If the results show that only a small proportion of the variance (typically less than 5%) lies between groups, then ordinary regression analysis is appropriate (Heck et al., 2013). *ICC* is computed as

follows:

$$ICC = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_w^2},$$

where σ_b^2 is the between-group variance and σ_w^2 is the within-group variance. The ICC value confirmed that 34% of the variance in member's satisfaction is explained by team membership, suggesting the need to use Hierarchical Linear Modelling to test hypotheses. Thus, we tested our hypothetical model through generalized structural equation modelling (GSEM) with Stata (version 16), including a random intercept in the equation predicting team member satisfaction to account for group membership. Additionally, variables were grand-mean centred, except for the outcome variables, a suitable method for mediational models with a nested structure (Kreft et al., 1995).

Results of GSEM analysis are presented in Table 2. Team virtuality presented a positive relationship with teamwork management ($B = .005$, $p < .01$). Teamwork management presented a positive effect on both team member satisfaction ($B = .54$, $p < .001$) and team innovation ($B = .41$, $p < .001$). To test hypothesis 1 and 2, we computed the indirect effect of team virtuality on team member satisfaction and team innovation, respectively. Results for team member satisfaction revealed that the indirect effect was significant (B = .003, $p < .05$, 95% CI = .0003; .005) providing support for hypothesis 1. Regarding team innovation, the indirect effect was also significant supporting H2 (B = .002, $p < .05$, 95% CI = .0002; .004).

Table 2. Generalized structural equation modelling analysis.

Predictor variables	Teamwork management	Member's satisfaction	Team innovation
Team tenure	−.01***	.01*	−.02***
Pandemic	.14*	.08	.55***
Team size			−.05***
Team virtuality	.005**	−.00	−.02***
Teamwork management		.54***	.41***
var(M1[Group])	.21		
var(e.teamwork management)	.22		
var(e.team innovation)	.95		
var(e.team member satisfaction)	.68		
Log-likelihood	−1197.54		

Note. The random intercept was included in the path analysis. *p < .05; **p < .01.; ***p < .001.

4 Discussion and Conclusions

In the present study, we examined how team virtuality influences team effectiveness, specifically, team member satisfaction and team innovation, through its impact on teamwork management processes. We draw on Conservation of Resources (COR) theory (Hobfoll, 1989; Hobfoll et al., 2018), which argues that individuals are motivated to protect, retain, and invest resources that enable them to function effectively at work. Virtuality constitutes a dual-edged condition in this framework: it generates resource gains (e.g., increased flexibility, broader access to information, reduced spatial constraints), but simultaneously introduces resource loss risks (e.g., reduced social cue richness, heightened ambiguity, greater coordination demands) (Carter et al., 2019; Zhang et al., 2024). COR theory posits that when individuals perceive potential or actual loss, they engage in compensatory resource-investment strategies to counteract those losses and secure future gains (Halbesleben et al., 2014).

Applying this logic to virtual teamwork, we argued that higher levels of virtuality would prompt team members to invest more heavily in teamwork management behaviours, such as work planning, resource coordination, clarifying responsibilities, and providing mutual support, to protect themselves from the potential depletion of cognitive and relational resources that virtual work can impose. This reasoning is consistent with broader team-effectiveness theorizing, which identifies coordination, communication structuring, and supportive behaviours as foundational processes that strengthen collective functioning in complex or uncertain environments (Mathieu et al., 2008; Rousseau et al., 2006).

Our findings support these predictions: as virtuality increased, so did the extent to which teams engaged in teamwork management behaviours, which acted as a critical mechanism linking virtuality to team outcomes. By structuring work, allocating resources efficiently, and maintaining shared support systems, teams were able to reduce ambiguity and coordination strain, conditions known to undermine satisfaction and impede creative work (Hülsheger et al., 2009; Mathieu et al., 2008). At the same time, effective teamwork management freed up cognitive resources and fostered psychological safety, both of which enable exploration, idea generation, and innovation implementation (Edmondson, 1999) and create a positive climate that, in turn, will contribute to increasing the levels of team member satisfaction.

Summing up, our study advances the literature on team virtuality by shedding light on the underlying mechanisms through which virtuality shapes team effectiveness (i.e., team member satisfaction and team innovation). Although previous research has examined the outcomes of virtuality, the pathways through which these effects occur have received far less attention. To the best of our knowledge, this is the first empirical study to test the indirect influence of team virtuality on multiple dimensions of team effectiveness via teamwork management.

Our results highlight that teamwork management constitutes a central process through which teams can buffer the challenges and leverage the opportunities associated with virtual collaboration. Effective planning, clear structuring of tasks, strategic allocation of resources, and mutual support are not merely beneficial, they are foundational for sustaining high levels of satisfaction among members and fostering the conditions that allow innovative ideas to emerge and be implemented (Rousseau et al., 2006).

From a practical standpoint, the findings highlight the importance of equipping managers and team leaders with the skills needed to coach and support teams working at higher levels of virtuality. Specifically, managers should help teams develop strong routines for planning and organizing their work, clarifying responsibilities, and monitoring progress. Such practices become especially crucial when face-to-face interaction is limited, as they provide the structure and predictability that virtual teams need to maintain engagement, collaborate effectively, and innovate.

The results of this study should be considered with a number of limitations. Although we used a two-wave data collection procedure, team virtuality and teamwork management were collected at the same time, despite the fact that we are testing a mediational process. However, it is important to highlight that team virtuality was assessed retrospectively, based on the preceding month's data, whereas the time reference for teamwork management was the current time. This temporal distinction helps to partially mitigate concerns about causal order. Secondly, the limited sample size at the team level ($n = 88$) may impact the statistical power and generalizability of the results. Future studies should aim to replicate these findings with larger samples. It would also be interesting to consider some moderators, such as emotional management or trust.

Acknowledgments. This work has been supported by the Fundação para a Ciência e a Tecnologia (FCT) under project grant UIDB/05037/2025.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

Batarseh, F.S., Usher, J.M., Daspit, J.J.: Collaboration capability in virtual teams: examining the influence on diversity and innovation. Int. J. Innov. Manag. **21**(04), 1750034 (2017)

Bliese, P.D.: Within-group agreement, non-independence, and reliability: implications for data aggregation and analysis. In: J., K., Klein, K.S.W.J. (eds.) Multilevel Theory, Research, and Methods in Organizations: Foundations, Extensions, and New Directions, pp. 349–381. Jossey-Bass (2000)

Bryk, A.S., Raudenbush, S.W.: Hierarchical Linear Models: Applications and Data Analysis Methods. Sage Publications, Inc (1992)

Carter, K.M., Mead, B.A., Stewart, G.L., Nielsen, J.D., Solimeo, S.L.: Reviewing work team design characteristics across industries: combining meta-analysis and comprehensive synthesis. Small Group Res. **50**(1), 138–188 (2019)

Chan, D.: Functional relations among constructs in the same content domain at different levels of analysis: a typology of composition models. J. Appl. Psychol. **83**, 234–246 (1998)

Cousins, K., Robey, D., I.: Zigurs.: managing strategic contradictions in hybrid teams. Eur. J. Inf. Syst. **16**(4), 460–478 (2007)

Dimas, I.D., Lourenço, P.R.: Intragroup conflict and conflict management approaches as determinants of team performance and satisfaction: two field studies. Negot. Confl. Manag. Res. **8**(3), 174–193 (2015)

Dixon, K.R., Panteli, N.: From virtual teams to virtuality in teams. Hum. Relat. **63**(8), 1177–1197 (2010)

Edmondson, A.: Psychological safety and learning behavior in work teams. Adm. Sci. Q. **44**(2), 350–383 (1999)

George, G., Lakhani, K.R., Puranam, P.: What has changed? The impact of covid pandemic on the technology and innovation management research agenda. J. Manag. Stud. **57**(8), 1754 (2020)

Gibson, C.B., Gilson, L.L., Griffith, T.L., O'Neill, T.A.: Should employees be required to return to the office? Organ. Dyn. **52**(2), 100981 (2023)

Hackman, J.R.: The design of work teams. In: Lorsch, J. (ed.) Handbook of Organizational Behavior, pp. 315–342. Prentice-Hall (1987)

Handke, L., Costa, P., Feitosa, J.: How virtual are we? Introducing the team perceived Virtuality scale. J. Bus. Psychol. **40**, 455–478 (2025)

Handke, L., Costa, P.L., Hincapie, M.X., Johnson, M.D.: Not even remotely close: how co-location imbalance affects subgroup formation in hybrid teams. J. Organ. Behav. **47**(2), 192–207 (2026)

Heck, R.H., Thomas, S.L., Tabata, L.N.: Multilevel and Longitudinal Modeling with IBM SPSS. Routledge, New York (2013)

Halbesleben, J.R., Neveu, J.P., Paustian-Underdahl, S.C., Westman, M.: Getting to the "COR" understanding the role of resources in conservation of resources theory. J. Manag. **40**(5), 1334–1364 (2014)

Hobfoll, S.E.: Conservation of resources: a new attempt at conceptualizing stress. Am. Psychol. **44**(3), 513–524 (1989)

Hobfoll, S.E., Halbesleben, J., Neveu, J.P., Westman, M.: Conservation of resources in the organizational context: the reality of resources and their consequences. Annu. Rev. Organ. Psych. Organ. Behav. **5**(1), 103–128 (2018)

Hülsheger, U.R., Anderson, N., Salgado, J.F.: Team-level predictors of innovation at work: a comprehensive meta-analysis spanning three decades of research. J. Appl. Psychol. **94**(5), 1128–1145 (2009)

James, L.R., Demaree, R.G., Wolf, G.: r_{wg}: an assessment of within-group interrater agreement. J. Appl. Psychol. **78**(2), 306–309 (1993)

Kreft, I.G., de Leeuw, J., Aiken, L.S.: The effect of different forms of centering in hierarchical linear models. Multivar. Behav. Res. **30**, 1–21 (1995)

Mahadevan, J.: Virtual Team Collaboration. Springer, A Guide for Individual Team Members/by Jasmin Mahadevan (2024)

Mathieu, J., Maynard, M.T., Rapp, T., Gilson, L.: Team effectiveness 1997-2007: a review of recent advancements and a glimpse into the future. J. Manag. **34**(3), 410–476 (2008)

Podsakoff, P.M., MacKenzie, S.B., Lee, J.-Y., Podsakoff, N.P.: Common method biases in behavioral research: a critical review of the literature and recommended remedies. J. Appl. Psychol. **88**(5), 879–903 (2003)

Purvanova, R.K., Kenda, R.: The impact of virtuality on team effectiveness in organizational and non-organizational teams: a meta-analysis. Appl. Psychol. **71**(3), 1082–1131 (2022)

Rousseau, V., Aubé, C., Savoie, A.: Le fonctionnement interne des équipes de travail: Conception et mesure. Can. J. Behav. Sci. / Revue Canadienne Des Sci. Du Comportement **38**(2), 120–135 (2006)

Standifer, R.L., Raes, A.M.L., Peus, C., Passos, A.M., Santos, C.M., Weisweiler, S.: Time in teams: cognitions, conflict, and team satisfaction. J. Manag. Psychol. **30**(6), 692–708 (2015)

Vera, D., Crossan, M.: Improvisation and innovative performance in teams. Organ. Sci. **16**(3), 203–224 (2005)

West, M.A., Anderson, N.R.: Innovation in top management teams. J. Appl. Psychol. **81**(6), 680–693 (1996)

Zheng, S., Yan, M., Liang, Y., Chen, Y., Wei, Q., Li, S.: Understanding the positive and negative effects of team virtuality: A theoretical review and research agenda. Hum. Resour. Manag. Rev. **34**(2), 101013 (2024)

A Daily-Grid Hazards-to-Dynamics Framework for Inpatient Flow

Cecilia Castro[1(✉)], Ana Jacinta Soares[1], Ana Paula Amorim[1], and Alexandre Magalhães[2]

[1] Centre of Mathematics, Universidade do Minho, Braga, Portugal
{cecilia,ajsoares,apamorim}@math.uminho.pt
[2] Neuro & Skeletal Circuits Group, INEB/i3S, Universidade do Porto, Porto, Portugal
alexandre.magalhaes@ineb.up.pt

Abstract. Routine hospital episode data record admissions and exits at date resolution, together with exit type and a limited set of baseline covariates. This paper develops a reproducible daily-grid method for deriving inpatient occupancy and cause-specific outflow series from such data. The statistical layer estimates grouped day-of-stay exit probabilities for discharge alive and in-hospital death, and the dynamical layer propagates these probabilities through a conservative cohort update under a half-open counting convention. The formulation separates directly observed episode-based quantities from probability tables and long-stay summaries that depend on pooling, completion, or tail closure. It also links the discrete construction to a McKendrick–von Foerster transport equation as a compatible continuous-time analogue, with exact accounting established for the grouped daily scheme used in practice. The resulting outputs provide cohort-level daily flow inputs for retrospective audits, comparison of operating regimes, and scenario analysis; when capacity data are available, the same series can be combined with them to study operational bed pressure. In an application to Brazilian administrative records for femur-fracture admissions, the reference sample contained 442,646 episodes after exclusion of 5,106 same-day admission-and-exit records. As an internal reconstruction check on this retained closed-episode cohort, the model-implied occupancy closely followed the empirical census, with a mean absolute error of 53.03, or 3.91% of mean empirical occupancy. Mass-balance residuals were at numerical roundoff level, confirming conservation of the cohort update. Sensitivity analyses over support length, sparse-cell pooling, tail handling, and same-day policy yielded stable reconstruction results.

Keywords: Administrative hospital data · Age-structured transport model · Competing risks · Conservative daily cohort discretization · Daily-grid identifiability · Inpatient flow modelling

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 192–209, 2027.
https://doi.org/10.1007/978-3-032-30518-3_13

1 Introduction

Hospitals make daily decisions about beds and patient flow under uncertainty, while routine administrative datasets record each stay as an episode defined by admission date, exit date, exit type, and a small set of admission-time covariates [1]. Without timestamps for intermediate clinical transitions, inpatient flow is reconstructed on a daily grid under explicit measurement and modelling conventions.

Admissions, length of stay, occupancy, and exits evolve jointly and have been studied through queueing, simulation, and forecasting models [2–4]. For date-recorded data, discharge alive and in-hospital death naturally define a grouped competing-risks problem [5,6], while age-of-stay-structured transport models provide a system-level description in which conservative discretizations preserve accounting links among admissions, occupancy, and outflows [7,8]. These strands provide complementary tools for capacity planning, aggregate prediction, and patient-level event modelling; this paper connects them through a daily-grid interface that maps grouped exit probabilities into an accounting-conservative inpatient population update [12].

This paper introduces a unified daily-grid framework that maps grouped competing-risks probabilities from date-recorded episodes into conservative inpatient flow dynamics and daily occupancy and outflow series. Its central identifiability distinction is that date-only episodes determine daily admissions, recorded exits, and half-open occupancy on the retained sample, whereas sub-daily exit times, intermediate clinical states, continuous-time hazard surfaces, sparse cells, and long-stay tails require modelling, pooling, completion, or tail-construction choices.

Admission-based and calendar-time formulations are handled within a common structure, clarifying how regime information enters the daily exit probabilities and their propagation through the cohort system. A McKendrick–von Foerster equation is included as a compatible continuous-time analogue of the grouped daily mechanism [7,9,10], while exact accounting results are established for the discrete formulation used in practice.

The applied analysis uses Brazilian administrative hospital records from the Hospital Information System of the Brazilian Unified Health System (SIH/SUS), accessed through the Department of Informatics of the Unified Health System (DATASUS), restricted to International Statistical Classification of Diseases and Related Health Problems, 10th Revision (ICD-10) code S72 hospitalizations for femur fractures [1,11]. The extract is used for reconstruction diagnostics and sensitivity checks on real episode data, focusing on internal consistency of the reconstructed daily series.

The remainder of the article is organized as follows. Section 2 presents the methodological framework. Section 3 describes the episode schema and computational workflow. Section 4 reports the SIH/SUS ICD-10 S72 extract, reconstruction diagnostics, and sensitivity checks. Section 5 summarizes the findings and possible extensions.

2 Methodological Framework

This section defines the daily-grid framework: retained episodes are converted into empirical daily series, grouped day-of-stay exit probabilities are estimated, and these probabilities are propagated through a conservative cohort update on the calendar grid.

2.1 Observed Episodes and Daily Series

We work with a retained analysis table containing one hospitalization episode per row. Unless otherwise stated, the development refers to the fully observed retained sample after preprocessing; same-day episodes and extraction-window censoring are discussed in Sect. 3.2. For episode $i = 1, \ldots, N$, we observe the admission day s_i, exit day e_i, exit type $c_i \in \{D, M\}$, where D denotes discharge alive and M denotes in-hospital death, and admission-time variables $\boldsymbol{x}_i$, used for stratification or optional modelling. Time is recorded on a daily grid. Under the half-open convention, episode i is present on day t if and only if $s_i \le t < e_i$, and its recorded day-resolution length of stay is $\ell_i := e_i - s_i$. This is a grouped quantity induced by recorded dates, not a directly observed continuous elapsed time. Same-day admission and exit give $\ell_i = 0$; these records are excluded from the default retained sample and handled in sensitivity analysis.

For stratified analyses, let g map admission-time variables to $\{1, \ldots, J\}$, and write $j(i) := g(\boldsymbol{x}_i)$. With $\mathbf{1}\{\cdot\}$ denoting the indicator function, the retained episode table directly determines the empirical daily series used for estimation and checking. Admissions into stratum j are

$$A_j(t) := \sum_{i=1}^{N} \mathbf{1}\{s_i = t,\ j(i) = j\},\ A(t) := \sum_{j=1}^{J} A_j(t). \tag{1}$$

Empirical cause-specific exits are

$$\widehat{D}_{\mathrm{emp},j}(t) := \sum_{i=1}^{N} \mathbf{1}\{j(i) = j,\ c_i = D,\ e_i = t\},$$

$$\widehat{M}_{\mathrm{emp},j}(t) := \sum_{i=1}^{N} \mathbf{1}\{j(i) = j,\ c_i = M,\ e_i = t\}.$$

The corresponding totals are

$$\widehat{D}_{\mathrm{emp}}(t) := \sum_{j=1}^{J} \widehat{D}_{\mathrm{emp},j}(t),\ \widehat{M}_{\mathrm{emp}}(t) := \sum_{j=1}^{J} \widehat{M}_{\mathrm{emp},j}(t).$$

Empirical occupancy is

$$\widehat{H}_{\mathrm{emp}}(t) := \sum_{i=1}^{N} \mathbf{1}\{s_i \le t < e_i\}. \tag{2}$$

These series are fully determined by the retained episode table and provide the empirical basis for probability estimation and reconstruction diagnostics.

2.2 Competing-Risks Quantities and Daily Estimands

Let S_{adm}, L, C, and $\boldsymbol{X}$ denote the random counterparts of s_i, ℓ_i, c_i, and $\boldsymbol{x}_i$. On the retained sample, $L \in \{1, 2, \dots\}$ is the grouped day-resolution length of stay. When regime stratification is used, calendar time is partitioned by a deterministic map $P(t) \in \{1, \dots, R\}$. Two conventions are distinguished. Under the admission-based convention, an episode is assigned the regime in force at admission, $P(S_{\mathrm{adm}})$. Under the calendar-time convention, an episode admitted on day s is exposed on grouped day of stay m to the regime in force on calendar day $s + m$. The primary estimands are one-day cause-specific exit probabilities on the grouped day-of-stay scale. Under the admission-based convention,

$$q_{k,j,r}(m) := \Pr\big(L = m+1,\, C = k \,\big|\, L > m,\, g(\boldsymbol{X}) = j,\, P(S_{\mathrm{adm}}) = r\big), \tag{3}$$

for $k \in \{D, M\}$ and $m \in \mathbb{N}_0$. Thus, $q_{k,j,r}(m)$ is the probability of exit by cause k at the end of grouped day m, among patients still present at the start of that grouped day. Under the calendar-time convention, the fine daily-grid probabilities are indexed by admission day:

$$q_{k,j,s}(m) := \Pr\big(L = m+1,\, C = k \,\big|\, L > m,\, g(\boldsymbol{X}) = j,\, S_{\mathrm{adm}} = s\big), \tag{4}$$

with the contemporaneous regime determined by $P(s + m)$. Regime-level calendar-time summaries may be obtained by risk-set-consistent pooling or smoothing over cells sharing the same contemporaneous regime, but their use as transportable inputs requires the assumption that, conditional on stratum and grouped day of stay, calendar time affects the exit mechanism only through that regime label. Continuous-time notation is used only as a compatible embedding of the grouped daily mechanism [5,12]. Under date-only observation, within-day exit times and continuous-time hazards are not identified. Suppressing indices, let $q(m) := \sum_k q_k(m)$. If $0 \le q(m) < 1$, one constant-within-day embedding sets

$$\Lambda(m) := -\log\{1 - q(m)\},\ \Lambda_k(m) = \begin{cases} \Lambda(m)\dfrac{q_k(m)}{q(m)}, & q(m) > 0, \\ 0, & q(m) = 0. \end{cases}$$

Under the constant-within-day embedding used here, or more generally under a within-day embedding with constant cause proportions, if cause-specific integrated hazards $\Lambda_k(m)$ are specified first and $\Lambda(m) := \sum_{k \in \{D,M\}} \Lambda_k(m)$, the induced daily probabilities are

$$q(m) = 1 - \exp\{-\Lambda(m)\},\ q_k(m) = \big(1 - \exp\{-\Lambda(m)\}\big)\frac{\Lambda_k(m)}{\Lambda(m)}, \quad k \in \{D, M\}, \tag{5}$$

with the ratio defined as 0 when $\Lambda(m) = 0$. These mappings are compatible representations of the observed daily probabilities, not additional empirical information. If $q(m) = 1$, the daily probability model remains well defined, but a finite continuous-time hazard embedding is singular.

2.3 Statistical Layer: Operational Estimands and Estimation Choices

The estimands in Sect. 2.2 enter the numerical framework through daily cause-specific exit-probability tables on the grouped day-of-stay support used by the cohort dynamics. The implementation is probability-based; integrated hazards are used only for reporting or for a compatible continuous-time embedding.

For retained episodes with $\ell_i \geq 1$, the grouped day-resolution length of stay is observed from the admission and exit dates. Under the admission-based convention, define $q_{j,r}(u) := \sum_{k\in\{D,M\}} q_{k,j,r}(u)$. For an episode i in stratum $j(i)$, admitted under regime $P(s_i)$, with length of stay ℓ_i and exit cause c_i, the grouped-data likelihood contribution is

$$\mathcal{L}_i = \left[\prod_{u=0}^{\ell_i-2} \big(1 - q_{j(i),P(s_i)}(u)\big)\right] q_{c_i,j(i),P(s_i)}\big(\ell_i - 1\big),$$

with the empty product interpreted as 1 when $\ell_i = 1$.

Under the calendar-time convention, the fine daily-grid likelihood is written in terms of admission-day-indexed probabilities.
With $q_{j,s}(u) := \sum_{k\in\{D,M\}} q_{k,j,s}(u)$, the corresponding contribution is

$$\mathcal{L}_i = \left[\prod_{u=0}^{\ell_i-2} \big(1 - q_{j(i),s_i}(u)\big)\right] q_{c_i,j(i),s_i}\big(\ell_i - 1\big).$$

A regime-pooled calendar-time implementation replaces the fine admission-day-indexed terms by pooled terms indexed by the contemporaneous regime: $P(s_i+u)$ in the survival factors and $P(s_i + \ell_i - 1)$ in the terminal factor. As noted in Sect. 2.2, this reduced representation requires a pooling or modelling assumption beyond direct cellwise identifiability.

The primary implementation uses the nonparametric daily estimators in Sect. 3.3, which match the date resolution of the data and the grouped update in Sect. 2.5. Alternative grouped-time, smoothed, piecewise-constant, or proportional-hazards estimators may also be used, provided that the dynamical layer receives daily probability tables on the chosen working grid.

If fitting is carried out at a finer covariate level, fitted probabilities must be aggregated before simulation. Under the admission-based convention, aggregation is over survivors sharing stratum and admission regime.

Under the calendar-time convention, aggregation must respect the surviving episode-day composition on the calendar-time by day-of-stay grid before any coarser regime-level pooling is applied. After this risk-set-consistent aggregation, the statistical layer is closed and the grouped cohort update takes the completed probability tables as fixed inputs.

2.4 Dynamical Layer: Continuous-Time Transport Analogue

The daily probability model admits a compatible continuous-time age-of-stay embedding in the McKendrick–von Foerster framework, but this embedding does not identify sub-daily inpatient dynamics from date-recorded data [7,9,10,13].

For each stratum $j \in \{1, \dots, J\}$, let $h_j(t, a) \geq 0$ denote the age-of-stay density at calendar time t and elapsed stay $a \geq 0$. Thus $h_j(t, a)\, da$ represents the number of inpatients in stratum j with elapsed stay in $[a, a + da)$.

After relabelling the daily calendar grid as consecutive integers, write $t^n = n$. Let $a_j(t) \geq 0$ be a boundary admission-rate embedding satisfying

$$\int_{t^n}^{t^{n+1}} a_j(u)\, du = A_j(t^{n+1}),\ n \in \mathbb{Z}.$$

The within-day shape of $a_j(t)$ is not identified by the data; only the daily admission totals are observed. The integral convention is chosen to match the daily boundary update used in the cohort scheme.

Let $\lambda_{k,j}(a, t) \geq 0$, $k \in \{D, M\}$, denote an aggregated cause-specific hazard surface compatible with the daily probability table. With regime labels, the calendar-time convention may be represented as $\lambda_{k,j}(a, t) = \lambda_{k,j,P(t)}(a)$, whereas the admission-based convention corresponds, along characteristics, to $\lambda_{k,j}(a, t) = \lambda_{k,j,P(t-a)}(a)$. If finer calendar-time indexing is retained, $\lambda_{k,j}(a, t)$ is interpreted as the corresponding cell-specific surface.

The continuous-time analogue is

$$\partial_t h_j(t, a) + \partial_a h_j(t, a) = -\big(\lambda_{D,j}(a, t) + \lambda_{M,j}(a, t)\big) h_j(t, a),\ t > 0,\ a > 0, \quad (6)$$

with boundary condition

$$h_j(t, 0) = a_j(t), \quad (7)$$

and initial condition $h_j(0, a) = h_{j,0}(a) \geq 0$. Exact accounting results in this paper are stated for the grouped daily scheme of Sect. 2.5; the continuous-time formulation is used only as a compatible analogue.

Define

$$H_j(t) = \int_0^\infty h_j(t, a)\, da,$$

and

$$D_j(t) = \int_0^\infty \lambda_{D,j}(a, t) h_j(t, a)\, da,\ M_j(t) = \int_0^\infty \lambda_{M,j}(a, t) h_j(t, a)\, da.$$

The corresponding totals are

$$H(t) := \sum_{j=1}^{J} H_j(t),\ D(t) := \sum_{j=1}^{J} D_j(t),\ M(t) := \sum_{j=1}^{J} M_j(t).$$

Proposition 1 (Continuous-Time Mass Balance). *Assume sufficient regularity and integrability so that $h_j(t, \cdot)$, $\lambda_{D,j}(\cdot, t) h_j(t, \cdot)$, and $\lambda_{M,j}(\cdot, t) h_j(t, \cdot)$ are integrable on $(0, \infty)$, with $h_j(t, a) \to 0$ as $a \to \infty$. Then, for each stratum j,*

$$\frac{d}{dt} H_j(t) = a_j(t) - D_j(t) - M_j(t)$$

for almost every $t \geq 0$.

Proof sketch. Integrate (6) over $a \in (0, \infty)$, use the boundary condition (7), and let the boundary term at $a \to \infty$ vanish under the stated assumptions.

2.5 Conservative Daily Cohort Discretization

The grouped daily exit mechanism induces a conservative cohort update on the same daily grid as the data. After relabelling the grid as consecutive integers, let $\Delta = 1$ day, $t^n := n\Delta$, $a^m := m\Delta$. For each stratum j, let $N_j^{n,m}$ denote the number, or expected number, of patients present on day t^n in grouped day-of-stay class m. This is a start-of-day presence state under the half-open convention: ordinary cells $m = 0, \ldots, m_{\max} - 1$ correspond to episodes admitted on day t^{n-m} and still present on day t^n. Admissions dated t^{n+1} enter through the boundary state $N_j^{n+1,0}$.

Let $\mathcal{M}_{\text{ret}}$ be the retained set of state indices. With a single retained tail cell, $\mathcal{M}_{\text{ret}} = \{0, \ldots, m_{\max}\}$, where $m = m_{\max}$ groups patients with day of stay at least $m_{\max}$. With truncation, $\mathcal{M}_{\text{ret}} = \{0, \ldots, m_{\max} - 1\}$, and survivor mass leaving the retained support is recorded as truncation loss. Grid-respecting retained tails are obtained by applying the same survivor-shift logic across additional tail cells. For grid-respecting retained tail cells that keep a unique grouped day of stay, the same probability-assignment principle applies to those retained cells. Occupancies are

$$H_j^n := \sum_{m \in \mathcal{M}_{\text{ret}}} N_j^{n,m}, \; H^n := \sum_{j=1}^{J} H_j^n.$$

When the retained state represents the intended grid exactly, H^n is directly comparable with $\widehat{H}_{\text{emp}}(t^n)$; under a reduced-tail closure, the comparison is a reduced-state reconstruction.

For each cause $k \in \{D, M\}$, let $q_{k,j}^{n,m} \in [0,1]$ be the one-day exit probability applied to cell (j,m) over $(t^n, t^{n+1}]$, and define $q_j^{n,m} := q_{D,j}^{n,m} + q_{M,j}^{n,m}$. On ordinary cells $m = 0, \ldots, m_{\max} - 1$, the admission day is uniquely determined by (n,m), namely t^{n-m}. Hence, the probabilities are assigned according to the working representation

$$q_{k,j}^{n,m} = \begin{cases} q_{k,j,t^{n-m}}(m), & \text{calendar-time,} \\ & \text{admission-day indexed,} \\ q_{k,j,P(t^n)}(m), & \text{calendar-time, regime-pooled,} \\ q_{k,j,P(t^{n-m})}(m), & \text{admission-based,} \end{cases} \quad m = 0, \ldots, m_{\max} - 1.$$

In the single-tail implementation, $N_j^{n,m_{\max}}$ aggregates all patients whose grouped day of stay is at least $m_{\max}$. Therefore $(n, m_{\max})$ does not determine a unique admission day, admission regime, or admission-day index. The tail probability is therefore a closure probability for the aggregated long-stay state. We write $q_{k,j}^{n,m_{\max}} := q_{k,j}^{n,\text{tail}}$, $q_j^{n,m_{\max}} := q_{D,j}^{n,\text{tail}} + q_{M,j}^{n,\text{tail}}$, $q_j^{n,m_{\max}} \leq 1$.

For a calendar-time regime-pooled tail closure, $q_{k,j}^{n,\text{tail}} = q_{k,j,P(t^n)}^{\text{tail}}$, because the contemporaneous regime is common to all patients in the tail at calendar day t^n. For admission-based or admission-day-indexed calendar-time representations, an unindexed single-tail state is compatible only with a marginal or composition-averaged tail closure.

Exact retention of admission-regime-specific or admission-day-specific long-stay effects requires either a grid-respecting tail representation or an augmented tail state carrying the relevant index.

For ordinary interior cells, $N_j^{n+1,m+1} = N_j^{n,m}\big(1-q_j^{n,m}\big),\ m = 0,\dots,m_{\max}-2$, and the boundary update is $N_j^{n+1,0} = A_j(t^{n+1})$, where $A_j(t)$ is defined in (1). In the single-tail case,

$$N_j^{n+1,m_{\max}} = N_j^{n,m_{\max}-1}\big(1-q_j^{n,m_{\max}-1}\big) + N_j^{n,m_{\max}}\big(1-q_j^{n,m_{\max}}\big).$$

This update is conservative for the implemented reduced state; exact recovery of all day-specific long-stay behaviour requires a grid-respecting tail representation or an equivalent exact tail construction.

The cause-specific removals over $(t^n, t^{n+1}]$ are

$$\widetilde{D}_j^n := \sum_{m\in\mathcal{M}_{\mathrm{ret}}} N_j^{n,m} q_{D,j}^{n,m},\ \widetilde{M}_j^n := \sum_{m\in\mathcal{M}_{\mathrm{ret}}} N_j^{n,m} q_{M,j}^{n,m}.$$

Thus, H^n is occupancy on day t^n, while $\widetilde{D}_j^n$ and $\widetilde{M}_j^n$ have exit date t^{n+1}.

At the first analysis day t^0, the state may be initialized directly from the episode table:

$$N_j^{0,m} = \sum_{i=1}^{N} \mathbf{1}\{j(i)=j,\ s_i \le t^0 < e_i,\ t^0 - s_i = m\},\ m = 0,\dots,m_{\max}-1.$$

If a single tail cell is retained, then

$$N_j^{0,m_{\max}} = \sum_{i=1}^{N} \mathbf{1}\{j(i)=j,\ s_i \le t^0 < e_i,\ t^0 - s_i \ge m_{\max}\}.$$

Grid-respecting tails are initialized analogously. If a pre-window warm start is used instead, the initialization choice and affected days should be reported.

Proposition 2 (Discrete Positivity and Mass Balance). *Assume that the initial state is nonnegative, $A_j(t^n) \ge 0$, and*

$$q_{D,j}^{n,m} \ge 0,\ q_{M,j}^{n,m} \ge 0,\ q_{D,j}^{n,m} + q_{M,j}^{n,m} \le 1$$

for all retained state components. Then all retained state components remain nonnegative.

If the retained state is closed, so that every non-exiting survivor is reassigned to a retained component, then, for each stratum j and step $n \ge 0$,

$$H_j^{n+1} - H_j^n = A_j(t^{n+1}) - \widetilde{D}_j^n - \widetilde{M}_j^n. \tag{8}$$

If the scheme is truncated at the last retained ordinary cell, then

$$H_j^{n+1} - H_j^n = A_j(t^{n+1}) - \widetilde{D}_j^n - \widetilde{M}_j^n - U_j^n, \quad \textit{where} \tag{9}$$

$$U_j^n := N_j^{n,m_{\max}-1}\big(1-q_j^{n,m_{\max}-1}\big)$$

is the truncation loss over $(t^n, t^{n+1}]$.

Proof. Nonnegativity follows because each update is a nonnegative combination of nonnegative state components and the nonnegative boundary input $A_j(t^{n+1})$. Summing the update equations over retained components gives the mass balance: in the closed case, survivor transfers cancel internally and only admissions and cause-specific removals remain. Under truncation, the uncancelled survivor flow is U_j^n, yielding (9).

Equation (8) is the accounting identity used in diagnostics for closed retained-tail implementations, including grid-respecting tails and conservative reduced closures. For truncated implementations, the adjusted balance (9) must be used. Because reduced-tail rows are closure quantities, applications should report sensitivity to $m_{\max}$ and retained-tail treatment.

2.6 Identifiability with Minimal Episode Variables

The framework distinguishes quantities determined directly by the retained episode table from quantities requiring pooling, completion, tail closure, or modelling assumptions. Given $(s_i, e_i, c_i, \boldsymbol{x}_i)$, or equivalently $(s_i, \ell_i, c_i, \boldsymbol{x}_i)$ after preprocessing, and a specified regime map $P(\cdot)$ when used, the empirical series in (1)–(2) are reconstructed directly. Under the half-open convention, the retained data determine daily admissions, cause-specific exits, and occupancy on the retained sample.

On the grouped day-of-stay scale induced by recorded dates, the main nonparametrically estimable quantities are one-day cause-specific exit probabilities for cells with nonempty risk sets. Under the admission-based convention, these are the stratum- and regime-specific probabilities in (3). Under the calendar-time convention, the directly indexed daily-grid quantities are the admission-day-specific probabilities in (4), or the equivalent calendar-time by day-of-stay cells.

Regime-pooled calendar-time probabilities can be computed as descriptive pooled ratios, but their use as transportable inputs assumes that, conditional on stratum and grouped day of stay, calendar time affects exits only through the contemporaneous regime label. The same distinction applies to long-stay cells: ordinary cells with positive exposure support day-specific estimation, whereas retained tail cells are pooled closure representations of long-stay behaviour. Once daily probabilities are defined on a chosen support, grouped survival, cumulative incidence, and length-of-stay quantities follow algebraically. Let α denote the relevant index, with $\alpha = r$ for the admission-based convention and $\alpha = s$ for admission-day-indexed calendar time, and define

$$q_{j,\alpha}(u) := \sum_{k \in \{D,M\}} q_{k,j,\alpha}(u). \text{ Then}$$

$$S_{j,\alpha}(m) = \prod_{u=0}^{m-1} \big(1 - q_{j,\alpha}(u)\big), \; S_{j,\alpha}(0) = 1, \text{ and}$$

$$F_{k,j,\alpha}(m) = \sum_{u=0}^{m-1} S_{j,\alpha}(u) q_{k,j,\alpha}(u), \ k \in \{D, M\}.$$

If probabilities are available on the full support of L, then

$$\Pr(L = m+1 \mid j, \alpha) = S_{j,\alpha}(m) q_{j,\alpha}(m), \ \mathbb{E}[L \mid j, \alpha] = \sum_{m=0}^{\infty} S_{j,\alpha}(m),$$

whenever the series converges.

For the regime-pooled calendar-time representation, the regime index is generally path-dependent along the stay. Writing

$$q_{j,r}(u) := \sum_{k \in \{D,M\}} q_{k,j,r}(u),$$

the grouped survival function for an episode admitted on day s is

$$S_{j,s}^{\text{pool}}(m) = \prod_{u=0}^{m-1} \left\{1 - q_{j,P(s+u)}(u)\right\}, \ S_{j,s}^{\text{pool}}(0) = 1.$$

The corresponding cause-specific cumulative incidence is

$$F_{k,j,s}^{\text{pool}}(m) = \sum_{u=0}^{m-1} S_{j,s}^{\text{pool}}(u) q_{k,j,P(s+u)}(u), \ k \in \{D, M\}.$$

Thus, under regime-pooled calendar-time indexing, there is no single fixed regime for the whole stay unless $P(s+u)$ remains constant over the relevant grouped days. The corresponding admission-day-specific length-of-stay quantities are

$$\overset{\text{pool}}{\Pr}(L = m+1 \mid j, s) = S_{j,s}^{\text{pool}}(m) q_{j,P(s+m)}(m), \ \mathbb{E}^{\text{pool}}[L \mid j, s] = \sum_{m=0}^{\infty} S_{j,s}^{\text{pool}}(m),$$

whenever the series converges. With finite support, the tail of the distribution and the mean beyond that support are determined by the chosen completion or tail closure.

Date-recorded episodes identify grouped daily quantities only; the mapping in (5) and the boundary-rate embedding in Sect. 2.4 are compatible representations, not additional empirical information. Once the analysis cells are fixed, the framework delivers exit probabilities and flow quantities at that aggregation level; richer covariate or state detail requires enlarging the state space.

Because admissions, exits, and occupancy are reconstructed directly from the retained episode table, the model-implied series can be compared with their empirical counterparts under the same state definition. These comparisons are internal reconstruction checks, with differences interpreted in light of numerical roundoff, sparse-cell pooling, finite-support completion, and the chosen tail treatment.

3 Data and Computational Workflow

This section describes the applied-data workflow: episode construction, preprocessing, daily probability estimation, cohort simulation, diagnostics, and reproducible outputs.

3.1 Data Source and Episode Schema

The framework uses one preprocessed row per hospitalization episode, with admission date, exit date, exit type, and admission-time variables for stratification or optional modelling. The applied extract was built from Brazilian SIH/SUS records restricted to ICD-10 S72 femur-fracture hospitalizations [1,11]. Source-level preprocessing removed duplicate hospitalization authorization identifiers, standardized dates and variables, and retained additional fields for traceability and descriptive quality control. The default stratified analysis used admission date, exit date, exit type, age, and sex.

For reproducibility, each computational run stores the analysis configuration, software session information, cleaned episode table, empirical series, probability tables, simulation outputs, reconstruction metrics, support diagnostics, figures, and article-ready summary tables.

3.2 Preprocessing, Quality Control, and Empirical Series

Preprocessing defines the retained sample and aligns the episode table with the half-open daily-grid convention of Sect. 2. Recorded length of stay is treated as $\ell_i := e_i - s_i$. Quality control checks missing or invalid dates, invalid exit types, missing age or sex, impossible ages, and inconsistent chronology such as $e_i < s_i$. All restrictions are applied consistently to the empirical series and to the estimated probability tables.

The default implementation uses fully observed episodes with recorded exit date and exit type. If an extraction window induces right-censoring, the likelihood in Sect. 2.3 and the estimators in Sect. 3.3 require censoring-aware extensions. In the applied SIH/S72 table, no missing admission dates, exit dates, exit types, age, or sex values remained after source-level preprocessing. However, because the analysis conditions on observed closed episodes, late-window estimates should be interpreted as conditional on the observed closed-episode extract rather than as censoring-adjusted estimates if later exits were unavailable in the source files.

Same-day admission-and-exit records have $\ell_i = 0$ under date arithmetic. The default analysis excluded them. A sensitivity analysis recoded them as one-day stays by shifting the recorded exit date to the next calendar day, so that they contribute to day-0 occupancy and day-0 exit probabilities. Under this recoding, empirical admissions, occupancy, and outflows are recomputed; hence the diagnostic target differs from the default retained-sample target.

The applied analysis used 10 age-sex strata. The reference run used a deterministic regime map with cut points 2020-03-01 and 2022-01-01, defining pre,

COVID, and post periods, and adopted the calendar-time regime-pooled representation. Admission-based and admission-day-indexed calendar-time representations remain available as alternative working choices.

After preprocessing, empirical daily admissions, cause-specific exits, and occupancy are reconstructed directly from admission dates, exit dates, exit types, and stratum assignments under the same half-open convention. These series provide the observed inflow for the dynamical layer, the risk sets and event counts for probability estimation, and the comparison series for reconstruction diagnostics.

3.3 Daily Exit Probabilities and Hazard-Scale Reporting

Daily exit probabilities are estimated on the grouped day-of-stay support used by the cohort dynamics. Ordinary cells are $m = 0, \ldots, m_{\max} - 1$. If a single retained tail is used, the additional cell $m = m_{\max}$ aggregates episode-days with grouped day of stay at least $m_{\max}$; this row is a pooled closure for the retained long-stay state, not a day-specific probability for each $m \geq m_{\max}$.

Let $W_i(m)$ be the row index assigned to episode i on grouped day of stay m:

$$W_i(m) = \begin{cases} P(s_i), & \text{admission-based convention,} \\ s_i, & \text{calendar-time, admission-day indexed,} \\ P(s_i + m), & \text{calendar-time, regime-pooled.} \end{cases}$$

For stratum j, row w, and ordinary grouped day $m = 0, \ldots, m_{\max} - 1$, define

$$Y_{j,w}(m) := \sum_{i=1}^{N} \mathbf{1}\{j(i) = j,\ \ell_i > m,\ W_i(m) = w\}, \tag{10}$$

$$Z_{k,j,w}(m) := \sum_{i=1}^{N} \mathbf{1}\{j(i) = j,\ c_i = k,\ \ell_i = m + 1,\ W_i(m) = w\},\ k \in \{D, M\}. \tag{11}$$

When $Y_{j,w}(m) > 0$,

$$\widehat{q}_{k,j,w}(m) = \frac{Z_{k,j,w}(m)}{Y_{j,w}(m)},\ \widehat{q}_{j,w}(m) = \sum_{k \in \{D,M\}} \widehat{q}_{k,j,w}(m).$$

For a single retained tail, probabilities must be estimated at the same resolution as the tail state propagated by the cohort update. Let v denote the retained tail-closure index. For an unindexed marginal tail, $v = *$ and long-stay exposure is collapsed over admission regimes or admission days. For a calendar-time regime-pooled tail, v may be the contemporaneous regime, since all patients in the tail at calendar day t^n are exposed to $P(t^n)$. Define

$$V_i(u) = \begin{cases} P(s_i + u), & \text{calendar-time regime-pooled tail closure,} \\ *, & \text{unindexed marginal single-tail closure.} \end{cases}$$

Then, for each retained tail index v,

$$Y^{\text{tail}}_{j,v} := \sum_{i=1}^{N} \sum_{u=m_{\max}}^{\ell_i - 1} \mathbf{1}\{j(i) = j,\ V_i(u) = v\}, \text{ and}$$

$$Z^{\text{tail}}_{k,j,v} := \sum_{i=1}^{N} \mathbf{1}\{j(i) = j,\ c_i = k,\ \ell_i - 1 \geq m_{\max},\ V_i(\ell_i - 1) = v\},\ k \in \{D, M\}.$$

If $Y^{\text{tail}}_{j,v} > 0$, the tail-closure row is

$$\widehat{q}^{\text{tail}}_{k,j,v} = \frac{Z^{\text{tail}}_{k,j,v}}{Y^{\text{tail}}_{j,v}},\ \widehat{q}^{\text{tail}}_{j,v} = \sum_{k \in \{D,M\}} \widehat{q}^{\text{tail}}_{k,j,v}.$$

Tail-closure rows with zero exposure are completed by the same fallback rule used for no-exposure ordinary rows, and this choice is recorded with the other completion settings.

In the cohort update,

$$q^{n,\text{tail}}_{k,j} = \begin{cases} \widehat{q}^{\text{tail}}_{k,j,P(t^n)}, & \text{calendar-time regime-pooled tail closure,} \\ \widehat{q}^{\text{tail}}_{k,j,*}, & \text{unindexed marginal single-tail closure.} \end{cases}$$

Tail counts indexed by admission regime or admission day may be used as diagnostic summaries, but they are valid simulation inputs only if the tail state carries the same index. Otherwise they must be collapsed into the marginal tail closure above. If grid-respecting tail states are retained, the ordinary day-specific definitions in (10)–(11) extend to the chosen tail support.

Sparse rows are completed by deterministic forward adjacent-bin pooling. Given $Y_{\min} \geq 1$, the algorithm starts at the first uncompleted day m and selects the shortest block $B = [m, m']$ satisfying

$$\sum_{u=m}^{m'} Y_{j,w}(u) \geq Y_{\min}.$$

The pooled numerator and denominator over B define a common probability for that block. A remaining positive-exposure suffix below threshold is pooled and reported; no-exposure cells carry forward the last completed value when available, otherwise use the prespecified fallback rule; and rows with no positive exposure use the same fallback rule. The applied choices for $m_{\max}$, $Y_{\min}$, tail treatment, and same-day handling are reported in Sect. 4.

The completed probability tables are passed directly to the cohort dynamics. Hazard-scale quantities are used only for reporting or for compatible continuous-time embeddings. For this purpose, let ω denote a generic completed row index: $\omega = w$ for ordinary cells and $\omega = v$ for retained tail-closure rows. For tail rows, set

$$\widehat{q}_{k,j,\omega}(m_{\max}) := \widehat{q}^{\text{tail}}_{k,j,\omega},\ \widehat{q}_{j,\omega}(m_{\max}) := \widehat{q}^{\text{tail}}_{j,\omega}.$$

Given completed probabilities $\widehat{q}_{j,\omega}(m)$ and $\widehat{q}_{k,j,\omega}(m)$, and $\epsilon > 0$, define

$$\widetilde{q}_{j,\omega}(m) := \min\{\max(\widehat{q}_{j,\omega}(m), 0),\, 1-\epsilon\},\ \ \widetilde{\Lambda}_{j,\omega}(m) := -\log\big(1-\widetilde{q}_{j,\omega}(m)\big),$$

and

$$\widetilde{\Lambda}_{k,j,\omega}(m) := \begin{cases} \widetilde{\Lambda}_{j,\omega}(m)\dfrac{\widehat{q}_{k,j,\omega}(m)}{\widehat{q}_{j,\omega}(m)}, & \widehat{q}_{j,\omega}(m) > 0, \\ 0, & \widehat{q}_{j,\omega}(m) = 0, \end{cases} \quad k \in \{D, M\}.$$

This preserves the cause proportions while replacing any unit total probability by a finite-hazard approximation. All completion and tail-closure choices are recorded with $m_{\max}$, $Y_{\min}$, ϵ, same-day handling, and retained-tail treatment.

3.4 Hazards-To-Dynamics Simulation and Internal Consistency Checks

After completion on the chosen working grid, the daily exit-probability tables are combined with the reconstructed admission series $A_j(t)$ and propagated through the cohort update of Sect. 2.5. The model-implied calendar-time series are

$$\widehat{H}_{\mathrm{sim}}(t^n) := \sum_{j=1}^{J} H_j^n,\ \ \widehat{D}_{\mathrm{sim}}(t^n) := \sum_{j=1}^{J} \widetilde{D}_j^{n-1},\ \ \widehat{M}_{\mathrm{sim}}(t^n) := \sum_{j=1}^{J} \widetilde{M}_j^{n-1},\ \ n \geq 1, \tag{12}$$

with $\widehat{H}_{\mathrm{sim}}(t^0) = \sum_{j=1}^{J} H_j^0$. Thus, $\widehat{H}_{\mathrm{sim}}(t^n)$ is occupancy on day t^n, whereas $\widehat{D}_{\mathrm{sim}}(t^n)$ and $\widehat{M}_{\mathrm{sim}}(t^n)$ are outflows with exit date t^n.

These series are compared with empirical occupancy and outflows reconstructed from episode dates under the same half-open convention. If the working grid is represented exactly, the comparison is direct. Under a conservative reduced-tail closure, it is a reduced-state reconstruction. Under truncation, survivor mass beyond the last retained cell must be reported.

For closed retained-tail implementations, the accounting residual is

$$\mathrm{Res}(t^n) := \widehat{H}_{\mathrm{sim}}(t^n) - \widehat{H}_{\mathrm{sim}}(t^{n-1}) - A(t^n) + \widehat{D}_{\mathrm{sim}}(t^n) + \widehat{M}_{\mathrm{sim}}(t^n),\ \ n \geq 1. \tag{13}$$

By Proposition 2, this residual is zero in exact arithmetic once the initialization, timing convention, and retained-tail update have been fixed. With a reduced tail closure, a zero residual verifies conservation of the implemented reduced scheme, not exact recovery of all day-specific long-stay behaviour.

For truncated implementations, the modified balance in (9) applies. With $U(t^n) := \sum_{j=1}^{J} U_j^{n-1}$, the adjusted residual is

$$\mathrm{Res}^*(t^n) := \mathrm{Res}(t^n) + U(t^n),\ \ n \geq 1,$$

and is expected to vanish in exact arithmetic.

The residuals are implementation diagnostics, not empirical fit measures. Reconstruction quality is assessed separately by comparing (12) with the empirical series from Sect. 3.2 under the same state definition. Because both the estimated probability tables and the comparison series are derived from the retained

sample, this is an internal reconstruction check rather than external predictive validation. For summary reporting, the workflow computes occupancy mean absolute error and root mean squared error, together with mean absolute errors for discharge and death outflows, over all calendar days after the initial analysis day. It also stores probability-table support diagnostics, completion summaries, model-state support by grouped day of stay, accounting residuals, sensitivity summaries, and figures.

4 Results and Diagnostics

This section reports the applied episode extract, the support of the estimated probability tables, the reconstruction diagnostics, and the sensitivity checks. All comparisons are internal reconstruction diagnostics on the retained episode data.

4.1 Applied Administrative Hospitalization Extract

The model-ready table contained 447,752 hospitalization episodes, with admission dates from 2017-01-01 to 2023-12-31 and recorded exit dates from 2017-01-02 to 2024-03-11. It included 428,968 discharges alive and 18,784 in-hospital deaths. The table was treated as an observed closed-episode extract, and the reconstruction diagnostics reported below are internal to the corresponding analysis cohort.

If later exits for admissions close to 2023-12-31 were unavailable in the source files, late-window estimates, especially for long-stay cells, should be interpreted as conditional on the observed closed-episode extract rather than as censoring-adjusted estimates.

The reference analysis used the default preprocessing rule, excluding 5,106 same-day admission-and-exit episodes. This left 442,646 retained episodes, with 424,164 discharges alive and 18,482 in-hospital deaths. The analysis used 10 age-sex strata, a regime-pooled calendar-time representation, $m_{\max} = 60$, $Y_{\min} = 25$, and a single retained tail cell. Here $m_{\max} = 60$ means ordinary day-of-stay cells $m = 0, \ldots, 59$, plus a retained calendar-time regime-pooled tail row for $m \geq 60$. The regime cut points were 2020-03-01 and 2022-01-01, defining pre-COVID, COVID, and post-COVID periods. The empirical occupancy series reconstructed from episode dates covered 2,627 calendar days, with mean empirical occupancy 1,356.01. Outflow summaries were computed over the 2,626 post-initial dates, giving mean daily discharge and death outflows of 161.52 and 7.04, respectively. The maximum recorded length of stay in the retained sample was 305 days.

4.2 Probability-Table Support

The reference probability table had 1,830 cells, corresponding to 10 strata, 3 regimes, and 61 day-of-stay states, including the retained tail cell. Positive exposure was present in 1,791 cells, or 97.87% of the table, and no cell had total estimated exit probability equal to one. The large majority of exposure was supported by raw cell estimates. Raw cells accounted for 99.887% of total episode-day exposure, while adjacent pooled cells and suffix pooled cells below $Y_{\min}$

accounted for 0.111% and 0.001%, respectively. The remaining 39 no-exposure cells were completed by carrying forward the last available value within the corresponding row. The exposure-weighted total one-day exit probability was 0.1243, decomposed into 0.1191 for discharge and 0.00519 for in-hospital death. Exposure was concentrated in the early and intermediate days of stay: the exposure-weighted median grouped day of stay was $m = 4$, and 90%, 95%, and 99% of exposure were accumulated by $m = 17$, $m = 23$, and $m = 43$. By $m = 30$, 97.42% of total exposure had already been observed. The retained tail cell $m \geq 60$ represented 0.40% of empirical episode-day exposure.

4.3 Reconstruction and Accounting Diagnostics

The model-implied occupancy and cause-specific outflow series were compared with the empirical series reconstructed directly from admission and exit dates under the same half-open convention. For occupancy, the reference implementation achieved $\mathrm{MAE}_{\mathrm{occ}} = 53.03$, $\mathrm{RMSE}_{\mathrm{occ}} = 67.38$. Relative to the mean empirical occupancy, these errors correspond to 3.91% and 4.97%. For daily outflows, the mean absolute errors were $\mathrm{MAE}_{\mathrm{dis}} = 26.69$, $\mathrm{MAE}_{\mathrm{death}} = 2.10$. These correspond to 16.52% of the mean daily discharge outflow and 29.80% of the mean daily in-hospital death outflow. The higher relative error for deaths is consistent with the lower volume and greater discreteness of the death outflow series. Accounting residuals computed from (13) were negligible, with mean and maximum absolute values 7.63×10^{-14} and 2.45×10^{-13}, respectively, consistent with exact discrete conservation up to numerical roundoff. The zero-residual diagnostic verifies conservation of the reduced closed-tail scheme, not exact recovery of all day-specific long-stay exit probabilities beyond $m_{\max}$.

4.4 Sensitivity Checks

Sensitivity was assessed with respect to finite support, sparse-cell pooling, retained-tail treatment, and same-day handling. The grid used $m_{\max} \in \{36, 60, 90\}$, $Y_{\min} \in \{12, 25, 50\}$, single-cell and grid-respecting retained-tail treatments, and two same-day policies: exclusion and recoding as one-day stays. This produced 36 configurations. The reconstruction metrics changed little across these settings. Occupancy MAE ranged from 52.99 to 53.29, and occupancy RMSE from 67.34 to 67.68. Daily discharge MAE ranged from 26.69 to 26.76, while daily death MAE ranged from 2.097 to 2.110. In relative terms, occupancy MAE remained between 3.91% and 3.92% of mean empirical occupancy, and occupancy RMSE between 4.97% and 4.98%. The corresponding ranges were 16.38%–16.52% for discharge outflow MAE and 29.49%–29.80% for death outflow MAE.

Recoding same-day episodes increased the retained sample from 442,646 to 447,752 records, but did not materially affect the reconstruction metrics. Similarly, varying $m_{\max}$, $Y_{\min}$, or the retained-tail treatment produced only small changes in the main errors. All closed-tail configurations preserved mass balance to numerical precision.

5 Conclusions

This paper proposed a reproducible daily-grid framework for converting administrative hospital episode data recorded at date resolution into daily occupancy and cause-specific outflow series. The framework links grouped-time competing-risks estimation with conservative daily cohort dynamics under a half-open timing convention, making the accounting links among admissions, occupancy, discharges, and in-hospital deaths explicit. Its main methodological contribution is to provide a transparent hazards-to-dynamics route, operationalised through grouped daily exit probabilities, while separating quantities directly supported by retained episode data from those requiring pooling, completion, tail closure, or modelling assumptions. Admission and exit dates determine the empirical daily series and support one-day exit probabilities on grouped day-of-stay cells with nonempty risk sets. Continuous-time hazards are used only as compatible embeddings, and long-stay behaviour is represented through an explicit retained-tail construction.

In the Brazilian femur-fracture hospitalization application, the reference reconstruction produced small relative occupancy error and preserved mass balance to numerical precision. Sensitivity checks over finite support, sparse-cell pooling, same-day handling, and retained-tail treatment showed stable reconstruction metrics. Because the empirical application uses one diagnosis group from one administrative data system, these results establish internal reconstruction consistency for this retained cohort; generalisability across hospitals, diagnoses, or health systems remains a question for external validation.

In practice, the workflow is most useful for analysts, planners, and health-services researchers who need to translate closed episode records into calendar-time summaries of inpatient flow. It can be run on routine administrative episode extracts to produce daily occupancy reconstructions, cause-specific outflow profiles, and day-of-stay exit probabilities by age, sex, diagnosis group, ward, hospital, or policy period, depending on the fields available in the administrative extract. These outputs can support retrospective flow audits, comparison of operating regimes, and scenario analyses in which alternative admission streams or exit-probability tables are propagated through the same conservative cohort update. They are intended to complement, rather than replace, real-time operational bed-management systems: the framework provides daily cohort occupancy and outflow inputs that can be combined with capacity information when operational bed pressure is the quantity of interest.

Natural extensions include uncertainty quantification, smoother or hierarchical estimation of daily exit probabilities, richer state definitions, multi-state models when additional timestamps are available, and external validation against independent operational data or additional clinical cohorts.

Acknowledgments. C. Castro, A.J. Soares, and A.P. Amorim acknowledge support from the FCT project UID/00013/2025 (https://doi.org/10.54499/UID/00013/ 2025).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. DATASUS (Brazilian Ministry of Health): Morbidade Hospitalar do SUS (SIH/-SUS). https://datasus.saude.gov.br/acesso-a-informacao/morbidade-hospitalar-do-sus-sih-sus/. Accessed 03 Mar 2026
2. Worthington, D., Utley, M., Suen, D.: Infinite-server queueing models of demand in healthcare: a review of applications and ideas for further work. J. Oper. Res. Soc. **71**(8), 1145–1160 (2020). https://doi.org/10.1080/01605682.2019.1609878
3. Bekker, R., uit het Broek, M., Koole, G.: Modeling COVID-19 hospital admissions and occupancy in the Netherlands. Eur. J. Oper. Res. **304**(1), 207–218 (2023). https://doi.org/10.1016/j.ejor.2021.12.044
4. Bos, H., Baas, S., Boucherie, R.J., Hans, E.W., Leeftink, G.: Bed census prediction combining expert opinion and patient statistics. Omega **133**, 103262 (2025). https://doi.org/10.1016/j.omega.2024.103262
5. Kalbfleisch, J.D., Prentice, R.L.: The Statistical Analysis of Failure Time Data, 2nd edn. Wiley, Hoboken (2002). https://doi.org/10.1002/9781118032985.
6. Prentice, R.L., Gloeckler, L.A.: Regression analysis of grouped survival data with application to breast cancer data. Biometrics **34**(1), 57–67 (1978). https://doi.org/10.2307/2529588
7. Perthame, B.: Transport Equations in Biology. Birkhäuser, Basel (2007). https://doi.org/10.1007/978-3-7643-7842-4
8. LeVeque, R.J.: Finite Volume Methods for Hyperbolic Problems. Cambridge University Press, Cambridge (2002). https://doi.org/10.1017/CBO9780511791253.
9. McKendrick, A.G.: Applications of mathematics to medical problems. Proc. Edinb. Math. Soc. **44**, 98–130 (1925). https://doi.org/10.1017/S0013091500034428
10. von Foerster, H.: Some remarks on changing populations. In: Stohlman, F. (ed.) The Kinetics of Cellular Proliferation, pp. 382–407. Grune & Stratton, New York (1959)
11. World Health Organization: International Statistical Classification of Diseases and Related Health Problems, 10th Revision (ICD-10), Version 2019. https://icd.who.int/browse10/2019/en. Accessed 03 Mar 2026
12. Andersen, P.K., Borgan, Ø., Gill, R.D., Keiding, N.: Statistical Models Based on Counting Processes. Springer, New York (1993). https://doi.org/10.1007/978-1-4612-4348-9.
13. Cushing, J.M.: An Introduction to Structured Population Dynamics. SIAM, Philadelphia (1998). https://doi.org/10.1137/1.9781611970005.

Convolutional Autoencoder Compression for LoRa Visual Transmission on Embedded Systems

João Mendes[1,2(✉)], Inês Sena[1,2], José Lima[1,2], and Ana I. Pereira[1,2]

[1] Research Centre in Digitalization and Intelligent Robotics (CeDRI), Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
{joao.cmendes,jllima,apereira}@ipb.pt

[2] Laboratório Associado para a Sustentabilidade e Tecnologia em Regiões de Montanha (SusTEC), Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
ines.sena@ipb.pt

Abstract. This work proposes and evaluates a lightweight convolutional autoencoder (CAE) designed for image compression in Internet of Things (IoT) nodes subject to severe bandwidth constraints. The system follows an encode–transmit–decode paradigm: a compact encoder, intended for deployment on an ESP32 microcontroller, compresses a 96×96 pixel grayscale image into a low-dimensional latent vector; this vector is then transmitted over LoRa as a single packet or multiple fragments; and a decoder at the receiver reconstructs the image from the received representation. Seven latent dimensions ($d \in \{8, 16, 32, 64, 128, 256, 512\}$) are systematically evaluated and compared against a Principal Component Analysis (PCA) baseline using Peak Signal-to-Noise Ratio (PSNR, measured in decibels) and the Structural Similarity Index Measure (SSIM) on the MNIST dataset. Results show that the CAE substantially outperforms PCA in all single-packet LoRa configurations ($d \leq 64$), achieving reconstructions that surpass commonly accepted quality thresholds while reducing the transmitted payload to a single 64-byte LoRa packet, a 144:1 reduction relative to the raw image. These findings demonstrate that non-linear learned representations are substantially more efficient than linear projections under the low-dimensional regime imposed by LoRa bandwidth constraints.

Keywords: Wireless Visual Sensor Networks · LoRa · Convolutional Autoencoder · Learned Image Compression · TinyML · Internet of Things

1 Introduction

The use of Wireless Sensor Networks (WSN) in the context of the Internet of Things (IoT) has intensified the requirement for efficient transmission of visual

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 210–223, 2027.
https://doi.org/10.1007/978-3-032-30518-3_14

information from embedded nodes with severely constrained computational and energy resources [1,2]. In sectors such as precision agriculture, environmental monitoring, and remote infrastructure management, the capacity to capture and communicate image data from low-cost devices is a growing, yet unresolved challenge [3,4].

Unlike scalar sensor data, images carry far more information per sample, and that comes at a cost. Wireless Visual Sensor Networks (WVSN) must deal with data volumes and structural complexity that standard WSN architectures were simply not built for, making the design of acquisition, processing, and transmission pipelines a genuinely difficult engineering problem [5].

In this context, LoRa (Long Range) technology is among the most promising wireless communication solutions for a wide range IoT applications [6]. LoRa modulation, based on Chirp Spread Spectrum (CSS), enables communication ranges up to 10 km in rural environments while maintaining "listening" mode energy consumption on the order of microwatts, making it especially suitable for battery- or solar-powered nodes [7]. However, the actual application bandwidth is typically below 2 kbps, and the maximum payload per packet is limited to approximately 115 bytes under common configurations (SF9, 125 kHz, 868 MHz band), resulting in severe restrictions on data transmission capacity [8,9].

Under these conditions, direct transmission of uncompressed images is entirely impossible. For example, a grayscale image of size 96×96 pixels requires 9,216 bytes, equivalent to more than 80 LoRa packets and a transmission time exceeding 40 s, without accounting for the duty cycle restrictions imposed by ETSI EN 300 220-2. Therefore, image compression is required for any system that transmits visual data over LoRa.

The main problem is that classic techniques of compression such as JPEG, which are widely available even on low-end microcontrollers [10], are designed for moderate compression with minimal perceptual loss, and its performance degrades significantly at very low bitrates, the exact regime in which LoRa systems operate. JPEG compression at payloads of 64 or 128 bytes introduces severe blocking artifacts and structural detail loss, compromising the use of the reconstructed image [11].

One alternative to these methods is the convolutional autoencoder (CAE). These models are trained to learn compact, informative representations of images [12,13]. CAEs learn nonlinear transformations of the input data, preserving semantically relevant features while using significantly fewer bytes [14].In the encoder, executed on the embedded device, a compression of the image into a low-dimensional latent vector. This is transmitted via LoRa and the decoder component, executed on the receiver server, reconstructs the image from the received representation.

All the feasibility of this method raises several research questions:

(i) What is the minimum latent vector size required to achieve acceptable visual quality, defined here as Peak Signal-to-Noise Ratio (PSNR, a logarithmic measure of reconstruction fidelity expressed in decibels) $\geq$ 25 dB and Struc-

tural Similarity Index Measure (SSIM, a perceptual metric ranging from 0 to 1) ≥ 0.90?

(ii) How does the convolutional autoencoder compare with Principal Component Analysis (PCA), a classical linear compression baseline [15], in terms of PSNR and SSIM for the same payload?

(iii) What latent dimension best balances reconstruction quality and transmission efficiency given the single-packet LoRa boundary ($d \leq 112$ bytes of useful payload)?

On top of this, executing the encoder on a low-cost microcontroller, specifically the ESP32-S3, featuring an Xtensa LX7 processor at 240 MHz and support for neural inference via TensorFlow Lite Micro introduces extra constraints on model size and complexity, requiring the evaluation of integer quantization (INT8) on reconstruction quality.

In this way, this work brings some original contributions. Starting with the proposal and implementation of a complete encoder/decoder pipeline based on a lightweight CAE custom-made to LoRa constraints; a systematic evaluation of seven latent vector dimensions ($d \in \{8, 16, 32, 64, 128, 256, 512\}$), characterizing the quality–payload trade-off; a quantitative comparison with a PCA baseline; A quantitative comparison with a PCA baseline across all evaluated dimensions, using PSNR and SSIM as reconstruction quality metrics; Identification of the optimal latent dimension under the single-packet LoRa constraint.

The remainder of this paper is organized as follows: Sect. 2 presents the state of the art; Sect. 3 describes the methodology; Sect. 4 discusses the experimental results; and Sect. 5 concludes the paper and present the main future work.

2 State of the Art

The transmission of image data over LoRa networks has been identified as a relevant technical challenge since the emergence of the technology. Chen et al. [16] proposed the Multi-Packet LoRa (MPLR) protocol, designed for reliable transmission of long messages in agricultural monitoring networks. By eliminating per-packet acknowledgement mechanisms typical of conventional Additive Links On-line Hawaii Area (ALOHA)-based approaches and replacing them with selective Automatic Repeat reQuest (ARQ), the protocol reduced image transmission time by 24% to 49% under packet loss rates ranging from 0% to 10%.

Ji et al. [17] introduced a LoRa-based visual monitoring system for agricultural IoT that transmits only modified regions of an image, decomposed into grid patches, making use of the temporal redundancy of static cameras. However, this approach is not generalizable to active scenes or sporadic image acquisition scenarios.

In the paper [18] Staikopoulos et al. presented a comprehensive survey of image transmission approaches over LoRa networks. In this study, they showed that neither the direct fragmentation with ARQ, pre-transmission classical compression, and resolution subsampling satisfies quality and efficiency requirements

under strict bandwidth constraints. Pointing the deep learning-based compression as an encouraging direction. Borsos [19] further confirmed this drawback by comparing JPEG-based approaches under LoRaWAN payload constraints in agricultural scenarios, concluding that the resulting image quality is frequently insufficient for visual diagnostic applications.

Showing that nonlinear learned transformations can outperform JPEG in terms of PSNR at equivalent bitrate, the work of Ballé et al. [14] introduced the first end-to-end optimized variational auto-encoder architecture. Bank et al. [12] present some systematized architectural variants of autoencoders and their application domains, demonstrating that convolutional autoencoders offer advantages over linear methods such as PCA by encoding complex correlation patterns that linear transformations fail to represent efficiently.

In [11], Jamil et al. identified the high computational cost of neural compression, which is normally 100 to 10,000 times greater than JPEG, as a key barrier to industrial adoption, motivating the development of lightweight architectures suitable for embedded deployment. Guerrisi et al. [20] demonstrated the feasibility of deploying neural encoders on severely resource-limited platforms for onboard satellite image compression, achieving PSNR values above 33 dB.

The Tiny Machine Learning (TinyML) paradigm, which refers to the deployment of machine learning models on extremely resource-constrained devices, has emerged as an active research field with exponential growth in publications since 2019 [21]. Capogrosso et al. [22] identified model quantization as one of the most prevalent optimization techniques for reducing memory footprint and inference latency on microcontrollers. Abadade et al. [23] present some use cases in plant disease classification and growth monitoring implemented on microcontrollers. In these experiments, it was shown that lightweight CNN models with INT8 quantization can achieve accuracies above 90% with inference latencies below 100 ms on ESP32 platforms. Also in the context of sensor data compression in embedded systems, Gomes et al. [24] prove that neural autoencoders are a viable solution with inference times compatible with real-time applications on low-end devices.

PCA is widely used as a baseline model for its mathematical foundation and efficient implementation. Based on this, some studies have been carried out, such as Jia et al. [15], who concluded that PCA obtains competitive performance only when the data correlation structure is approximately linear, which is not typically the case for image data. While the leading principal components capture low-frequency structures effectively, they are less efficient in representing high-frequency features and local texture patterns relative to nonlinear representations learned by convolutional autoencoders. This limitation becomes especially noticeable at low latent dimensions ($d \leq 64$), which correspond exactly to the regime of interest for single-packet LoRa transmission.

Based on all of the works, and, despite considerable advances in LoRa-based image transmission, neural image compression, and TinyML, there is, to the best of the authors' knowledge, no work that systematically proposes and evaluates a complete neural encoder/decoder pipeline specifically optimized for LoRa chan-

nel constraints, including encoder deployment on a microcontroller and direct comparison with a PCA baseline across a wide range of latent dimensions. This work intends to address this gap.

3 Methodology

The experimental testing of the proposed pipeline was carried out using the MNIST dataset [25], which consists of 70,000 grayscale images of handwritten digits with an original resolution of 28×28 pixels, 60,000 for training and 10,000 for testing. Although MNIST is traditionally used for classification tasks, its adoption in this work is justified as a controlled validation environment: the dataset exhibits sufficient visual variability to assess the compression and reconstruction capabilities of the models, while its public availability guarantees reproducibility and facilitates direct comparison with future work.

To address the input requirements of the proposed encoder, all images were resized using bilinear interpolation, going from 28×28 to 96×96 pixels, and normalized to the range $[0, 1]$. These dimensions were chosen because of the convolutional architecture, which applies four downsampling operations with a stride of 2, reducing the spatial resolution to 6×6.

As shown in Fig. 1, the architecture is an asymmetric CAE composed of two distinct components: a lightweight encoder designed for execution on limited-resource embedded systems, and a more expressive decoder intended for execution on the receiver server. The encoder consists of four sequential convolutional blocks, each comprising a Conv2D layer with a 3×3 kernel and stride 2 (progressively downsampling the input from 96×96 to 6×6 pixels), followed by Batch Normalization and ReLU activation. The number of filters per block follows the sequence [16, 32, 48, 64], keeping the total number of encoder parameters below 50,000, consistent with ESP32-S3 memory bounds.

To achieve the last vector, a global average pooling layer is used. In this operation, the spatial dimension is reduced from $(6 \times 6 \times 64)$ to a 64-element vector. In the continuity of the process, a fully connected layer without activation allows unconstrained real-valued representations prior to quantization. In the decoding part executed on the server, all of this process is undone. For that, four transposed convolution blocks with stride 2, Batch Normalization, and ReLU activation, followed by a final transposed convolution with Sigmoid activation, are used to reconstruct the image in the range $[0, 1]$.

The asymmetry between encoder and decoder is a purposeful design choice: the encoder omits residual connections to minimize inference memory consumption, while the decoder incorporates them to improve reconstruction quality.

The model was trained using the Adam optimizer with an initial learning rate of 1×10^{-3} and a composite loss function combining Mean Squared Error (MSE) and Structural Similarity (SSIM), defined as:

$$\mathcal{L} = 0.7 \cdot \mathcal{L}_{MSE} + 0.3 \cdot \mathcal{L}_{SSIM}.$$

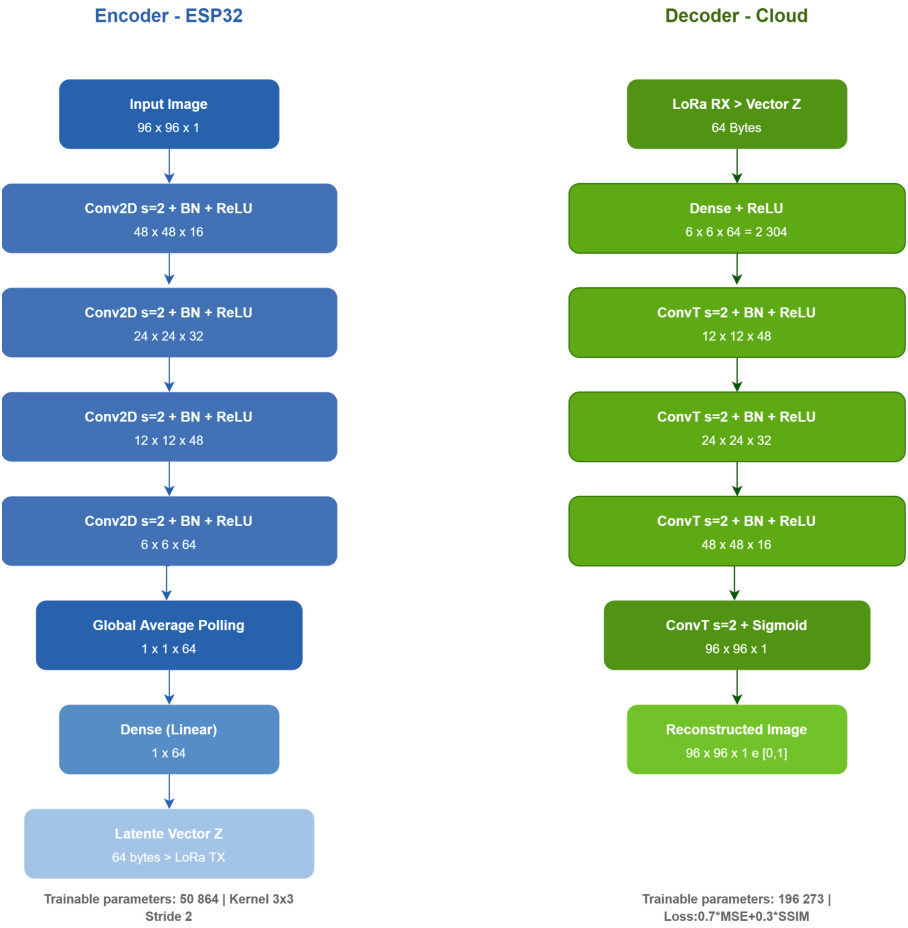

Fig. 1. Asymmetric convolutional autoencoder used in this work.

This formulation balances pixel-wise fidelity with perceptual structure preservation, and was empirically found to produce sharper reconstructions than pure MSE minimization.

All the training hiperparameters are present in the Table 1 namely adaptive early stopping, including a warm-up period of 15 epochs during which stopping criteria were disabled, followed by a patience of 20 epochs and a minimum improvement threshold of 1×10^{-5}. The maximum number of epochs was set to 200. The learning rate was reduced by a factor of 0.5 whenever validation loss did not improve for 8 consecutive epochs (ReduceLROnPlateau). Training was conducted on an NVIDIA RTX 3060 Laptop GPU (6 GB VRAM) with a batch size of 64 and a fixed random seed of 42 to maintain reproducibility.

As mentioned before the baseline adopted for comparison is PCA [15], a widely used linear compression technique. For each evaluated dimension d, PCA

Table 1. Training hyperparameters of the convolutional autoencoder.

Hyperparameter	Value
Optimizer	Adam
Initial learning rate	1×10^{-3}
Loss function	0.7· MSE +0.3· SSIM
Batch size	64
Warm-up (epochs)	15
Early stopping patience	20
Minimum delta	1×10^{-5}
ReduceLROnPlateau factor	0.5
ReduceLROnPlateau patience	8
Maximum epochs	200
Random seed	42

is fitted on the training set after pixel-wise z-score normalization, projecting each image into a vector of d principal components. Reconstruction is obtained by inverse projection followed by denormalization. This comparison allows isolating the effective gain of nonlinear learned representations over linear decompositions at identical payload sizes.

The experimental space includes seven latent vector dimensions $d \in \{8, 16, 32, 64, 128, 256, 512\}$, covering ultra-compact to high-dimensional representations. When the latent vector size exceeds the maximum LoRa payload (approximately 115 bytes under SF9, 125 kHz, 868 MHz), it must be split into multiple fragments, each transmitted in a separate packet. The mapping between latent dimension, number of fragments, and resulting air-time is summarised in Table 2.

Table 2. LoRa transmission configurations for each latent dimension (SF9, 125 kHz, 868 MHz).

d	Payload (bytes)	Fragments	Air-time (s)	Regime
8	8	1	0.52	Single-packet
16	16	1	0.52	Single-packet
32	32	1	0.52	Single-packet
64	64	1	0.52	Single-packet
128	128	2	1.05	Multi-fragment
256	256	3	1.57	Multi-fragment
512	512	5	2.61	Multi-fragment

Prior to transmission, the FP32 latent vector is quantized to UINT8 using min-max normalization based on training statistics, reducing the payload from $4d$ bytes to d bytes, with an additional overhead of 4 bytes for dequantization parameters.

Reconstruction quality is evaluated using PSNR and the Structural Similarity Index Measure (SSIM) [26]. Both metrics are computed between the original input image and the image reconstructed by the decoder, over the full test set (10 000 images), using batches of 256, and reported as mean values. PSNR is computed as:

$$\mathrm{PSNR} = 10 \cdot \log_{10} \left(\frac{\mathrm{MAX}^2}{\mathrm{MSE}} \right),$$

where MAX is the maximum possible pixel value (1.0 after normalisation) and MSE is the mean squared error between original and reconstructed images. Values above 25 dB are typically considered acceptable for lossy compression [11]. SSIM evaluates perceptual similarity in terms of luminance, contrast, and structure, ranging from 0 (no similarity) to 1 (identical). For PCA, the cumulative explained variance is additionally reported as an interpretable indicator of retained information.

Both metrics are computed over the full test set (10,000 images), using batches of 256, and reported as mean values. Additionally, for PCA, the cumulative explained variance is reported as an interpretable indicator of retained information.

The LoRa transmission layer was simulated in software, modelling payload constraints, fragmentation protocol (ALFP), and air-time based on physical channel parameters: Spreading Factor (SF), Bandwidth (BW), and Coding Rate (CR). Under the adopted configuration (SF9, BW = 125 kHz, 868 MHz band), the maximum useful payload per packet is approximately 115 bytes. No physical LoRa hardware was used at this stage, as simulation is sufficient to validate the compression pipeline and quantify quality–transmission trade-offs.

Upcoming work will include integration with real hardware (SX1276 and ESP32-S3 with TFLite Micro), verification on agricultural datasets, and measurement of response time and energy consumption in field conditions. All code was implemented in Python 3.10 using TensorFlow 2.12, scikit-learn, and pandas, and will be publicly released upon publication to ensure full reproducibility.

4 Results and Discussion

Table 3 presents the complete experimental results, including PSNR and the SSIM metrics for both the CAE and the PCA baseline, as well as convergence epochs and explained variance for all evaluated latent dimensions. The analysis show three key patterns, the first one is the consistent superiority of the CAE over PCA in the single-packet LoRa regime. The second is the existence of a sweet spot at $d = 64$ that maximizes reconstruction quality per transmitted byte. And the last one is the pronounced diminishing returns beyond $d = 64$,

making multi-fragment configurations of marginal relevance for the proposed system.

Table 3. Complete experimental results. Single-packet LoRa configurations achieving acceptable quality (PSNR $\geq$ 25 dB) are highlighted.

d	CAE				PCA			
	PSNR (dB)	SSIM	Epochs	Frags.	PSNR (dB)	SSIM	Var. Expl. (%)	Frags.
8	20.62	0.849	64	1	15.39	0.505	27.51	1
16	23.28	0.900	70	1	16.71	0.556	38.50	1
32	25.95	0.937	139	1	18.15	0.616	52.13	1
64	**27.74**	**0.955**	137	1	20.12	0.679	68.08	1
128	27.41	0.953	122	2	22.95	0.745	84.21	2
256	27.29	0.950	78	3	27.81	0.832	95.76	3
512	28.35	0.960	137	5	36.16	0.936	99.60	5

Figure 2 provides a visual comparison of the reconstructed images produced by the CAE and PCA at six latent dimensions, from $d = 8$ to $d = 256$, for a representative MNIST sample. The progressive improvement in CAE reconstruction quality is clearly visible, with recognizable digit structure already recovered at $d = 32$ and near-original fidelity achieved at $d = 64$. In contrast, PCA reconstructions remain blurry and structurally degraded across all single-packet configurations, confirming the quantitative results in Table 3.



Fig. 2. Visual comparison of CAE and PCA reconstructions across latent dimensions $d \in \{8, 16, 32, 64, 128, 256\}$ for a single test sample. Green borders indicate single-packet LoRa configurations ($d \leq 64$); red borders indicate multi-fragment configurations ($d > 64$).

The CAE consistently outperforms PCA across all dimensions up to $d = 256$, with the best values observed in the single-packet regime. Looking at the dimension of $d = 32$, its possible to check that the CAE achieves 25.95 dB, compared to 18.15 dB for PCA, representing a 7.8 dB gain and an approximate sixfold reduction in mean squared error. At $d = 64$, the advantage remains significant at 7.6 dB (27.74 dB versus 20.12 dB). This consistent superiority reveals the limitations of linear methods when applied to image data, where pixel correlations are fundamentally nonlinear [12].

Also, the SSIM metric has the same behavior, this is confirmed by the values $d = 64$, where CAE achieves an SSIM of 0.955 compared to 0.679 for PCA, indicating significantly better preservation of perceptual structure at identical compression rates. Notably, PCA at $d = 64$ explains only 68.1% of the total variance, which accounts for its lower reconstruction quality.

With a PSNR of 27.74 dB and an SSIM of 0.955, the dimension of $d = 64$ is the best configuration. This surpasses commonly accepted quality thresholds (PSNR $\geq$ 25 dB, SSIM $\geq$ 0.90) while requiring only 64 bytes (a single LoRa packet). Increasing to $d = 128$ (two fragments) only a substantial improvement of 0.33 dB (+1.2%) while doubling transmission air-time from 0.52 s to 1.05 s and introducing fragmentation overhead. In $d = 256$ the results have a slight degradation of 0.45 dB compared to $d = 64$, suggesting an approach to a suboptimal local minimum, as reflected by the lower number of training epochs (78 vs 137). In $d = 512$, the CAE outperforms $d = 64$ more clearly (+0.61 dB), but the number of LoRa packages increases to five and the transmission time to 2.61 s.

In the Fig. 3, it's possible to see six representative reconstructions at $d = 64$ for both methods. The CAE consistently preserves stroke geometry and contrast in all digit classes, with PSNR values ranging from 25.1 dB to 35.3 dB. PCA reconstructions at the same payload exhibit visible blurring and loss of fine structure, with PSNR values ranging from 17.5 dB to 23.1 dB, underscoring once again the advantage of nonlinear learned representations.

When looking at the last dimension of approximately $d = 512$, there is a reversal of trends; here, the PCA surpasses the CAE, reaching approximately 36.16 dB compared to 28.15 dB. This reversal is essentially expected, as it is a nearly flawless regime for the PCA, explaining about 99.6% of its variance. In the case study, this transmission is practically impossible due to the high fragmentation that would be necessary, as well as the increase in message airtime.

Table 4 summarises the differential performance gain of CAE over PCA across all evaluated dimensions, confirming the consistent advantage of nonlinear representations in the single-packet regime.

Looking at the number of epochs until early stopping its possible to see the complexity of the compression task. Latent dimensions $d = 32$ and $d = 64$ require the largest number of epochs (139 and 137, respectively), indicating the difficulty of learning compact yet informative representations. At Lower dimensions, the convergence tends to be faster due to reduced reconstruction requirements. While at the dimension of $d = 256$ unexpectedly fast convergence, apparently indicating convergence toward a suboptimal local minimum.



Fig. 3. Qualitative reconstruction results at $d = 64$ for six test samples, comparing the proposed CAE with the PCA baseline.

Table 4. Differential performance gain of CAE over PCA.

d	ΔPSNR (dB)	ΔSSIM	Interpretation
8	+5.23	+0.344	CAE superior, low quality
16	+6.55	+0.344	CAE superior, moderate quality
32	+7.81	+0.321	CAE superior, acceptable quality
64	+7.62	+0.276	Optimal trade-off
128	+4.46	+0.207	Reduced CAE advantage
256	−0.52	+0.118	PCA competitive
512	−7.81	+0.024	PCA dominates

This study shows some limitations, starting with the validation. This was carried out exclusively on MNIST, whose grayscale handwritten digit images differ significantly from real agricultural images in terms of texture, illumination variability, and acquisition artifacts. Second, LoRa transmission was simulated rather than experimentally validated with real hardware, and therefore channel effects such as packet loss and noise were not considered. Third, the encoder was not quantized to INT8 nor deployed on an actual ESP32-S3 platform, and thus inference delay and energy consumption were not measured.

5 Conclusion and Future Work

This work proposed and evaluated a CAE-based image compression and reconstruction pipeline specifically designed for transmission over LoRa channels

within resource-constrained embedded systems. The system follows an encode–transmit–decode paradigm: a lightweight encoder compresses the captured image at the sensor node into a low-dimensional latent vector, which is transmitted via LoRa in a single packet or multiple fragments depending on its size, and a decoder at the receiver reconstructs the image from the received representation.

The experimental evaluation, conducted on the MNIST dataset resized to 96×96 pixels, covered seven dimensions ($d \in \{8, 16, 32, 64, 128, 256, 512\}$) and included a comparison with one of the most traditional classical linear compression (PCA), using PSNR along with the SSIM as quality metrics.

The results show that the convolutional autoencoder regularly outperforms PCA across all single-packet transmission configurations ($d \leq 64$), achieving PSNR gains ranging from +5.2 dB ($d = 8$) to +7.8 dB ($d = 32$). In particular, $d = 64$ appears as the optimal operating point of the system. With only 64 bytes of payload transmittable within a single LoRa packet without fragmentation, the neural encoder achieves 27.7 dB PSNR as well as 0.955 SSIM, surpassing commonly accepted quality thresholds (25 dB and 0.90, respectively) [11].

To contextualise these results in terms of practical transmission feasibility, it is instructive to contrast the proposed approach with the uncompressed baseline. A raw 96×96 grayscale image requires 9 216 bytes, which—under the adopted LoRa configuration (SF9, BW = 125 kHz, CR = 4/5, 868 MHz)—necessitates approximately 81 packets and yields an estimated total air-time of approximately 42 s (computed as 81×0.52 s per packet), not accounting for inter-packet gaps or the duty-cycle restrictions imposed by ETSI EN 300 220-2, which would extend the effective wall-clock transmission time considerably. For PCA at $d = 64$, the payload is identical (64 bytes, single packet, 0.52 s air-time), but reconstruction quality is substantially lower (20.12 dB PSNR, SSIM 0.679) than the CAE (27.74 dB PSNR, SSIM 0.955), confirming that the benefit of the CAE lies not in further payload reduction but in superior reconstruction quality at the same transmission cost. By contrast, the proposed CAE at $d = 64$ encodes the image into a single 64-byte packet transmitted in 0.52 s, representing a 144:1 reduction in payload size relative to the raw image. This quantitative comparison demonstrates that, under LoRa constraints, compression is not simply an optimisation strategy but a central enabler of image transmission, and that the proposed pipeline achieves this with reconstruction quality that surpasses accepted perceptual thresholds (27.7 dB PSNR, SSIM 0.955).

Beyond $d = 64$, diminishing returns become evident: increasing to $d = 128$ (two fragments, doubled air-time) yields only a marginal improvement of 0.3 dB, while $d = 256$ results in a slight degradation compared to $d = 64$. These data confirm that non-linear representations learned by convolutional autoencoders are substantially more efficient than linear PCA projections in the low-dimensional regime imposed by LoRa constraints.

The limitations of this study indicate clear directions for future work. First, validation should be extended to real agricultural datasets—such as PlantVillage [27] and in-field acquired images, to evaluate performance under realistic conditions involving color information, texture variability, and illumination changes.

Second, the encoder should be quantized to INT8 using TensorFlow Lite Micro and deployed on actual ESP32-S3 hardware to measure inference latency, energy consumption per frame, and memory footprint in actual scenarios.

Third, integration with physical LoRa hardware (e.g., SX1276 modules) will enable evaluation under real transmission conditions, including packet loss, RSSI variability, and compliance with duty-cycle constraints defined by ETSI EN 300 220-2. Finally, extending the pipeline to RGB images and higher resolutions constitutes a critical step toward real-world deployment in precision agriculture applications, where color information plays a key role in disease diagnosis and crop maturity assessment.

Future work will explore these directions, including a dedicated study on olive maturity classification in real field conditions using the proposed framework.

Acknowledgments. This work was supported by FCT – Fundação para a Ciência e Tecnologia, I.P., through the projects CeDRI (UID/05757/2025, DOI: 10.54499/UID/05757/2025) and UID/PRR/05757/2025 (DOI: 10.54499/UID/PRR/05757/2025), as well as by the Associate Laboratory SusTEC (LA/P/0007/2020, DOI: 10.54499/LA/P/S0007/2020). Additional support was provided by the NORTE2030-FEDER-01182700 project Olive Grove to Fork 4.0.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Use of Artificial Intelligence Tools. Artificial intelligence–based tools were used solely to improve the clarity and grammar of the English writing. Specifically, Grammarly and ChatGPT were employed for language refinement.

References

1. El Khediri, S.: Wireless sensor networks: a survey, categorization, main issues, and future orientations for clustering protocols. Computing **104**(8), 1775–1837 (2022)
2. Fu, Z., Al-Shareeda, M., Ali, M., Manickam, S., Karuppayah, S.: Wireless sensor networks in the internet of things: review, techniques, challenges, and future directions. Indon. J. Electr. Eng. Comput. Sci. **31**, 1190–1200 (2023)
3. Mowla, M.N., Mowla, N., Shahen Shah, A.F.M., Rabie, K.M., Shongwe, T.: Internet of things and wireless sensor networks for smart agriculture applications: a survey. IEEE Access **11**, 145813–145852 (2023)
4. Aldhaheri, L., et al.: Lora communication for agriculture 4.0: opportunities, challenges, and future directions. IEEE Internet Things J. **12**(2), 1380–1407 (2025)
5. De, S., Jalajamony, H.M., Adhinarayanan, S., Joshi, S., Upadhyay, H., Fernandez, R.: Multimedia transmission over lora networks for IoT applications: a survey of strategies, deployments, and open challenges. Sensors **25**(23) (2025)
6. Lavanya, P., Subba Reddy, I.V.: A comprehensive review on lora implementation in IoT application domains. In: AIP Conference Proceedings, vol. 2516, p. 140001 (2022)
7. Chen, C., Luo, J., Xu, Z., Xiong, R., Shen, D., Yin, Z.: Enabling large-scale low-power lora data transmission via multiple mobile lora gateways. Comput. Netw. **237**, 110083 (2023)

8. Semtech Corporation. Sx1276/77/78/79 datasheet: 137 MHz to 1020 MHz low power long range transceiver (2015). Revision 7
9. Yousuf, A., Rochester, E., Ghaderi, M.: A low-cost LoRaWAN testbed for IoT: implementation and measurements, pp. 361–366 (2018)
10. Wallace, G.K.: The jpeg still picture compression standard. Commun. ACM **34**(4), 30–44 (1991)
11. Jamil, S., Piran, M.J., Rahman, M., Kwon, O.-J.: Learning-driven lossy image compression: a comprehensive survey. Eng. Appl. Artif. Intell. **123**, 106361 (2023)
12. Berahmand, K., Daneshfar, F., Salehi, E.S., Li, Y., Xu, Y.: Autoencoders and their applications in machine learning: a survey. Artif. Intell. Rev. **57** (2024)
13. Tschannen, M., Agustsson, E., Lucic, M.: Deep generative models for distribution-preserving lossy compression (2018)
14. Ballé, J., Laparra, V., Simoncelli, E.P.: End-to-end optimized image compression (2017)
15. Jia, W., Sun, M., Lian, J., Hou, S.: Feature dimensionality reduction: a review. Complex Intell. Syst. **8** (2022)
16. Chen, T., Eager, D., Makaroff, D.: Efficient image transmission using lora technology in agricultural monitoring IoT systems. In: IEEE iThings/GreenCom/CPSCom/SmartData, pp. 937–944 (2019)
17. Ji, M., Yoon, J., Choo, J., Jang, M., Smith, A.H.: Lora-based visual monitoring scheme for agriculture IoT. In: IEEE Sensors Applications Symposium (SAS), pp. 1–6 (2019)
18. Staikopoulos, A., Kanakaris, V., Papakostas, G.A.: Image transmission via lora networks—a survey. In: IEEE International Conference on Image, Vision and Computing (ICIVC) (2020)
19. Borsos, D.: Image transmission with LoRaWAN in agriculture. In: Kovács, T.A., Nyikes, Z., Berek, T., Daruka, N., Tóth, L. (eds.) HCC 2022. Advanced Sciences and Technologies for Security Applications, pp. 235–246. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-47990-8_21
20. Guerrisi, G., Del Frate, F., Schiavon, G.: Convolutional autoencoder algorithm for on-board image compression. In: IEEE IGARSS, pp. 151–154 (2022)
21. Tsoukas, V., Gkogkidis, A., Boumpa, E., Kakarountas, A.: A review on the emerging technology of tinyML. ACM Comput. Surv. **56**(10), 259 (2024)
22. Capogrosso, L., Cunico, F., Cheng, D.S., Fummi, F., Cristani, M.: A machine learning-oriented survey on tiny machine learning. IEEE Access **12**, 23406–23426 (2024)
23. Abadade, Y., Temouden, A., Bamoumen, H., Benamar, N., Chtouki, Y., Hafid, A.S.: A comprehensive survey on tinyML. IEEE Access **11**, 96892–96922 (2023)
24. Gomes, A., et al.: Autoencoders for embedded sensor data compression: a case study on vehicular IoT systems. In: IEEE ETFA (2025)
25. LeCun, Y., Cortes, C., Burges, C.J.: The MNIST database of handwritten digits (1998). http://yann.lecun.com/exdb/mnist/
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. **13**(4), 600–612 (2004)
27. Hughes, D.P., Salathé, M.: An open access repository of images for computer-based plant disease diagnostics. arXiv preprint arXiv:1511.08060 (2015)

Classification of Postal Distribution Centres

Sandra M. Aleixo[1,2(✉)], Célia Fernandes[1], Paulo Ramos[1], Margarida Ramires[1], and Rute Milho[3]

[1] Department of Mathematics, ISEL – Instituto Superior de Engenharia de Lisboa, IPL – Instituto Politécnico de Lisboa, Rua Conselheiro Emídio Navarro, 1, 1959-007 Lisboa, Portugal
{sandra.aleixo,celia.fernandes,paulo.ramos}@isel.pt, A50051@alunos.isel.pt

[2] CEAUL – Centro de Estatística e Aplicações da Universidade de Lisboa, Faculdade de Ciências da Universidade de Lisboa, Bloco C6, Piso 4, 1749-016 Lisboa, Portugal

[3] CTT-Correios de Portugal, S.A., Strategy and Operations Development, Modelling, Networks and Optimisation Directorate, Av. Marechal Gomes da Costa, 13, 1849-001 Lisboa, Portugal
rute.m.milho@ctt.pt

Abstract. This work was conducted in the context of industrial mathematics, in partnership with CTT-Correios de Portugal, S.A., highlighting the strong link between academia and the business world. It addresses the classification of Postal Distribution Centres using multivariate statistical methods combined in a less common way. First, a correlation analysis was conducted on the variables selected for the study, revealing several statistically significant associations. To mitigate collinearity, Principal Component Analysis was applied to derive uncorrelated components, which were then used as inputs to Cluster Analysis. Both hierarchical and partitional clustering methods were applied, and solution quality was assessed using graphical inspection, silhouette coefficients, and within- and between-cluster sums of squares. Ward's method, k-means, and k-medoids produced consistent and similar groupings; in this paper, the k-means results are considered. Linear Discriminant Analysis and Multinomial Logistic Regression were used to assign new Postal Distribution Centres to the previously identified groups. The second classifier was used to corroborate the classification results obtained by LDA, since some LDA assumptions are violated and one of the clusters is small. The classification results from both methods showed a high degree of agreement. To avoid inflated accuracy caused by using the same data throughout all stages of the pipeline (standardisation, PCA, CA, LDA/MLR), the leave-one-out cross-validation method was applied across the entire pipeline, rather than only at the classification stage. As a result, the accuracy value obtained is a more realistic estimate of the model's performance. All analyses were conducted in R.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 224–241, 2027.
https://doi.org/10.1007/978-3-032-30518-3_15

Keywords: Industrial mathematics · Postal distribution centres · Principal component analysis · Cluster analysis · Linear discriminant analysis · Multinomial logistic regression · Cross-validation

1 Introduction

This work builds on a study conducted at CTT-Correios de Portugal, S.A., the Portuguese postal operator founded in 1520 and one of the oldest postal institutions worldwide. CTT operates multiple Postal Distribution Centres (PDCs), where postal items are sorted and prepared for last-mile delivery. From each PDC, mail carriers depart on delivery routes ("Giros"), which may be fixed (predefined routes) or flexible (routes that vary according to daily traffic conditions). Traffic refers to the volume of items handled and can be tracked (e.g., express or registered mail) or untracked (e.g., non-priority mail). Each route comprises an active path, corresponding to distances within areas containing distribution points, and a dead path, corresponding to the remaining travel distance (e.g., from the PDC to the start of the route and from the end of the route back to the PDC). Distribution points (DPs) denote the delivery locations. These operational concepts motivate the variables considered in this study [7].

The objectives are twofold: (i) to classify the existing PDCs into homogeneous groups and (ii) to use models that enable the assignment of a new PDC to one of the identified groups. Several multivariate statistical techniques were applied using R (version 4.5.2) [23].

First, the variables provided by the company were examined through a correlation analysis, and several statistically significant associations were found. To reduce dimensionality and mitigate collinearity, Principal Component Analysis (PCA) was applied. After Varimax rotation, a set of uncorrelated principal components (PCs) was retained, explaining a large proportion of the original variance. These components were then used as inputs to Cluster Analysis (CA) to form groups of PDCs with similar characteristics, in line with the company's operational perspective. Linear Discriminant Analysis (LDA) and Multinomial Logistic Regression (MLR) were used as classifiers for assigning new PDCs to the previously defined clusters. Finally, a fully consistent leave-one-out cross-validation procedure was applied to the entire modelling pipeline: standardisation, PCA, CA, and LDA/MLR. Related work in industrial mathematics and other real-world data applications has addressed similar problems; see, for example, [1,5,6,9].

The remainder of the paper is organised as follows. Section 2 summarises the main methodological concepts. Section 3 presents the empirical analysis and results. Section 4 concludes with the main findings and practical implications.

2 Methodology

In this study, PDCs were classified using a multivariate framework: PCA, CA, and LDA/MLR. PCA was applied to obtain a set of uncorrelated components

that served as inputs to CA, which identifies groups of PDCs with similar operational profiles. Finally, LDA and MLR were fitted to the resulting cluster labels to evaluate cluster separability and to assign new PDCs to the predefined groups, with the aim of achieving high expected classification accuracy.

2.1 Principal Component Analysis

PCA was used to reduce the dimensionality of a dataset by transforming an initial set of correlated variables into a smaller set of uncorrelated variables (PCs) that retain most of the information in the original data. Each PC is a linear combination of the original variables. Components are ordered by decreasing explained variance: the first accounts for the maximum possible variance, the second accounts for the maximum remaining variance subject to being orthogonal to the first, and so on. In this study, PCA was used in three steps: (i) assessing data adequacy, (ii) selecting the number of components to retain, and (iii) deriving variables for subsequent CA.

Assessment of Data Adequacy. Before applying PCA, data suitability was assessed using the Kaiser-Meyer-Olkin (KMO) statistic and Bartlett's test of sphericity. The KMO statistic measures the proportion of variance that may be attributed to common factors. Values range from 0 to 1, with higher values indicating that PCA is appropriate. As a rule of thumb, values between 0.7 and 0.8 are considered acceptable, values above 0.8 are considered good, and values below 0.5 suggest that PCA is not suitable for the data. Bartlett's test of sphericity tests the null hypothesis that the correlation matrix is the identity matrix. Rejecting the null hypothesis indicates that at least some variables are significantly correlated, which supports the use of PCA.

Determining the Number of Principal Components. If the adequacy criteria are met, PCA can be fitted and the number of components to retain can be selected using standard criteria. The following criteria were considered: (i) the cumulative percentage of variance explained (retaining enough components to explain at least 70% of the total variance) and (ii) the Kaiser criterion (retaining components with eigenvalues greater than or equal to one).

Obtaining Variables for CA Using PCA. To improve interpretability, Varimax rotation was applied to the retained components. The Varimax orthogonal rotation method aims to ensure that, for each principal component, only a few loadings are substantial, with all others close to zero; that is, it maximises the variation among the loadings within each principal component. Differences in the variables' contributions are amplified by increasing the loadings of those that contribute most to the component and decreasing the loadings of those that contribute least. Loadings with an absolute value of 0.5 or greater are considered relevant. The resulting scores of the rotated, uncorrelated components

were then used as inputs to the subsequent CA. For further details on PCA, see, for example, [11,13,14,19].

2.2 Cluster Analysis

CA partitions a set of objects into groups (clusters) based on measurements of p variables, with the aim of producing clusters that are internally cohesive and well separated from each other. In practice, CA involves two key steps: selecting an appropriate clustering solution (including the number of clusters) and assessing its quality.

Clustering Methods. Objects are grouped so that members of the same cluster are more similar to each other than to members of different clusters. Clustering can be performed using hierarchical or partitional (non-hierarchical) methods. Hierarchical methods rely on a pairwise distance matrix, while partitional methods operate directly on the data matrix. In this study, multicollinearity was mitigated by using uncorrelated component scores as inputs. Euclidean distance, computed on standardised variables, was used as the dissimilarity measure. Among hierarchical approaches, single linkage, complete linkage, average linkage, and Ward's method were considered; only the methods that yielded the most coherent solutions were reported. For partitional clustering, the k-means and k-medoids algorithms were applied. The **k-means** method seeks to partition the observations into k clusters by minimising the within-cluster sum of squares. The algorithm alternates between (i) assigning each observation to the cluster with the nearest centroid and (ii) updating centroids as the mean of the assigned observations, until convergence. To visualise the clustering results, a two-dimensional graphical representation based on the first two principal components was used.

Cluster Validation. Cluster quality was assessed using the **silhouette coefficient** for both hierarchical and partitional solutions. For observation i, the silhouette value S_i compares its average dissimilarity to objects within its assigned cluster with its average dissimilarity to objects in the nearest alternative cluster. Values of S_i close to 1 indicate well-clustered observations; values near 0 suggest borderline assignments between two clusters; and negative values indicate potential misclassification. For further details on CA, see, for example, [8,12,17,18].

2.3 Linear Discriminant Analysis

LDA is used to assess whether the variables considered can reliably separate the groups previously obtained by CA and to assign new observations to those groups. LDA constructs up to $m = \min\{k-1, p\}$ linear discriminant functions for k groups and p discriminant variables, aiming to maximise between-group separation while controlling within-group variability.

Assumption Assessment. Classical LDA relies on approximately multivariate normal distributions within each group and equal within-group covariance matrices across groups. These conditions were assessed using standard diagnostic procedures, including the Royston test for multivariate normality [20–22,24,26] and Box's M test for homogeneity of covariance matrices [3,4,25]. In practice, LDA is often reasonably robust to moderate departures from these assumptions when group sizes are not too small and the separation between groups is adequate. In addition, the discriminant variables are not linear combinations of one another; LDA assumes that the predictor variables are not highly multicollinear. This concern is mitigated when using uncorrelated component scores as inputs.

Model Estimation. LDA constructs m discriminant functions as linear combinations of the independent variables. The coefficients are chosen to maximise between-group separation relative to within-group dispersion, yielding uncorrelated discriminant scores. Discriminant-score plots provide a convenient way to visualise group separation and identify potential overlap or outliers.

Classification Rule. Classifying observations–either those used to fit the discriminant functions or new cases–into predefined groups is one of the main objectives of discriminant analysis. Classification can be performed using classification functions. If group sizes are unequal or prior probabilities differ, these functions can be adjusted by adding a term that reflects group priors. For each observation, a classification score is computed for each group by evaluating the corresponding classification function, and the observation is assigned to the group with the highest score. Classification functions can be used both to assess the classification performance of the discriminant model and to assign new observations to the predefined groups.

Evaluation of Classification Performance. To assess classification performance, the rates of correct and incorrect classification were estimated. One common approach is the leave-one-out (jackknife) method, where one observation is removed, the discriminant functions are re-estimated using the remaining data, and the omitted observation is then classified. Repeating this procedure for all observations provides a nearly unbiased estimate of the misclassification rate and is often preferred for small to moderate sample sizes. See, for example [16] for further details.

2.4 Multinomial Logistic Regression

MLR is a supervised classification method used when the response variable has more than two nominal categories. In this study, it is used to assess whether the variables considered can reproduce the predefined group structure and to assign new observations to those groups. Unlike LDA, MLR does not require the explanatory variables to follow a multivariate normal distribution within each

group, nor does it require equality of covariance matrices across groups. For this reason, it provides a less restrictive framework for multiclass classification problems.

Assumption Assessment. MLR relies on fewer distributional assumptions than LDA. In particular, it does not assume multivariate normality of the explanatory variables within groups or homogeneity of covariance matrices across groups. However, some conditions should still be considered. The response variable must consist of mutually exclusive nominal categories, the observations should be independent, and severe multicollinearity among the explanatory variables should be avoided. In addition, each category should contain a sufficient number of observations to allow stable estimation of the model parameters. When component scores are used as explanatory variables, the risk of multicollinearity is reduced because these scores are uncorrelated by construction.

Model Estimation. MLR models the relationship between a set of explanatory variables and a categorical response variable with more than two possible outcomes. One category is taken as the reference category, and the model estimates the relative effect of the explanatory variables on the probability of belonging to each of the remaining categories. The estimated model therefore provides, for each observation, a set of predicted probabilities, one for each possible group. These probabilities summarise the evidence in favour of assigning the observation to each category.

Assessment of the Fitted Model. The adequacy of the fitted MLR model can be assessed using standard model-based criteria and diagnostic measures. These may include likelihood-based tests, information criteria, pseudo-coefficients of determination, and the statistical significance of the estimated coefficients. Such measures provide information on the overall contribution of the explanatory variables to group separation. However, when the response categories correspond to groups obtained from a previous unsupervised procedure, these results should be interpreted as evidence of reproducibility of the predefined group structure rather than as independent validation of that structure.

Classification Rule. Classification is based on the predicted probabilities obtained from the fitted MLR model. For each observation, the model estimates the probability of membership in each predefined group. The observation is then assigned to the group with the highest predicted probability. This rule can be applied both to the observations used to fit the model and to new observations, provided that the same explanatory variables are available and processed in a consistent way.

Evaluation of Classification Performance. To assess classification performance, the predicted group labels are compared with the original group memberships. This comparison is usually summarised through a confusion matrix, from which the rates of correct and incorrect classification can be obtained. A robust internal assessment can be obtained using a leave-one-out (jackknife) procedure. In this approach, one observation is removed from the dataset, the MLR model is fitted using the remaining observations, and the omitted observation is then classified according to the highest predicted probability. This procedure is repeated until each observation has been omitted once. The global confusion matrix is then constructed by aggregating the out-of-sample predictions obtained across all iterations. This provides a less optimistic estimate of the misclassification rate and is particularly useful when an independent validation sample is not available. For further details on MLR, see, for example, [10].

3 Practical Application to CTT Data

This work is part of a study conducted at CTT-Correios de Portugal, S.A., which operates multiple PDCs. The main goals are to (i) group existing PDCs into homogeneous, well-separated clusters and (ii) use models that enable the assignment of a new PDC to one of the identified groups. These goals were addressed using multivariate statistical techniques.

3.1 Variables

The dataset provided by CTT includes the following variables:

- PDC identifier (`pdc`);
- Average number of postmen per kilometre travelled (`postmen`);
- Population density, inhabitants per km^2 (`pop_density`);
- Fast transport, by car or moped, per kilometre travelled (`fast_transp`);
- Slow transport, on foot or by bicycle, per kilometre travelled (`slow_transp`);
- Distribution points per kilometre travelled (`dist_points`);
- Traffic volume, number of objects, per kilometre travelled (`traffic_vol`);
- Total distance travelled, km (`route`).

The average number of postmen per kilometre travelled is defined as the average number of postmen working on a given route divided by the length of that route. This variable captures distribution efficiency in relation to the human effort required to cover the route: a high number of postmen per kilometre may indicate lower efficiency, whereas a low number may indicate excessive workload.

Population density per km^2 is the number of inhabitants per unit area (km^2) in the region considered. It serves as a proxy for demand for postal services, since more densely populated areas typically require more deliveries. It also supports route planning and resource allocation by helping to adjust the number of postmen and vehicles to the underlying population.

Fast transport (car or moped) per kilometre travelled reflects the distance covered using faster modes of transport. This variable assesses the capacity for rapid coverage and delivery efficiency, and supports route planning by indicating when motorised transport is prioritised, for example, on longer routes or in less walkable areas.

Slow transport (on foot or by bicycle) per kilometre travelled reflects the distance covered on foot or by bicycle. It helps identify areas where delivery speed is constrained–often dense urban centres or narrow streets–and supports optimisation in locations where vehicles have limited access.

Distribution points per kilometre travelled represents the number of delivery points (residences, businesses, post office boxes) relative to the distance travelled. It measures delivery concentration: high values indicate dense delivery patterns, whereas low values suggest dispersed areas, which can reduce efficiency.

Traffic volume (number of objects) per kilometre travelled is the average number of letters or parcels delivered per kilometre travelled. This variable quantifies workload relative to distance and is useful for scaling resources, predicting delivery times, and planning routes. High values may indicate more productive routes, whereas low values may reflect more dispersed delivery patterns.

The distance travelled (km) represents the total distance travelled by the routes associated with each PDC. It enables classification of distribution centres according to the size and complexity of their routes and supports strategic planning, such as grouping similar PDCs for logistics optimisation.

3.2 Correlation Analysis

Since CA is sensitive to strong correlations among inputs, which can influence the results, pairwise associations among the original variables were first examined. Figure 1 shows the correlogram computed using Spearman's rank correlation, which is appropriate when normality cannot be assumed for the variables under study.

The correlogram presented in Fig. 1 indicates that most variables are strongly correlated. Therefore, applying CA directly to the original variables could bias the clustering solution. To mitigate this issue, PCA was used to derive uncorrelated component scores, which were then used as inputs.

3.3 Principal Component Analysis

The suitability of the data for PCA was assessed using the KMO statistic and Bartlett's test of sphericity. The KMO value was 0.84, and Bartlett's test strongly rejected the null hypothesis ($p < 0.001$), indicating that PCA is appropriate given the presence of significant correlations among the original variables.

PCA was then performed on standardised variables, since they are measured on different scales. Table 1 summarises component importance.

Only the first two components have eigenvalues greater than one (Kaiser criterion), and together they explain approximately 83% of the total variance.

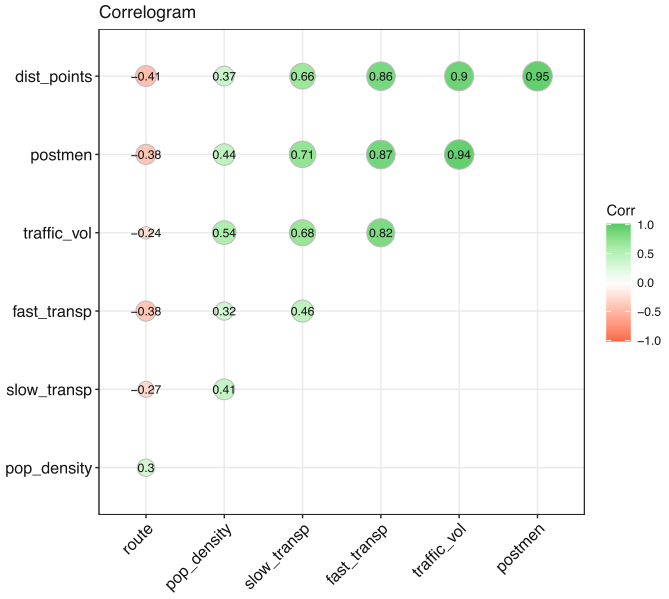

Fig. 1. Spearman rank correlations for the original variables.

Table 1. PCA component importance.

Importance of components	PC1	PC2	PC3	PC4	PC5	PC6	PC7
Standard deviation	2.162	1.059	0.808	0.614	0.316	0.215	0.162
Proportion of variance	0.668	0.160	0.093	0.054	0.014	0.007	0.004
Cumulative proportion	0.668	0.828	0.922	0.975	0.990	0.996	1.000

Therefore, the first two components were retained, and an orthogonal Varimax rotation was applied to improve interpretability.

Table 2 reports the loadings of the rotated principal components (RPCs). The percentage of variance explained by the first rotated principal component is approximately 65%, and the percentage of total variance explained by the first two rotated principal components is approximately 83%.

All variables have relevant weights (absolute loading ≥ 0.5) on at least one rotated component. The rotated components were interpreted as follows: RC1 has high negative loadings on all variables except `route`, and thus captures operational intensity and service density (more postmen, higher traffic, more distribution points, higher population density, and more slow-transport activity); RC2 has a strong positive loading on `route` and a moderate positive loading on `pop_density`, contrasting long/extensive routes with more compact delivery environments.

Table 2. Loadings after Varimax rotation.

	postmen	pop_density	fast_transp	slow_transp	dist_points	traffic_vol	route
RC1	-0.964	-0.753	-0.673	-0.889	-0.930	-0.964	0.196
RC2	-0.180	0.447	-0.289	0.000	-0.218	-0.108	0.918

3.4 Clustering PDCs Using PCA Scores

CA was performed on the scores of the two rotated components. Both hierarchical and partitional methods were considered, for different numbers of clusters, and the results were analysed to identify the best method.

A good cluster analysis should produce groups that are internally cohesive, well separated from each other, appropriately defined in number, and meaningful and useful for practical decision-making. Therefore, for each method, a graphical analysis of the cluster representation was carried out to assess separation between groups; the silhouette plot was also analysed to verify whether the number of negative silhouette values was reasonable; the average silhouette width was calculated to determine whether it was acceptable; and, finally, the practical meaningfulness of the clusters obtained was assessed.

Comparing the results, three approaches that produced coherent and operationally meaningful solutions were identified: Ward's method, k-means, and k-medoids. Across these methods, the most consistent solutions were obtained with seven clusters. The k-medoids solution differed from k-means in only five PDC assignments.

In what follows, only the k-means solution is reported, as it provided the most stable and interpretable clustering in this study. Table 3 reports cluster sizes and centroids in the rotated-component space.

Table 3. Number of PDCs per cluster and cluster centroids (rotated PC space).

Clusters	1	2	3	4	5	6	7
Sizes	9	29	64	5	32	61	13
RC1 means	0.547	-0.870	0.276	-2.279	0.528	0.543	-2.769
RC2 means	2.839	-0.471	-0.816	2.669	0.851	-0.085	0.377

Based on the rotated loadings (Tab. 2) and the cluster centroids (Tab. 3), the clusters can be characterised as follows:

– Clusters 4 and 7 have similarly high negative centroid values on RC1, but they differ in their centroid values on RC2. PDCs in both clusters have many postmen, high traffic, and many distribution points; they are located in areas with higher population density and rely more on slow transport. However, PDCs in cluster 4 have longer routes, whereas those in cluster 7 have more compact routes.

- Clusters 3, 5, and 6 have centroid values on RC1 that are positive and similar in magnitude, with smaller absolute values than those in clusters 4 and 7; they differ mainly in their centroid values on RC2. PDCs in these three clusters have fewer postmen, lower traffic, are located in areas with lower population density, and make less use of slow transport. They differ primarily in route length: routes are long in cluster 5, medium in cluster 6, and short in cluster 3.
- In cluster 1, the centroid value on RC1 is small, whereas the centroid value on RC2 is positive and the highest among all clusters. This corresponds to PDCs with very long routes located in areas with average population density.
- In cluster 2, the centroid values for the two RCs are both negative and similar in magnitude. This corresponds to PDCs with intermediate values in terms of the number of postmen, traffic volume, population density, and transport modes, but with more compact routes.

For confidentiality reasons, the real identifiers of the PDCs are not disclosed, and no geographical analysis of the cluster composition is presented. According to CTT, the results of this analysis are interesting and consistent with what was expected.

Figure 2 visualises the seven clusters and Fig. 3 shows the corresponding silhouette plot. The average silhouette width is approximately 0.41, with few observations showing negative silhouette values, suggesting that most PDCs are well assigned to their clusters.

The k-means objective diagnostics were also evaluated. It was verified that the within-cluster sums of squares are relatively small, and the ratio of between-cluster to total sum of squares is high (0.86), indicating substantial separation.

Overall, the seven-cluster k-means solution is supported by standard validity diagnostics and provides a coherent operational segmentation of PDCs, according to CTT's assessment.

3.5 Linear Discriminant Analysis and Multinomial Logistic Regression Classifiers

LDA was fitted using the rotated PC scores as discriminant variables and the k-means cluster labels as groups. The aim was to use this analysis to enable the classification of new PDCs into the predefined clusters. Since this method is based on assumptions, the initial step was to verify those assumptions. In this sense, multivariate normality within clusters and homogeneity of covariance matrices across clusters were assessed. Multivariate normality was supported in five of the seven clusters, whereas equality of covariance matrices was not. Despite these limitations, it was decided to proceed with the classification using LDA. The high classification performance reported below indicates that the violation of the assumptions had no major practical consequences.

However, MLR has advantages over LDA in contexts where its assumptions are potentially violated. In particular, it does not require the assumption of multivariate normality of the predictors or the equality of covariance matrices

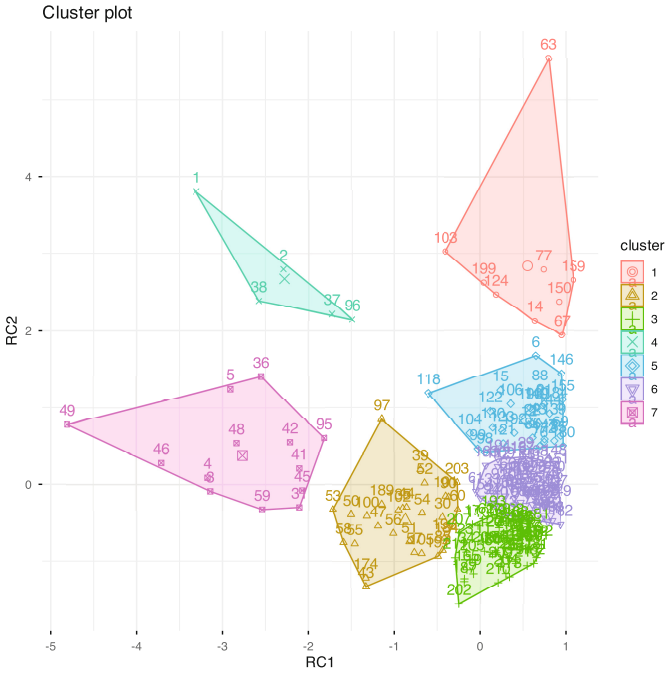

Fig. 2. Rotated PCA (RC1–RC2) visualisation of the k-means clusters.

between groups, making it more flexible in modelling real data. Additionally, because it does not depend on the estimation of group-specific covariance matrices, it tends to exhibit greater stability in situations where some clusters are small, as observed in this study. In this sense, the complementary use of this approach helps mitigate the limitations of LDA and reinforces the robustness of the proposed classification rule.

To evaluate classification accuracy, it is important to use an adequate cross-validation method. In this study, it is critical to mention that the data analysed correspond to all PDCs of the organisation, and therefore do not constitute a sample in the classical statistical sense. Given that each PDC has specific operational characteristics, the creation of independent training/test subsets could compromise the representativeness of the analysis. Therefore, the leave-one-out (jackknife) cross-validation method was adopted, which maximises the use of available information and allows direct evaluation of the stability of the assignment rule. Thus, classification accuracy was evaluated using this cross-validation procedure for both methods. The leave-one-out confusion matrices are presented in Table 4.

Inspection of the LDA confusion matrix shows that clusters 4 and 7 have no misclassified PDCs. The correct classification rates for clusters 1, 2, 3, 5, and 6 are 0.8889, 0.8276, 0.9844, 0.9375, and 0.9836, respectively. The overall correct classification rate is approximately 0.9531, indicating that the derived rule is

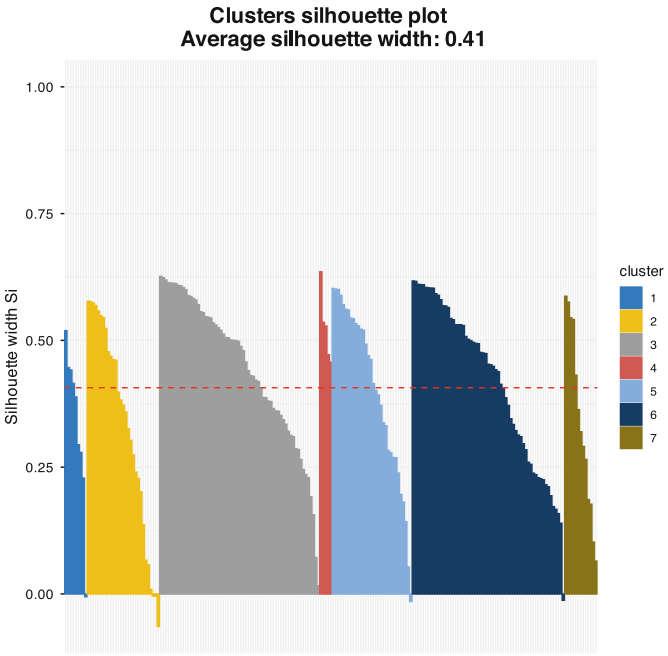

Fig. 3. Silhouette plot.

effective for assigning PDCs to clusters and can be used to classify new PDCs with low expected error.

Inspection of the MLR confusion matrix shows that the overall correct classification rate is approximately 0.9624. Clusters 1 and 4 have no misclassified PDCs. The correct classification rates for clusters 2, 3, 5, 6, and 7 are 0.9310, 0.9531, 0.9688, 0.9836, and 0.9231, respectively.

It can be noted that there are 5 PDCs that are in the same cluster for both classifiers.

To classify a new PDC using the LDA classifier, the RPC scores were computed by standardising the original variables and then applying the rotated loadings. The seven LDA classification functions were then evaluated, and each PDC was assigned to the cluster with the highest score.

To classify a new PDC using the MLR classifier, the RPC scores were first obtained by standardising the original variables and applying the rotated loadings. The fitted MLR model was then used to compute the predicted membership probabilities for the seven clusters, and each PDC was assigned to the cluster with the largest predicted probability.

Table 4. LOOCV (jackknife) confusion matrices for LDA and MLR classifications.

Prediction	LDA							MLR						
	1	2	3	4	5	6	7	1	2	3	4	5	6	7
1	8	0	0	0	1	0	0	9	0	0	0	0	0	0
2	0	24	4	0	0	1	0	0	27	0	1	0	0	1
3	0	0	63	0	0	1	0	0	1	61	0	0	2	0
4	0	0	0	5	0	0	0	0	0	0	5	0	0	0
5	0	0	0	0	30	2	0	0	0	0	0	31	1	0
6	0	0	1	0	0	60	0	0	0	1	0	0	60	0
7	0	0	0	0	0	0	13	0	1	0	0	0	0	12

3.6 Classification Improvement Using Leave-One-Out Validation of the Overall Pipeline: External Cross-Validation

LDA and MLR classifiers do not provide an independent validation of the clusters because the same dataset was used to derive the models and to assess their classification performance. Therefore, the high apparent accuracy should not be interpreted as a low expected error for new PDCs, since leave-one-out was applied only to the classifier, while standardisation, PCA, rotation, and clustering were kept fixed and were fitted using the full dataset.

The originally reported evaluation could lead to an optimistic estimate of predictive performance. It should be noted that only the final LDA and MLR classifications were subjected to leave-one-out validation, while earlier stages of the pipeline (standardisation, PCA, Varimax rotation, and clustering) were estimated using the full dataset. To address this concern, a fully consistent leave-one-out cross-validation procedure was applied to the entire modelling pipeline. In this framework, each iteration of the cross-validation loop re-estimates all data-dependent steps exclusively on the training subset, including standardisation, PCA, Varimax rotation, k-means clustering, LDA, and MLR. Note that cluster solutions are made comparable across resamples via label-switching resolution (i.e., label matching/alignment by permuting cluster labels to maximise agreement). This overall validation ensures that every observation is evaluated strictly out of sample with respect to the full modelling process, rather than only the final classification stage. For further details on the global pipeline, see, for example, [2,15,27].

Using the procedure presented above, the leave-one-out confusion matrices corresponding to the k-means clustering and to both classifiers are presented in Table 5. It is important to note that the cluster assignments obtained under the full leave-one-out cross-validation framework are not necessarily identical to those derived in the earlier analysis, where validation was restricted to the classification stage. This is because all components of the pipeline, including clustering, are re-estimated at each iteration. In fact, differences in cluster assignments across leave-one-out iterations are expected when the full pipeline is re-

estimated. All stages of the pipeline procedure are data-dependent, so small perturbations in the sample propagate through the feature space and may alter cluster centroids and memberships. Additionally, cluster labels are inherently subject to permutation invariance (label switching) and are therefore aligned across iterations using an optimal matching procedure. Consequently, observed differences reflect the intrinsic variability of the data-driven pipeline rather than methodological inconsistency.

Table 5. LOOCV (jackknife) confusion matrices for k-means, LDA and MLR classifications.

Prediction	K-1	2	M 3	E 4	A 5	N 6	S 7	L 1	D 2	A 3	4	5	6	7	M 1	L 2	R 3	4	5	6	7
1	11	0	0	0	0	0	0	10	0	0	0	1	0	0	8	0	0	0	3	0	0
2	0	76	0	0	0	0	0	0	76	0	0	0	0	0	0	76	0	0	0	0	0
3	0	0	21	0	0	0	0	0	0	21	0	0	0	0	1	0	19	0	0	1	0
4	0	0	0	42	0	0	0	0	10	0	32	0	0	0	0	0	0	39	0	1	2
5	2	0	0	0	5	0	0	4	0	0	0	3	0	0	2	0	0	0	5	0	0
6	0	2	0	0	0	45	0	0	10	0	0	0	37	0	0	2	0	1	0	44	0
7	0	0	0	0	0	0	9	0	0	0	1	0	0	8	0	0	0	0	0	0	9

Under full leave-one-out cross-validation of the global pipeline, the confusion matrix for the k-means baseline yields 209 correctly classified observations out of 213, corresponding to an overall accuracy of 98.1%, indicating strong stability of the clustering structure under small perturbations of the data. Based on the confusion matrices of both classifiers, the calculated accuracies were 87.8% for LDA and 93.9% for MLR. The gap between the accuracies of LDA and MLR suggests that the separation between clusters is not fully linear in the transformed feature space, with the more flexible multinomial model better capturing the underlying structure.

The percentages of correct classification for the global pipeline validation process, based on the matrices in Table 5, are lower than those reported in the analysis of Table 4, but still indicate a high level of agreement. This reduction confirms that the earlier performance estimates were indeed optimistic due to partial reuse of information in the preprocessing and clustering stages.

Importantly, the revised results provide a more realistic estimate of the expected generalisation performance of the global pipeline for new PDC observations.

4 Conclusions

This study addressed the classification of CTT PDCs using several multivariate statistical techniques, with the outputs of each step used as inputs to the next.

First, the aim was to group existing PDCs into homogeneous clusters based on the operational variables under study. Since the original variables exhibited strong correlations, PCA was applied to derive a small set of uncorrelated variables with minimal loss of information. The RPC scores were then used as inputs to CA, yielding a seven-cluster solution with good internal cohesion and satisfactory separation. This solution was supported by standard validity diagnostics and was operationally meaningful for the company.

The obtained clusters can be characterised as follows: two clusters consist primarily of PDCs with many postmen, high traffic, and many distribution points, located in areas with higher population density and relying more on slow transport; three clusters comprise PDCs with fewer postmen and lower traffic, located in areas with lower population density, and making less use of slow transport; one cluster consists of PDCs with very long routes in areas with average population density; and one cluster comprises PDCs with intermediate values for the number of postmen, traffic volume, population density, and transport modes, but with more compact routes.

LDA results indicate that the rotated PCs provide strong discrimination among the seven clusters, with a high overall correct classification rate and a low expected misclassification rate. Therefore, the derived LDA classification functions constitute a practical tool for assigning new PDCs to the established groups.

To assess the robustness of the classification rule obtained, the LDA results were compared with those of an alternative MLR model. This approach was considered particularly suitable given the violation of some LDA assumptions, namely homogeneity of covariance matrices across groups and multivariate normality, as well as the small size of some clusters. The results obtained through MLR were globally consistent with those obtained using LDA, showing a high degree of agreement in the assignment of PDCs to the different groups. This convergence between methodologies, based on distinct assumptions, reinforces the robustness of the identified segmentation and suggests that the captured cluster structure is stable.

To prevent overestimating accuracy due to reusing the same data across all stages of the global pipeline (standardisation, PCA, CA, LDA/MLR), leave-one-out cross-validation was implemented throughout the entire process rather than only during the classification step. This approach yields a more reliable estimate of the model's true performance.

Overall, the work illustrates how established multivariate methods can be combined into an effective decision-support framework for operational segmentation in an industrial setting, reinforcing the value of collaboration between academic modelling and business expertise.

Acknowledgments. This work was funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under CEAUL Research Unit, UID/00006/2025, DOI: https://doi.org/10.54499/UID/00006/2025, and by the European Union – NextGenerationEU through the project UID/PRR/00006/2025, DOI: https://doi.org/10.54499/UID/PRR/00006/2025. The authors thank CTT – Correios de Portugal,

S.A., for participating in the study and for authorising the publication of the information used in this paper.

Disclosure of Interest. The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this article. All analyses were conducted independently, and no external funding sources affected the study design, data collection, or interpretation of results. During the preparation of this manuscript, the authors used ChatGPT-5 (OpenAI) for the purpose of improving the clarity and readability of the manuscript. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

References

1. Aleixo, S.M., Pereira, B., Fernandes, T., Rodrigues, J., Pina, M.: Cluster analysis for two HoReCa channels of a large grocery retail company. In: Gervasi, O., Murgante, B., Garau, C., Taniar, D., Rocha, A., Faginas Lago, M. (eds.) Computational Science and Its Applications - ICCSA 2024, LNCS, vol. 14815, pp. 337–356. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-65154-0_21
2. Ambroise, C., McLachlan, G.J.: Selection bias in gene extraction on the basis of microarray gene-expression data. Proc. Natl. Acad. Sci. **99**(10), 6562–6566 (2002)
3. Box, G.E.P.: A general distribution theory for a class of likelihood criteria. Biometrika **36**, 317–346 (1949)
4. Box, G.E.P.: Non-normality and tests on variances. Biometrika **40**, 318–335 (1953)
5. Brentari, E., Dancelli, L., Manisera, M.: Clustering ranking data in market segmentation: a case study on the Italian McDonald's customers' preferences. J. Appl. Stat. **43**(11), 1959–1976 (2016)
6. Bulut, H., Oner, Y.: The evaluation of socio-economic development of development agency regions in Turkey using classical and robust principal component analyses. J. Appl. Stat. **44**(16), 2936–2948 (2017)
7. CTT – Correios de Portugal, S.A (2025). https://www.ctt.pt/
8. Everitt, B.S., Landau, S., Leese, M., Stahl, D.: Cluster Analysis. John Wiley & Sons, United Kingdom (2011)
9. Gardner-Lubbe, S.: Linear discriminant analysis for multiple functional data analysis. J. Appl. Stat. **48**(11), 1917–1933 (2021)
10. Gelman, A., Hill, J.: Data Analysis Using Regression and Multilevel/Hierarchical Models. Cambridge University Press (2006)
11. Hair, J.F., Jr., Black, W.C., Babin, B.J., Anderson, R.E.: Multivariate Data Analysis. Pearson Education Limited, London (2013)
12. Hennig, C., Meila, M., Murtagh, F., Rocci, R.: Handbook of Cluster Analysis. Chapman and Hall/CRC, Boca Raton, Florida (2016)
13. Johnson, R.A., Wichern, D.W.: Applied Multivariate Statistical Analysis. Pearson Prentice Hall, Hoboken, New Jersey (2007)
14. Jolliffe, I.T.: Principal Component Analysis. Springer-Verlag, New York (2002)
15. Kuhn, M., Johnson, K.: Applied Predictive Modeling. Springer (2013)
16. McLachlan, G.J.: Discriminant Analysis and Statistical Pattern Recognition. John Wiley & Sons, Hoboken, New Jersey (2005)
17. Rai, P., Singh, S.: A survey of clustering techniques. Int. J. Comput. Appl. **7**, 1–5 (2010)

18. Reddy, C.K., Aggarwal, C.C.: Data Clustering: Algorithms and Applications. Chapman and Hall/CRC, Boca Raton, Florida (2014)
19. Rencher, A.C., Christensen, W.F.: Methods of Multivariate Analysis. John Wiley & Sons, Hoboken, New Jersey (2012)
20. Royston, P.: Some techniques for assessing multivariate normality based on the Shapiro-Wilk W. Appl. Stat. **32**, 121–133 (1983)
21. Royston, P.: Approximating the Shapiro-Wilk W-test for non-normality. Stat. Comput. **2**, 117–119 (1992)
22. Royston, P.: A toolkit for testing for non-normality in complete and censored samples. The Statistician **42**(1), 37–43 (1993)
23. The R Project for Statistical Computing (2025). https://www.r-project.org/
24. Shapiro, S.S., Wilk, M.B.: An analysis of variance test for normality (complete samples). Biometrika **52**(3/4), 591–611 (1965)
25. Srivastava, M.S.: Some tests concerning covariance matrices in high-dimensional data. J. Jpn. Stat. Soc. **35**(2), 251–272 (2005)
26. Thode, H.: Testing for Normality. CRC Press, New York (2002)
27. Varma, S., Simon, R.: Bias in error estimation when using cross-validation for model selection. BMC Bioinform. **7**(91), 1–8 (2006)



Lightweight Transformer-Based Event Embeddings with Log-Frequency Session Weighting for Purchase Prediction in E-Commerce Clickstream Data

Deniz Guzelcay[1(✉)], Sefa Çayır[1], Konuralp Bol[1], Halil İbrahim Simsek[2], Onur Cilgin[2], and Mehmet S. Aktas[1]

[1] Yildiz Technical University, Istanbul, Turkey
{deniz.guzelcay,sefa.cayir,konuralp.bol}@std.yildiz.edu.tr, aktas@yildiz.edu.tr
[2] Trendyol R and D Center, Istanbul, Turkey
{halilibrahim.simsek,onur.cilgin}@trendyol.com

Abstract. E-commerce platforms generate large volumes of clickstream data in the form of short user sessions. Predicting whether the next event in a session will be a purchase is a useful task for recommendation, marketing analytics, and real-time decision support. Many existing methods address this problem with complex sequential deep learning models, but such models often require substantial computational resources and large training datasets. This limits their practicality in resource-constrained settings. In this work, we study a lightweight alternative based on event embeddings and frequency-aware session representation. Each clickstream event token, such as a page view or user action, is encoded as a dense vector using a lightweight transformer-based embedding model. To reflect repeated behavior within a session, we apply a logarithmic frequency-weighting scheme so that recurring events contribute more strongly to the final session representation without dominating it. The weighted event embeddings are then aggregated into a fixed-length session vector and used as input to supervised classifiers for next-event purchase prediction. The approach is evaluated on both a public clickstream dataset and a large real-world e-commerce dataset. Experimental results show that frequency-weighted session representations perform better than unweighted representations across different embedding models and classifiers. The gains are more evident in highly imbalanced settings where purchase events are rare. These findings indicate that lightweight transformer-based embeddings combined with log-frequency weighting offer a practical approach for next-event purchase prediction in large-scale clickstream environments.

Keywords: Clickstream Data Analysis · Purchase Prediction · Event Embeddings · Session Representation · Lightweight Transformers · Frequency Weighted Aggregation

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 242–259, 2027.
https://doi.org/10.1007/978-3-032-30518-3_16

1 Introduction

E-commerce platforms now produce large amounts of interaction data as users browse, search, view products, and add items to their carts. These actions are typically stored as clickstream logs. When viewed over short periods, they form user sessions that can give clues about what the user intends to do next. This makes session data useful for several tasks, including recommendation, advertising, and behavior analysis. A particularly important question is whether a session will result in a purchase. Since purchases are directly related to revenue, predicting them has become a common problem in both data mining and recommender systems [1,2]. The present work approaches this problem from the angle of efficient session representation, looking for a balance between predictive accuracy and the computational cost of deployment in large-scale platforms.

Research on e-commerce has examined both purchase intent prediction and user behavior modeling. Early studies largely relied on statistical methods and manually designed behavioral features derived from clickstream logs, such as visit counts, session duration, and navigation depth [1]. These approaches were computationally efficient, but they often fell short in capturing the more complex patterns of user navigation. To improve predictive performance, later research began applying machine learning models to deeper feature sets [2]. Over time, representation learning methods and neural architectures also became more common for modeling the relationships among behavioral signals. Neural recommender models and factorization-based approaches, were used in click-through rate and conversion prediction and showed promising results [3,4]. At the same time, sequential modeling techniques, including recurrent neural networks and transformer-based architectures, were increasingly adopted to capture temporal dependencies in session-based recommendation and next-event prediction tasks [5].

These developments improved predictive performance, but they also raised practical concerns. Many deep sequential models deliver strong results, yet they are not always easy to use in real deployments. Training and deploying them often requires considerable computing power, enough labeled data, and fairly involved inference pipelines. In production environments, where fast and efficient predictions are essential, these requirements can become a real limitation. Their complexity can also make deployment and maintenance harder, especially on large commercial platforms. At the other end, lightweight session modeling methods are easier to work with, but they often oversimplify user behavior. When they rely on basic pooling or hand-crafted features, they may overlook the meaning of different event types and miss important behavioral signals in the interaction sequence [6,7]. As a result, there is a need for session representation methods that can do both: capture meaningful patterns in user behavior and still remain efficient and scalable in real-world settings. This need shapes the central question of the present study, namely how to build session representations that are informative enough for next-event purchase prediction while remaining computationally light. The remainder of this section reviews related embedding-based work and then introduces the proposed approach in more detail.

Several of the past studies have looked into using embedding techniques to model user behavior in clickstream data. For instance, researchers have used graph-based and natural language processing (NLP) inspired embedding methods such as Word2Vec, Node2Vec, and DeepWalk. These methods help represent user navigation sessions and analyze behavioral patterns in e-commerce platforms [8]. In later research, researchers also looked at embedding-based representations in system-level applications. These included predictive prefetching and smart data delivery in data-heavy client-server systems [9]. Several other studies have investigated the extent to which weighting methods can enhance embedding quality. One example is event frequency, inverse session frequency (EF-ISF) weighting, which helps in finding behavioral patterns [10] and aging-based temporal weighting strategies for improving session classification performance [11]. While these studies demonstrated the effectiveness of embedding techniques for representing clickstream data and modeling user behavior, they primarily focused on improving representation quality or clustering and classification performance through various embedding and weighting strategies. By contrast, our focus is on purchase prediction. To address it, we use a lightweight session representation that combines computationally efficient transformer-based event embeddings with log-frequency weighting to capture behavioral intensity within user sessions.

Within this context, the task we focus on is predicting whether the next event in a user session is a purchase, based on the navigation behavior observed up to that point. The central research problem is how to construct informative and efficient session-level representations from clickstream events that support this prediction task. More specifically, we examine whether repeated user actions within a session are associated with purchase intent and how this information can be incorporated into the session representation. To tackle this research problem, we need a way to represent the semantic features of event tokens. We also need to capture the intensity of behavior shown in repeated interactions during sessions.

Our approach to this challenge is to combine lightweight transformer-based event embeddings with a logarithmic frequency-weighting strategy at the session level. In the proposed method, each clickstream event token representing a page or user action is converted into a dense numerical vector. This transformation uses pretrained transformer-based embedding models, such as Sentence-BERT and related architectures [12–14]. Unlike large transformer models used in sequential recommendation tasks, the embedding models in this study are easy to run and fit for settings with limited computational resources. After event embeddings are generated, session-level representations are built by combining event vectors. This uses a logarithmic frequency-weighting method that boosts the impact of repeated events. It also limits excessive dominance with logarithmic scaling. In this way, the representation of the session captures both what the events represent and how strongly certain behaviors appear within the session.

To assess the approach, we run experiments on a public clickstream dataset and on a large-scale real-world e-commerce dataset. Event embeddings generated from lightweight transformer models are aggregated into session-level vec-

tors using both standard unweighted aggregation and the proposed log-frequency weighted aggregation strategy. These session vectors are then used as input features for supervised classification models, including Random Forest [15] and XGBoost [16]. To ensure reliable evaluation and prevent data leakage across user sessions, a stratified group cross-validation strategy is used during model training and testing. Model performance is evaluated using widely adopted classification metrics such as F1-score and accuracy, which are particularly suitable for evaluating classification tasks involving imbalanced datasets [17].

The main contributions of this study can be summarized as follows. First, we introduce a lightweight session representation framework that converts clickstream event tokens into numerical vectors using efficient transformer-based embedding models. Second, we propose a logarithmic frequency-weighting method that incorporates session-level event frequency information in creating session representations. Third, we demonstrate, through extensive experiments on both public and real-world datasets, that the frequency-weighted session representation significantly improves purchase prediction performance compared to standard unweighted aggregation methods. Finally, the results indicate that our approach offers an effective and efficient alternative to complex sequential models for purchase prediction in large-scale e-commerce settings.

The rest of this paper is organized as follows. Section 2 presents a review of the literature on session-based user behavior modeling, embedding-based clickstream representations, and the prediction of purchase intent. Section 3 introduces the proposed frequency-weighted session embedding framework, including the generation of event embeddings using lightweight transformer models and the log-frequency weighting method for creating session representations. Section 4 describes the experimental study, which includes the datasets, preprocessing steps, experimental setup, and evaluation methods. Additionally, this section discusses the experimental results, interpretation of the findings, and possible threats to validity. Finally, Sect. 5 concludes the paper and outlines several directions for future research.

2 Literature Review

Modeling user behavior from session-based interaction data has become a central research problem in recommender systems, clickstream analytics, and e-commerce intelligence. In many online platforms, user interactions occur in short browsing sessions without persistent user identifiers or long-term historical records. Predicting user intent from limited session-level behavior signals has gained a lot of interest in both research and business. This problem has been investigated through several modeling approaches, ranging from sequential neural architectures and embedding-based methods to graph-based models, frequency-aware session representations, and transformer-based sequential recommenders. Each of these approaches aims to capture different aspects of user behavior. They look at the order of events, the meaning behind their connections, and the different levels of importance of interactions during a session [7].

Early work on session-based recommendation mainly used recurrent neural networks to model the sequence of user interactions. A key contribution in this area was the GRU4Rec model introduced by Hidasi et al. [18]. It showed that gated recurrent neural networks can model time patterns in user sessions effectively. Later studies built on this idea and suggested various ways to improve sequential modeling. Li et al. [19], for example, proposed the Neural Attentive Recommendation Machine (NARM), which uses an attention mechanism to focus on the most informative items within a session. Quadrana et al. [20] later proposed hierarchical recurrent neural networks to capture patterns at both the session level and the user level. Later surveys highlighted the importance of sequential patterns in recommendation systems and provided a detailed classification of sequence-aware recommender algorithms [7]. These architectures work well for modeling user interactions over time. However, they often require a lot of training data and computing power.

To move beyond the limits of purely sequential models, several studies have explored embedding-based representation learning methods for modeling session data. Improvements in natural language processing inspire this approach. Sessions can be seen like sentences. Interaction events serve as words. In this context, embedding algorithms such as Word2Vec are often used to learn dense vector representations of items. These representations depend on how frequently the items appear together in sessions [21]. For instance, Greenstein-Messica and Rokach [22] proposed a session-based recommendation approach that learns item embeddings directly from session sequences. Similarly, graph embedding techniques such as DeepWalk [23] and Node2Vec [24] have been applied to interaction graphs constructed from session transitions. These approaches help show hidden similarities between items and uncover behavioral patterns across sessions. However, many embedding-based methods treat all events in a session equally. They do not consider that some interactions might be more important than others.

Several studies addressed this issue by adding weighting methods to session representations. Similar to TF-IDF in information retrieval, these methods assign more weight to important events and less to interactions that happen often but provide little information. In clickstream and behavioral modeling tasks, frequency-based weighting can also detect signals of behavioral intensity. This may indicate a stronger user interest or intent to buy. Event frequency-inverse session frequency is one example and has been explored for representing behavioral patterns and improving session classification performance [11]. The limitation is that these methods usually work as stand-alone statistical weighting schemes rather than as part of the embedding process itself. As a result, they can highlight important interactions, but they may still miss semantic relationships between different kinds of events.

More recent research has attempted to overcome these limitations by integrating attention mechanisms into sequential recommendation models. Attention-based architectures dynamically estimate the importance of events within a session and allow neural models to focus on interactions that are most predictive of user intent. For example, Li et al. [19] proposed an attention-

based recurrent model that learns both global session representations and local item importance weights. Similarly, self-attention mechanisms have been widely adopted in sequential recommendation models to capture contextual dependencies between interaction events [25]. These models improve representation learning by allowing the model to selectively focus on informative elements within the sequence. However, attention-based architectures typically rely on complex neural networks that require substantial computational resources for training and inference.

Graph-based approaches have also emerged as an effective solution for modeling session-based interactions. In graph-based session models, items are represented as nodes and transitions between items are modeled as edges in a directed graph structure. Graph neural networks can then be used to learn representations that capture complex relationships between interaction events. One of the most influential works in this area is SR-GNN proposed by Wu et al. [26], which models sessions as directed graphs and applies graph neural networks to capture item transition patterns. Graph-based approaches have demonstrated strong performance in session-based recommendation tasks because they explicitly model structural relationships between items. However, graph neural architectures often introduce additional computational complexity due to graph construction and message-passing operations.

More recently, transformer-based architectures have been introduced for modeling sequential user behavior in recommendation systems. These approaches leverage self-attention mechanisms to capture long-range dependencies between events in user sessions. For example, Sun et al. [5] proposed BERT4Rec, a transformer-based recommendation model that learns bidirectional representations of item sequences. Kang and McAuley [25] proposed SASRec, which applies self-attention to model sequential user behavior. In addition, industrial recommendation systems like Transformers4Rec have been created to connect natural language processing methods with large-scale recommendation processes [27]. While transformer-based architectures often achieve state-of-the-art performance, they generally require large training datasets and significant computational resources.

Several of our previous studies focused on embedding-based ways of modeling user navigation in clickstream environments. Earlier work examined graph-based and NLP-inspired methods, such as Word2Vec, Node2Vec, and DeepWalk, for representing user navigation sessions and analyzing behavioral patterns on e-commerce platforms. Subsequent studies examined embedding-based predictive prefetching strategies for intelligent data delivery mechanisms in data-intensive client-server systems. Other studies investigated the integration of this work by considering weighting mechanisms within embedding representations, including event frequency–inverse session frequency weighting and aging-based temporal weighting strategies for improving session classification [11]. Although these studies supported the use of embedding and weighting methods for user behavior modeling, they focused on representation learning and session classification rather than purchase prediction.

Overall, the current research shows that session-based behavioral modeling improves with sequential neural architectures, embedding-based representations, graph-based interaction modeling, and attention-based weighting methods. However, several challenges persist. Many leading methods depend on deep neural architectures that require a lot of computing power. This makes them unsuitable for real-time use in large-scale e-commerce settings. Additionally, embedding-based representations often do not clearly include signals of behavioral intensity, such as how often events happen within sessions. On the other hand, statistical weighting methods point out important interactions but overlook the semantic relationships between events. To tackle these issues, this study suggests a lightweight session representation framework. It combines transformer-based event embeddings with a logarithmic frequency-weighting method. This approach allows for the development of efficient and informative session representations for predicting purchases in large-scale clickstream environments.

3 Methodology

This section presents the proposed framework for constructing frequency-weighted session representations from clickstream event sequences for purchase intent prediction. The primary goal of the framework is to transform user navigation events into informative numerical representations that capture both the semantic properties of interaction events and the behavioral intensity reflected by repeated actions within a session. Instead of training computationally expensive sequential deep learning architectures such as recurrent neural networks or transformer-based sequence models, the proposed approach utilizes pretrained lightweight transformer-based embedding models to generate semantic representations of event tokens. These representations are then aggregated using a frequency-aware weighting strategy to construct session-level embeddings. As illustrated in Fig. 1, the proposed framework consists of three main stages:

1. Event representation learning using pretrained transformer-based embeddings
2. Frequency-based weighting of session events
3. Purchase prediction using supervised classification models

Event Token Embedding Generation: User navigation sessions consist of sequences of interaction events such as page visits, product views, searches, and cart interactions. In the proposed framework, each event type is treated as a token similar to words in natural language processing. This representation enables the application of embedding techniques to capture semantic relationships between interaction events. To generate numerical representations of event tokens, we employ pretrained lightweight transformer-based embedding models. Specifically, three Sentence-BERT (SBERT) based models are utilized in this study to balance computational efficiency and representation quality:

- *all-MiniLM-L6-v2* [28]
- *intfloat/e5-small-v2* [13]
- *BAAI/bge-small-en-v1.5* [14]

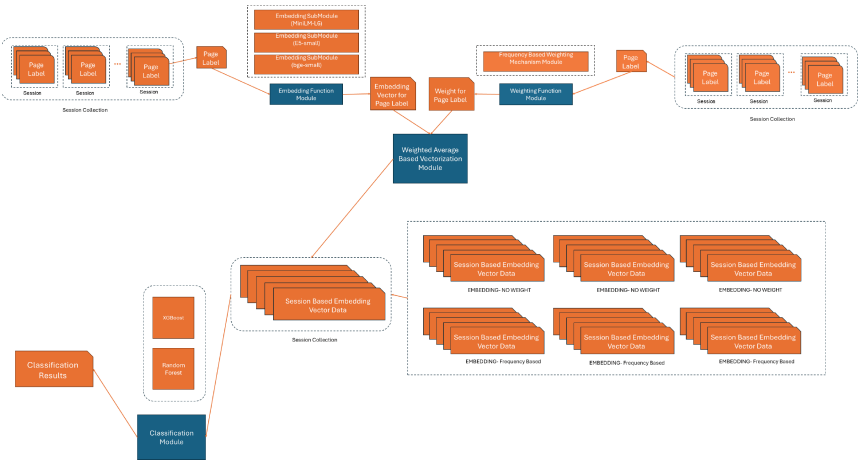

Fig. 1. Overview of the proposed framework.

Given an event token e_i, the embedding model $M(\cdot)$ generates a dense vector representation $v_i \in \mathbb{R}^D$ as follows:

$$v_i = M(e_i) \tag{1}$$

These embeddings position related events near each other in the vector space. This helps capture the similarities in interaction patterns.

Log-Frequency Weighting Mechanism: Embedding models are useful for capturing semantic relationships between event types, but standard aggregation methods like simple averaging still assume that every event contributes in the same way to the final session representation. In practice, this is often not the case. When users interact repeatedly with certain event types, those actions may carry stronger behavioral meaning or indicate higher intent. To reflect this more clearly, we incorporate a logarithmic frequency-weighting strategy into the session representation. Let $freq(e_i)$ denote the number of occurrences of event e_i within session S. The importance weight assigned to event e_i is defined as

$$w_i = \log(1 + freq(e_i)) \tag{2}$$

This logarithmic transformation increases the importance of repeated events while limiting the influence of events that occur very often.

Frequency-Weighted Session Embedding Construction: Once the event embeddings and frequency weights are available, the final session representation is formed through weighted aggregation of the event vectors. For a session $S = \{e_1, e_2, ..., e_n\}$, the session embedding V_S is computed as the normalized weighted sum of event embeddings:

$$V_S = \frac{\sum_{i=1}^{n} w_i \cdot v_i}{\sum_{i=1}^{n} w_i} \tag{3}$$

This formulation gives repeated events more influence on the final representation while keeping the values numerically stable.

Frequency-Weighted Session Embedding Algorithm: A summary of the proposed frequency-weighted procedure for constructing session embeddings is provided in Algorithm 1.

Algorithm 1. Frequency Weighted Session Embedding

Input: Session $S = \{e_1, e_2, ..., e_n\}$
Output: Session embedding vector V_S
Initialize $V_S \leftarrow 0$
Initialize $W \leftarrow 0$
for each event e_i in session S **do**
    $v_i \leftarrow M(e_i)$
    $w_i \leftarrow \log(1 + freq(e_i))$
    $V_S \leftarrow V_S + w_i \cdot v_i$
    $W \leftarrow W + w_i$
end for
$V_S \leftarrow V_S / W$
return V_S

The algorithm operates in two main stages. First, each event token in the session is transformed into a dense semantic vector using the selected embedding model $M(\cdot)$. These vectors capture semantic relationships between different interaction types in the clickstream. Second, a logarithmic frequency weighting is applied. This allows repeated events to contribute more to the session representation. However, it does not let very frequent events dominate the session. After that, the event vectors are combined using normalized weighted aggregation. This creates the final session embedding. It creates a clear representation that captures both semantic similarity and also the frequency of interactions. It also provides downstream machine learning models with a better basis for using session-level behavioral signals to predict purchases from large-scale clickstream datasets.

Purchase Prediction Framework: The session embeddings created in the earlier stage give short numerical descriptions of how users navigate. These representations then serve as an input feature for supervised learning algorithms meant to predict purchase intent. The framework transforms clickstream sessions of varying length into fixed-length vectors, which allows traditional machine learning methods to be used more efficiently without explicit sequence modeling. In this way, it provides a computationally simpler alternative to complex sequential deep learning models, while still retaining meaningful information about user navigation behavior.

Using frequency-weighted embeddings helps the framework retain behavioral cues that are often lost in traditional session modeling. Repeated interactions with certain event types may signal stronger intent or deeper interest in particular products or categories. When semantic event representations are combined

with frequency-based weighting, the resulting session vectors capture both the nature of the interactions and their intensity. This leads to a more informative feature space for supervised learning models. It allows classifiers to better differentiate between sessions that result in purchases and those that do not. As a result, the proposed framework makes it easier to predict purchase intent while keeping computational efficiency suitable for large-scale clickstream environments.

Computational Complexity Analysis: An important advantage of the proposed framework is its computational efficiency compared to sequential deep learning models that explicitly model event order within sessions. Let n denote the number of events in a session and d represent the embedding dimensionality. The generation of event embeddings requires a constant-time inference operation for each event using the pretrained embedding model. Therefore, the embedding stage has a computational complexity of $O(n)$ per session. The frequency-weighting step requires computing the frequency of each event within the session and applying the logarithmic transformation. Since these operations are performed once per event, this step also has a complexity of $O(n)$. The weighted aggregation of event embeddings requires a linear combination of n vectors of dimension d, resulting in a complexity of $O(n.d)$ per session. Consequently, the overall computational complexity of constructing a session embedding using the proposed framework is $O(n.d)$ which scales linearly with the number of events in a session. This linear complexity makes the proposed approach significantly more efficient than sequential neural architectures such as recurrent neural networks or transformer-based sequence models, which typically exhibit quadratic or higher computational costs with respect to sequence length. Therefore, the proposed framework provides a computationally efficient and scalable solution for large-scale clickstream analysis environments.

4 Experimental Study

This section presents the experimental evaluation of the proposed frequency-weighted session embedding framework. The objective of the experimental study is to investigate whether incorporating session-level event frequency information into embedding aggregation improves purchase prediction performance compared to conventional unweighted session representations. The experimental evaluation focuses on three main research questions:

- **RQ1:** Does frequency-weighted aggregation improve purchase prediction performance compared to standard unweighted embedding aggregation?
- **RQ2:** Is the proposed approach robust across different lightweight transformer embedding models?
- **RQ3:** How does the effectiveness of frequency weighting change under different class imbalance conditions?

These questions were investigated through experiments using both a public clickstream dataset and a large-scale proprietary e-commerce dataset.

Datasets: Two clickstream datasets with similar structural properties but different behavioral characteristics were used in the experiments. The first dataset contains user activity logs from a mobile coffee shop application published on Kaggle. Each record corresponds to a user session and includes information about the user identifier, page interactions, timestamp information, and other contextual attributes. A sample record from the dataset is shown in Table 1. This example illustrates how a single interaction event is represented using user-level contextual attributes.

Table 1. Sample User Data

User Information	Page Information	Day of the Week	Hour
0096610	002	03 (Wednesday)	23

Because user actions in sessions are recorded as sequential events, this dataset provides a suitable framework for session-based modeling. The dataset contains 276,566 sessions and 46 distinct event types performed by 96,200 users. The diversity of interaction types makes the dataset appropriate for evaluating embedding-based behavioral modeling approaches.

The second dataset consists of large-scale user clickstream data obtained from the real-life e-commerce platform. Each record contains a session identifier, a sequence duration array describing time intervals between interactions, and a sequence array representing the ordered list of user actions. A representative example of a session record from the real-life e-commerce website dataset is shown in Table 2. The table illustrates how a session is represented as a sequence of interaction events and corresponding temporal intervals.

Table 2. Sample from real-life e-commerce dataset

Field	Value
session_id	0000585b-f366-417b-aead-23d14b03a433
sequence_duration	0, 4, 1402, 2, 11, 4, 39, 3, 47, 3, 19, 4, 2, 3, 45, 4
sequence	PRODUCT, OTHER, PRODUCT, PRODUCT, PRODUCT, PRODUCT, PRODUCT, MAINPAGE, PRODUCT, OTHER, OTHER, SEARCH, SEARCH, SEARCH, SEARCH, SEARCH

Unlike the public dataset, the real-life e-commerce website dataset contains pre-labeled purchase outcomes, enabling direct supervised learning experiments without additional annotation. The statistical comparison of both datasets is presented in Table 3. As shown in the table, the datasets differ in several important ways, especially in terms of the severe class imbalance in the proprietary dataset and the longer average session length in the coffee shop dataset.

Table 3. Statistical Comparison of Datasets

Metric	Real-life e-commerce website	Coffee Shop
Total Sessions	547,715	276,566
Purchase Ratio (Balance)	1.08% (Highly Imbalanced)	54.38% (Balanced)
Vocabulary Size	37 Events	46 Events
Avg. Session Length	16.6 Events	156.6 Events

Using both balanced and highly imbalanced datasets allows us to test how robust and generalizable the proposed representation method is under different behavioral distributions.

Data Preprocessing: Clickstream datasets often include interactions that take place after a purchase. If we keep these events during training, the model may focus on the purchase itself instead of the behaviors that lead to it. To prevent this issue, we added a preprocessing step to end sessions right after the purchase event. For sessions with a purchase label, we removed all events that happened after the purchase. This step ensures the model only sees behaviors before the purchase and prevents data leakage during training.

Experimental Setup: The experimental prototype was built in Python 3.14 using popular data science libraries. We used *Pandas* and *NumPy* for data preprocessing and session manipulation. Event embeddings were created with pretrained models from the *SentenceTransformers* framework. We reduced embedding dimensionality to 64 using a dense projection layer built with *PyTorch*. Two widely used tree-based classification algorithms were used: Random Forest and XGBoost. These classifiers are reliable in high-dimensional feature spaces and perform well on tabular prediction tasks.

Cross-Validation Strategy: To get a reliable performance estimate, we used a 10-Fold Stratified Group Cross-Validation strategy with *StratifiedGroupKFold*. This evaluation strategy addresses two important issues in clickstream prediction: Maintaining class distribution across folds, Preventing data leakage across users. By grouping sessions by user identifier, interactions from the same user cannot appear in both training and testing partitions.

Evaluation Metrics: Two standard classification metrics were used: F1-score (Purchase class), Accuracy. Because purchase prediction represents an imbalanced classification task in real-world e-commerce environments, the F1-score for the purchase class was used as the primary evaluation metric.

Experimental Results: The experimental results obtained on the proprietary real-life e-commerce website dataset are summarized in Table 4. The table compares the performance of weighted and unweighted session representations across three embedding models and two classifiers.

To further highlight the effect of frequency weighting, Table 5 presents the relative improvement in F1-score obtained by the proposed method compared to the unweighted baseline.

Table 4. Performance on Proprietary(Real-life e-commerce website) Dataset

Model	Clf.	Method	F1-Score(Purchase)	Accuracy
MiniLM	RF	Unweighted	0.3759	0.9681
		Weighted	**0.6139**	**0.9924**
	XGB	Unweighted	0.3736	0.9673
		Weighted	**0.6828**	**0.9910**
E5-small	RF	Unweighted	0.3762	0.9681
		Weighted	**0.6100**	**0.9923**
	XGB	Unweighted	0.3732	0.9673
		Weighted	**0.6858**	**0.9911**
BGE-small	RF	Unweighted	0.3759	0.9681
		Weighted	**0.6142**	**0.9924**
	XGB	Unweighted	0.3728	0.9673
		Weighted	**0.6828**	**0.9910**

Table 5. Relative F1-Score Improvement on Proprietary Dataset

Model	Classifier	Improvement (%)
MiniLM	Random Forest	+63.31%
	XGBoost	**+82.76%**
E5-small	Random Forest	+62.15%
	XGBoost	**+83.76%**
BGE-small	Random Forest	+63.39%
	XGBoost	**+83.15%**

Similar experiments were conducted on the coffee shop dataset. The results are presented in Table 6.

Relative improvements for the balanced dataset are summarized in Table 7.

Discussion: The results suggest that frequency-aware weighting helps session embeddings perform better in purchase prediction. This pattern is particularly noticeable in the proprietary real-world e-commerce dataset, where purchase events are very rare. In that setting, repeated interactions with certain event types appear to signal stronger purchase intent. The logarithmic weighting scheme captures some of this intensity, which seems to help the classifier distinguish purchase sessions from non-purchase ones more effectively. The gain is smaller in the coffee shop dataset, which has a more balanced class distribution. Even so, the weighted version still performs better than the unweighted baseline across all embedding models and classifiers considered in the study.

Behavioral Interpretation: Beyond performance improvements, the experimental findings provide important insights into user behavior. When users interact multiple times during a session, it usually indicates they are more

Table 6. Performance on Coffee Shop Dataset

Model	Clf.	Method	F1-Score(Purchase)	Accuracy
MiniLM	RF	Unweighted	0.6487	0.6673
		Weighted	**0.7374**	**0.7359**
	XGB	Unweighted	0.6485	0.6672
		Weighted	**0.7605**	**0.7506**
E5-small	RF	Unweighted	0.6484	0.6671
		Weighted	**0.7369**	**0.7355**
	XGB	Unweighted	0.6489	0.6674
		Weighted	**0.7607**	**0.7511**
BGE-small	RF	Unweighted	0.6485	0.6672
		Weighted	**0.7366**	**0.7353**
	XGB	Unweighted	0.6484	0.6670
		Weighted	**0.7597**	**0.7502**

Table 7. Relative F1-Score Improvement on Coffee Shop Dataset

Model	Classifier	Improvement (%)
MiniLM	Random Forest	+13.67%
	XGBoost	**+17.27%**
E5-small	Random Forest	+13.65%
	XGBoost	**+17.23%**
BGE-small	Random Forest	+13.59%
	XGBoost	**+17.16%**

engaged with specific products or categories. For example, users might look at product pages several times, do multiple searches, or go back to certain categories before making a purchase. The suggested frequency-weighted representation captures this behavior by giving more importance to repeated events while using logarithmic scaling to avoid excessive influence. As a result, the created session embeddings reflect both the relationships among interaction events and the signals of behavioral intensity that indicate user intent.

Threats to Validity of the Results: Several potential threats to validity should be considered when interpreting the experimental results. First, although two datasets were used, both datasets originate from e-commerce or mobile application environments. Therefore, the generalizability of the proposed method to other domains of sequential user behavior remains an open question. Second, the experiments rely on pretrained embedding models without domain-specific fine-tuning. While this choice ensures computational efficiency, fine-tuned models may potentially achieve higher predictive performance. Third, the evaluation was conducted using offline experiments. Future work should evaluate the app-

roach in online production environments to assess its real-time performance and scalability.

5 Conclusion and Future Work

This study addressed the problem of predicting purchase intent from session-based clickstream data in e-commerce environments. Modeling user behavior from short interaction sessions remains a challenging task because clickstream sequences often contain limited contextual information while exhibiting complex behavioral patterns. Existing approaches frequently rely on computationally expensive sequential neural architectures, which may limit their applicability in large-scale production systems. To overcome these limitations, this work proposed a lightweight session representation framework that integrates transformer-based event embeddings with a logarithmic frequency-weighting mechanism.

Returning to the three research questions stated in Sect. 4, the experimental evidence can be summarized as follows. Regarding RQ1, the log-frequency weighted aggregation produced higher F1-scores than the unweighted baseline in every configuration tested in this study. As reported in Tables 5 and 7, the relative F1 improvements ranged from approximately 13% on the balanced coffee shop dataset to over 80% on the highly imbalanced proprietary dataset. Regarding RQ2, the gain was consistent across the three lightweight transformer embedding models considered (MiniLM, E5-small, and BGE-small); the differences between these embedding models were small relative to the gap between the weighted and unweighted variants. This suggests that the weighting scheme, rather than the choice of embedding model, is the main driver of the observed improvement. Regarding RQ3, the effect of frequency weighting was clearly larger on the imbalanced dataset, where purchase events are rare, than on the balanced dataset. A plausible interpretation is that, when positive examples are scarce, the additional signal carried by repeated intent-bearing events becomes more discriminative for the classifier.

The proposed framework constructs session-level representations by first transforming clickstream event tokens into dense semantic vectors using pretrained lightweight transformer embedding models. These event representations are then aggregated through a log-frequency weighting strategy that emphasizes repeated interactions within a session while preventing extreme dominance of highly frequent events. By integrating semantic event representations with frequency-aware behavioral signals, the proposed approach generates compact session embeddings that capture both interaction semantics and behavioral intensity without requiring explicit sequence modeling.

A detailed experimental study used both a publicly available clickstream dataset and a large real-world e-commerce dataset. The experiments compared the proposed frequency-weighted aggregation method with traditional unweighted embedding aggregation across several transformer embedding models and classification algorithms. The results consistently showed that adding

frequency-based weighting significantly improves purchase prediction performance. In particular, the proposed method achieved notable gains in F1-score and overall accuracy across different embedding models and classifiers. These improvements were especially significant in the proprietary e-commerce dataset, which had a serious class imbalance. This suggests that frequency-aware session representations effectively capture behavioral signals related to purchase intent.

Overall, the findings suggest that combining lightweight transformer-based event embeddings with frequency-aware weighting offers an effective and efficient alternative to complex sequential recommendation models. The proposed framework enables scalable modeling of user navigation behavior while preserving important patterns found in clickstream interactions. As a result, this approach delivers practical benefits for large-scale e-commerce platforms that require efficient and clear user behavior modeling.

Despite these promising results, several directions remain for future research. First, the current framework focuses exclusively on frequency-based weighting to capture behavioral intensity signals. Future work may investigate hybrid weighting strategies that jointly incorporate frequency, temporal recency, and contextual information to better model evolving user intent within sessions. Second, the present study formulates purchase prediction as a binary next-event classification task. Extending the framework to multi-class next-event prediction problems may enable modeling a broader range of user interaction outcomes, including product exploration, cart additions, and recommendation engagement.

Another promising direction is to enhance the quality of event representations through domain-adaptive embedding techniques. While this study uses pretrained lightweight transformer models, future work could investigate domain-specific fine-tuning strategies that adjust embedding models to the features of clickstream interaction data. Such adjustments could further improve the semantic representation of user navigation events and boost downstream predictive performance.

Acknowledgements. The authors would like to thank the Trendyol Research and Development Center for providing the industrial collaboration and data infrastructure that supported this study. This research paper used generative AI tools, including Google Gemini and ChatGPT, solely for improving language clarity, grammar, and overall readability of the manuscript. These tools were not used to generate scientific ideas, experimental designs, analyses, or research results. The authors retain full responsibility for the accuracy, originality, and integrity of the work.

References

1. Moe, W.W., Fader, P.S.: Dynamic conversion behavior at e-commerce sites. Manage. Sci. **50**(3), 326–335 (2004)
2. Sakar, C.O., Polat, S.O., Katircioglu, M., Kastro, Y.: Real-time prediction of online shoppers' purchasing intention using multilayer perceptron and lstm recurrent neural networks. Neural Comput. Appl. **31**, 6893–6908 (2019)

3. Guo, H., Tang, R., Ye, Y., Li, Z., He, X.: Deepfm: a factorization-machine based neural network for ctr prediction, arXiv preprint arXiv:1703.04247 (2017)
4. Zhang, S., Yao, L., Sun, A., Tay, Y.: Deep learning based recommender system: a survey and new perspectives. ACM Comput. Surv. **52**(1), 1–38 (2019)
5. Sun, F., et al.: Bert4rec: sequential recommendation with bidirectional encoder representations from transformer. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, pp. 1441–1450 (Nov 2019)
6. Jannach, D., Ludewig, M., Lerche, L.: Session-based item recommendation in e-commerce: on short-term intents, reminders, trends and discounts. User Model. User-Adap. Inter. **27**(3), 351–392 (2017)
7. Quadrana, M., Cremonesi, P., Jannach, D.: Sequence-aware recommender systems. ACM Comput. Surv. **51**(4), 1–36 (2018)
8. Taşgetiren, N., Kılıç, T., Aktas, M.S.: On the use of graph embedding techniques for clustering user browsing navigational behaviours. Int. J. Web Grid Serv. **20**(4), 482–504 (2024)
9. Buyuktanir, T., Aktas, M.S.: On the use of embedding techniques for modeling user navigational behavior in intelligent prefetching strategies. Concurrency Comput. Pract. Exper. **37**(3), e8356 (2025)
10. Tasgetiren, N., Safak, I., Aktas, M.S.: An event frequency-inverse session frequency based weighting mechanism for data embedding methodologies. Softw. Pract. Exper. **55**(6), 983–1002 (2025)
11. Taşgetiren, N., cSafak, I., Aktaş, M.S.: Aging-based weighting for session classification in user behavior analysis, vol. 15886, pp. 16–34 (Jun. 2025)
12. Reimers, N., Gurevych, I.: Sentence-BERT: sentence embeddings using siamese BERT-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 3982–3992 (Nov. 2019)
13. Wang, L., et al.: Text embeddings by weakly-supervised contrastive pre-training, arXiv preprint arXiv:2212.03533 (Dec. 2022)
14. S. Xiao, Z. Liu, P. Zhang, N. Muennighoff, D. Lian, and J.-Y. Nie, "C-pack: Packed resources for general chinese embeddings," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Jul. 2024, pp. 641–649
15. Breiman, L.: Random forests. Mach. Learn. **45**(1), 5–32 (2001)
16. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794
17. Sokolova, M., Lapalme, G.: A systematic analysis of performance measures for classification tasks. Information Processing & Management **45**(4), 427–437 (2009)
18. Hidasi, B., Karatzoglou, A., Baltrunas, L., Tikk, D.: Session-based recommendations with recurrent neural networks, arXiv preprint arXiv:1511.06939 (Nov. 2015)
19. Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., Ma, J.: Neural attentive session-based recommendation. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, pp. 1419–1428 (Nov 2017)
20. Quadrana, M., Karatzoglou, A., Hidasi, B., Cremonesi, P.: Personalizing session-based recommendations with hierarchical recurrent neural networks. In: Proceedings of the Eleventh ACM Conference on Recommender Systems, pp. 130–137 (Aug. 2017)
21. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space, arXiv preprint arXiv:1301.3781 (Jan. 2013)

22. Greenstein-Messica, A., Rokach, L., Friedman, M.: Session-based recommendations using item embedding. In: Proceedings of the 22nd International Conference on Intelligent User Interfaces. Association for Computing Machinery, pp. 629–633 (2017)
23. Perozzi, B., Al-Rfou, R., Skiena, S.: DeepWalk: online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 701–710 (2014)
24. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 855–864 (Aug. 2016)
25. Kang, W.-C., McAuley, J.: Self-attentive sequential recommendation. In: 2018 IEEE International Conference on Data Mining (ICDM), pp. 197–206. IEEE (Nov. 2018)
26. Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., Tan, T.: Session-based recommendation with graph neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33(01), pp. 346–353 (2019)
27. de Souza Pereira Moreira, G., Rabhi, S., Lee, J.M., Ak, R., Oldridge, E.: Transformers4Rec: Bridging the gap between NLP and sequential/session-based recommendation. In: Proceedings of the 15th ACM Conference on Recommender Systems, pp. 143–153 (Sep. 2021)
28. Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., Zhou, M.: MiniLM: deep self-attention distillation for task-agnostic compression of pre-trained transformers. Adv. Neural. Inf. Process. Syst. **33**, 5776–5788 (2020)

A Modified Bias-Corrected Estimator for the Log-Logistic Shape Parameter

Ayana Mateus[1,2(✉)] and Frederico Caeiro[1,2]

[1] Center for Mathematics and Applications (NOVA Math), Campus de Caparica, Portugal
[2] NOVA School of Science and Technology (NOVA FCT), Campus de Caparica, Portugal
{amf,fac}@fct.unl.pt

Abstract. In this article, we propose a modified bias reduced estimator of the shape parameter of the three-parameter log-logistic distribution. The new estimator is based on the tail subsample and allows for the independent estimation of only the shape parameter. Consequently, the new estimator can be used in conjunction with the best linear unbiased estimators for the location and scale parameters. Further, we discuss the choice of the threshold that defines the tail subsample. The theoretical and simulation results demonstrate that the proposed estimator is nearly unbiased and clearly outperforms other approaches, particularly in small and moderate sample sizes. Finally, the performance of the new estimator is illustrated with an application to a real dataset.

Keywords: Bias correction · Log-logistic model · Shape-parameter

1 Introduction

The log-logistic distribution is a flexible heavy tailed model, widely applied in fields such as survival analysis, hydrology, insurance, economics, and reliability. The variable X follows a three parameter log-logistic distribution, with shape parameter $\alpha > 0$, location parameter μ, and scale parameter $\sigma > 0$, if its distribution function (DF) can be expressed as follows

$$F(x) = 1 - \frac{1}{1 + \left(\frac{x-\mu}{\sigma}\right)^{\alpha}} = \frac{1}{1 + \left(\frac{x-\mu}{\sigma}\right)^{-\alpha}}, \quad x > \mu. \tag{1}$$

This model is a particular member of the four-parameter kappa family of distributions [22]. The location parameter defines the lower bound of the support of the distribution. When $\mu = 0$ one obtains the two parameter log-logistic model, which is also known as the Fisk [14] or the Pareto-type III distribution [3]. The scale parameter corresponds to the median of $X - \mu$ and the parameter α controls the tail heaviness.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 260–271, 2027.
https://doi.org/10.1007/978-3-032-30518-3_17

The r-th moment about the location parameter, μ, is:

$$E[(X-\mu)^r] = \sigma^r B\left(1-\frac{r}{\alpha}, 1+\frac{r}{\alpha}\right), \quad r<\alpha,$$

where B represents the complete Beta function. This moment exists only if $r < \alpha$. Hence, the shape parameter determines the existence of moments: the mean and the variance only exist if $\alpha > 1$ and $\alpha > 2$, respectively. This moment constraint has important implications for parameter estimation methods based on moments.

Several methods for estimating the parameters of the three-parameter log-logistic model have already been proposed by different researchers (see [1,8,27,28]). Balakrishnan et al. [5] obtained the best linear unbiased estimators (BLUE) for the location and scale parameters in (1), assuming a known shape parameter. Regarding the estimation of the parameter α, Ahsanullah and Alzaatreh [2] observed that the log-logistic model belongs to the class of Pareto-type distributions and derived an estimator, based on the largest order statistics. The selection of the number of largest order statistics is not trivial, as it involves the traditional bias-variance trade-off. More recently, Mateus and Caeiro [23] and Caeiro and Mateus [13] improved the aforementioned shape parameter estimator by proposing alternative bias-reduced estimators. The present work focuses on further reducing the bias in the estimation of the shape parameter of the three-parameter log-logistic distribution.

This paper is structured as follows: Sect. 2 provides a review of several estimators for the parameter α, introduces the proposed bias-reduced estimator and discusses the choice of the sample fraction used in inference. Section 3 presents a Monte Carlo simulation study to evaluate the finite sample performance of the proposed estimator. An application to a real dataset is given in Sect. 4. Finally, Sect. 5 provides a summary and perspectives of future research.

2 Estimation of the Shape Parameter

In the following, let $(X_1, X_2, \ldots, X_n)$ denote a sample of n independent and identically distributed random variables with a common three-parameter log-logistic, with DF in (1). Furthermore, let $(X_{1:n} \le X_{2:n} \le \cdots \le X_{n:n})$ denote the associated sample of order statistics.

2.1 A Review of Several Estimators from the Literature

Ahsanullah and Alzaatreh [2] observed that the log-logistic model exhibits a Pareto-type tail, with a tail index equal to α. More precisely, the survival function $1 - F(x)$ is characterized by:

$$1 - F(x) = x^{-\alpha} l(x), \tag{2}$$

where $\alpha > 0$ and $l(x)$ is slowly-varying, as $x \to \infty$. The parameter $\xi = \alpha^{-1}$, known as the extreme value index, is the primary focus in tail inference. A wide

range of estimators for ξ have been developed in the literature [6,7,18,24]. One of the most widely used estimator for a positive extreme value index was introduced by Hill [21]. This estimator corresponds to the mean of the log-excesses, above the threshold $X_{n-k:n} > 0$,

$$H(k) = \frac{1}{k} \sum_{i=1}^{k} \ln X_{n-i+1:n} - \ln X_{n-k:n}, \qquad 1 \leq k \leq n-1, \tag{3}$$

where k represents the number of the largest order statistics (exceedances) used from the right tail of the distribution.

This estimator is only invariant under scale transformations. In the case of data from a model with a location parameter, such as the three-parameter log-logistic distribution, different locations lead to different bias patterns (see [4,16,20] for further details). This problem led Ahsanullah and Alzaatreh [2] to use a particular case of the location-invariant Hill-type estimator introduced in [15],

$$\hat{H}(k) = \frac{1}{k} \sum_{i=1}^{k} \ln \frac{X_{n-i+1:n} - X_{1:n}}{X_{n-k:n} - X_{1:n}}, \quad 1 \leq k \leq n-2. \tag{4}$$

Thus, this approach yielded the following scale- and location-invariant estimator for the shape parameter:

$$\hat{\alpha}^{LIH}(k) = \frac{1}{\hat{H}(k)}. \tag{5}$$

The choice of the threshold k is critical and involves the well known trade-off between the bias and the variance of the estimator ([9,11,17,26]). Ahsanullah and Alzaatreh [2] suggested $k = \lfloor n/10 \rfloor$, if $n > 100$, where $k = \lfloor x \rfloor$ designate the floor function. It is known that, to minimize the mean square error (MSE), k should be intermediate, that is, the threshold should satisfies the following conditions:

$$k = k_n \to \infty \quad \text{and} \quad \frac{k}{n} \to 0, \quad \text{as} \qquad n \to \infty. \tag{6}$$

Mateus and Caeiro [23] obtained the distributional behaviour

$$\hat{\alpha}^{LIH}(k) \overset{d}{=} \alpha \left(1 - \frac{Z_k}{\sqrt{k}} + O_p\left(\frac{\sqrt{k}}{n}\right) - \frac{k}{2n}(1 + o_p(1)) \right), \tag{7}$$

where Z_k represents an asymptotic standard normal random variable, and found that the optimal threshold, k_0^{LIH}, which minimizes the asymptotic MSE (AMSE) of $\hat{\alpha}^{LIH}(k)$, is

$$k_0^{LIH} = (2n^2)^{1/3}.$$

Since the location-invariant Hill-type, in (4), is strongly affected by bias, other alternative bias reduced estimators appeared in the literature. Mateus and Caeiro [23] adapted the generalized Jackknife (GJ) estimator for the extreme

value index, introduced in [19], thereby deriving the following location-invariant GJ estimator for the shape parameter α:

$$\hat{\alpha}^{GJ}(k) = \frac{\hat{H}(k)}{\hat{M}(k) - (\hat{H}(k))^2}, \quad 1 \leq k \leq n-2, \tag{8}$$

with

$$\hat{M}(k) = \frac{1}{k}\sum_{i=1}^{k}\left(\ln \frac{X_{n-i+1:n} - X_{1:n}}{X_{n-k:n} - X_{1:n}}\right)^2. \tag{9}$$

While this estimator has a null dominant term of asymptotic bias, it exhibits a larger asymptotic variance than the location-invariant Hill-type estimator $\hat{\alpha}^{LIH}(k)$, in (4). More recently, Caeiro and Mateus [13] eliminated the main term of asymptotic bias of $\hat{H}(k)$, in (4), without increasing its asymptotic variance. They derived the following bias-reduced (BR) estimator for the shape parameter α:

$$\hat{\alpha}^{BR}(k) = \frac{1}{\hat{H}^{BR}(k)}, \quad 1 \leq k \leq n-2, \tag{10}$$

where

$$\hat{H}^{BR}(k) = \hat{H}(k) \times \left(1 - \frac{k}{2n}\right), \quad 1 \leq k \leq n-2. \tag{11}$$

This estimator is obtained by estimating the main element of asymptotic bias of $\hat{H}(k)$, which is $\alpha^{-1}\frac{k}{2n}$, by $\hat{H}(k)\frac{k}{2n}$, and subsequently subtracting it from the original estimator, that is,

$$\hat{H}^{BR}(k) = \hat{H}(k) - \hat{H}(k)\frac{k}{2n}.$$

Since only the dominant bias term is removed, the $\hat{H}^{BR}(k)$ estimator tends to exhibit a positive bias for large values of k.

2.2 An Improved Bias-Reduced Estimator

To further decrease the effect of bias, we now propose an improved bias-reduced (IBR) estimator for the parameter α. The estimator is constructed by using bias correction techniques similar to the one used in [10,12]. More precisely, since $e^x - 1 = x + o(x)$, as $x \to 0$ and $e^x - 1 > x$, for $x > 0$, the main bias term of $\hat{H}(k)$ is estimated by $\hat{H}(k)(\exp\left(\frac{k}{2n}\right) - 1)$ and subsequently subtracted from the $\hat{H}(k)$ estimator, leading to the improved bias reduced (IBR) estimator:

$$\hat{\alpha}^{IBR}(k) = \frac{1}{\hat{H}^{IBR}(k)}, \quad 1 \leq k \leq n-2 \tag{12}$$

where

$$\hat{H}^{IBR}(k) = \hat{H}(k) \times \left[2 - \exp\left(\frac{k}{2n}\right)\right], \quad 1 \leq k \leq n-2. \tag{13}$$

For a detailed analysis of the asymptotic properties of the aforementioned estimators, see Caeiro and Mateus [13]. Furthermore, the asymptotic normality of (13) follows from the results established in the same paper.

2.3 Threshold Selection

For small and moderate sample sizes, the determination of the optimal threshold should not be made based on asymptotic results. Here we apply the algorithm proposed in [23]:

Algorithm [Empirical Threshold Selection]:

1. Generate 5000 samples of size n from a log-logistic distribution, with n taking values between 50 and 2000, with step 50.
2. Let $\hat{\alpha}(k,i,n)$ denote the estimates based on the i-th sample of size n. For each sample size,
   - compute $\hat{\alpha}(k,i,n)$, $k = 1,2,\ldots,n-2$, $i = 1,\ldots,5000$.
   - compute the empirical Mean Squared Error, as a function of k.
   - Determine the value $\hat{k}_0(n)$ that minimizes the empirical Mean Squared Error.
3. Finally, perform a Power Regression with $\hat{k}_0(n)$ the response variable and the sample size as the predictor variable. The regression coefficients are the vector values (a_1, a_2). The empirical threshold is then given by

$$\hat{k}_E = \hat{k}_E(n) = \lfloor a_1 n^{a_2} \rfloor, \tag{14}$$

where $\lfloor x \rfloor$ is the floor function (the integer part of x).

Remark: The agreement between the empirical thresholds and the power regression line is dependent on the parameters (sample size, step and replicates), defined in Step 1.

After applying this Algorithm to samples of the log-logistic model, with parameters $(\alpha,\mu,\sigma) = (1,1,1)$, the following empirical thresholds, for the estimators $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$ and $\hat{\alpha}^{IBR}(k)$, we obtained:

$$k_E^{GJ} = \lfloor 1.11 n^{0.92} \rfloor, \tag{15}$$

$$\hat{k}_E^{BR} = \lfloor 1.09 n^{0.86} \rfloor, \tag{16}$$

and

$$\hat{k}_E^{IBR} = \lfloor 0.79 n^{0.99} \rfloor. \tag{17}$$

Figure 1 displays the empirical optimal sample fractions and their corresponding fitted regression curves, for the reduced bias estimators of α. Clearly, there is a good alignment between the empirical data and the fitted curves.

3 Monte Carlo Simulation Study

In this section, we employ the Monte Carlo simulation method to analyse the finite-sample performance of the proposed bias-reduced estimator for the shape parameter α, in (12), against the estimators defined in (5), (8) and (10). The simulation study is described as follows:

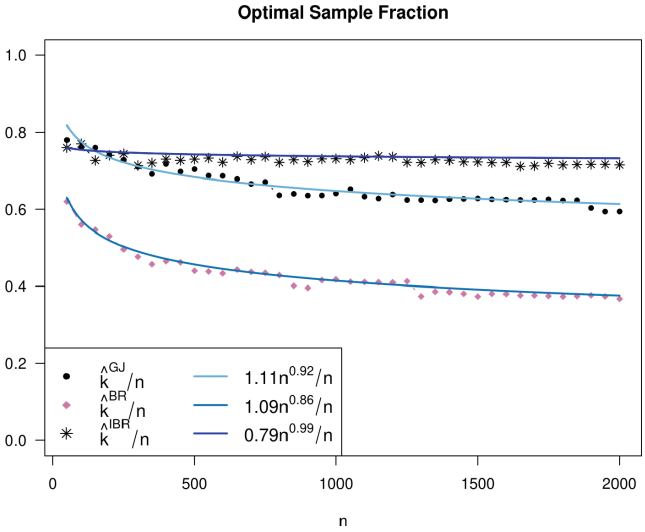

Fig. 1. Empirical optimal sample fraction for $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$ and $\hat{\alpha}^{IBR}(k)$.

1. For each of the parameter combinations listed below, we generated 5000 samples from the three-parameter log-logistic model using the inverse transformation technique. The quantile function of the log-logistic model is specified by
$$Q(u) = \mu + \sigma \left(\frac{u}{1-u} \right)^{1/\alpha}, \tag{18}$$
with $0 < u < 1$.
2. Let $\hat{\alpha}^{M}(k)$, with $M \in \{LIH, GJ, BR, IBR\}$, generically represent any of the aforementioned estimators under study, for the parameter α. For each generated sample, the estimates $\hat{\alpha}_i^{M}(k)$, $i = 1, 2, \ldots, 5000$, were calculated for every k. We assess the estimators accuracy and efficiency using three standard statistical measures: Bias, Standard Deviation (SD), and Root Mean Squared Error (RMSE). These metrics are defined as follows:
$$Bias(\hat{\alpha}^{M}(k)) = \bar{\alpha}^{M}(k) - \alpha, \tag{19}$$
$$SD(\hat{\alpha}^{M}(k)) = \frac{1}{5000} \sum_{i=1}^{5000} (\hat{\alpha}_i^{M}(k) - \bar{\alpha}^{M}(k))^2 \tag{20}$$
$$RMSE(\hat{\alpha}^{M}(k)) = \sqrt{\frac{1}{5000} \sum_{i=1}^{5000} (\hat{\alpha}_i^{M}(k) - \alpha)^2}, \tag{21}$$
with
$$\bar{\alpha}^{M}(k) = \frac{1}{5000} \sum_{i=1}^{5000} \hat{\alpha}_i^{M}(k), \tag{22}$$

and $1 \leq k \leq n-2$.

3. We also determined the simulated metrics at their optimal level. For simplicity of notation, hereafter we shall denote

$$\hat{\alpha}_0^{LIH} = \hat{\alpha}^H(k_0^{LIH}),$$

$$\hat{\alpha}_E^{GJ} = \hat{\alpha}^{GJ}(\hat{k}_E^{GJ}),$$

$$\hat{\alpha}_E^{RB} = \hat{\alpha}^{RB}(\hat{k}_E^{RB}),$$

$$\hat{\alpha}_E^{IRB} = \hat{\alpha}^{IRB}(\hat{k}_E^{IRB}).$$

4. We used samples of size n, with $n \in \{25, 50, 100, 250, 500, 1000, 2500, 5000\}$, across the following four scenarios:
   - Scenario I: $(\alpha, \mu, \sigma) = (0.5, 1.5, 1)$;
   - Scenario II: $(\alpha, \mu, \sigma) = (0.75, 1.5, 1)$;
   - Scenario III: $(\alpha, \mu, \sigma) = (1.5, 1.5, 1)$;
5. All numerical computations and statistical evaluations were conducted in R programming environment [25].

The simulated metrics for all values of k are presented in Figs. 2, 3, 4, and 5 for $n = 500$. Similar patterns were observed for the other sample sizes considered in this study. Since the simulated bias is nearly identical to the Mbias across all parameter combinations, we report only the simulated bias. This bias is generally positive for $\hat{\alpha}^{GJ}(k)$ and non-positive for the remaining estimators. Notably, $\hat{\alpha}^{IBR}(k)$ yields a near-zero absolute bias over a wide range of k-values. Regarding the standard deviation (SD), for a fixed k, we observe:

$$SD(\hat{\alpha}^{LIH}(k)) \leq SD(\hat{\alpha}^{BR}(k)) \leq SD(\hat{\alpha}^{IBR}(k)) \leq SD(\hat{\alpha}^{GJ}(k)).$$

If different values of k are allowed for each estimator, these rankings may vary. Furthermore, Fig. 5 shows that the proposed estimator, $\hat{\alpha}^{IBR}(k)$, can achieve the lowest RMSE. Table 1 presents the aforementioned metrics, computed at the optimal level for each estimator. For each metric and each sample size, the best results of the simulation are highlighted in bold. The results indicate that the proposed estimator $\hat{\alpha}^{IBR}(k)$ consistently outperforms the considered alternatives across all criteria.

4 Application to Real Data

In this section, we present an application using a real-world dataset to illustrate the behaviour of the estimators for the log-logistic shape parameter. This dataset, originally studied by Ahmad et al. [1], consists of the annual maximum flow (in m^3/s) for the River Spey at Kinrara from 1952 to 1982. The observations are as follows: 89.8, 109.1, 202.2, 146.3, 212.3, 116.7, 109.1, 80.7, 127.4, 138.8, 283.5, 85.6, 105.5, 118.0, 387.8, 80.7, 165.7, 111.6, 134.4, 131.5, 102.0, 104.3, 242.5, 214.8, 144.6, 114.2, 98.3, 102.8, 104.3, 196.2, 143.7.

Table 1. Simulated Bias, SD and RMSE, for the estimators at their optimal thresholds.

	25	50	100	250	500	1000	2500	5000
	$(\alpha, \mu, \sigma) = (0.5, 1.5, 1)$							
$Bias(\hat{\alpha}_0^{LIH})$	–0.0875	-0.0750	–0.0588	–0.0445	–0.0367	–0.0298	–0.0227	–0.0178
$Bias(\hat{\alpha}_E^{GJ})$	0.1685	0.0962	0.0584	0.0349	0.0239	0.0173	0.0122	0.0098
$Bias(\hat{\alpha}_E^{BR})$	–0.0370	–0.0270	–0.0234	–0.0165	–0.0138	–0.0108	–0.0082	-0.0067
$Bias(\hat{\alpha}_E^{IBR})$	**0.0126**	**0.0067**	**0.0010**	**–0.0016**	**–0.0020**	**–0.0018**	**–0.0016**	**–0.0012**
$SD(\hat{\alpha}_0^{LIH})$	0.1317	0.1015	0.0811	0.0625	0.0502	0.0406	0.0307	0.0248
$SD(\hat{\alpha}_E^{GJ})$	0.2424	0.1542	0.1057	0.0693	0.0507	0.0368	0.0245	0.0178
$SD(\hat{\alpha}_E^{BR})$	**0.1074**	0.0779	0.0566	0.0393	0.0294	0.0221	0.0149	0.0111
$SD(\hat{\alpha}_E^{IRB})$	0.1141	**0.0769**	**0.0538**	**0.0340**	**0.0243**	**0.0169**	**0.0107**	**0.0076**
$RMSE(\hat{\alpha}_0^{LIH})$	0.1580	0.1262	0.1002	0.0768	0.0622	0.0503	0.0382	0.0305
$RMSE(\hat{\alpha}_E^{GJ})$	0.2952	0.1818	0.1207	0.0776	0.0561	0.0407	0.0274	0.0203
$RMSE(\hat{\alpha}_E^{BR})$	**0.1136**	0.0825	0.0612	0.0426	0.0325	0.0246	0.0170	0.0129
$RMSE(\hat{\alpha}_E^{IBR})$	0.1148	**0.0772**	**0.0538**	**0.0341**	**0.0244**	**0.0170**	**0.0108**	**0.0077**
	$(\alpha, \mu, \sigma) = (0.75, 1.5, 1)$							
$Bias(\hat{\alpha}_0^{LIH})$	–0.1337	–0.1132	–0.0884	–0.0668	–0.0551	–0.0446	–0.0341	–0.0268
$Bias(\hat{\alpha}_E^{GJ})$	0.2640	0.1473	0.0884	0.0526	0.0359	0.0259	0.0183	0.0147
$Bias(\hat{\alpha}_E^{BR})$	-0.0646	–0.0434	–0.0360	–0.0249	–0.0207	–0.0162	–0.0123	–0.0100
$Bias(\hat{\alpha}_E^{IBR})$	**0.0040**	**0.0044**	**-0.0008**	**–0.0031**	**–0.0034**	**–0.0028**	**-0.0025**	**–0.0019**
$SD(\hat{\alpha}_0^{LIH})$	0.1965	0.1521	0.1216	0.0938	0.0752	0.0609	0.0461	0.0372
$SD(\hat{\alpha}_E^{GJ})$	0.3692	0.2319	0.1586	0.1040	0.0761	0.0553	0.0368	0.0267
$SD(\hat{\alpha}_E^{BR})$	**0.1592**	0.1164	0.0848	0.0589	0.0442	0.0332	0.0224	0.0166
$SD(\hat{\alpha}_E^{IBR})$	0.1684	**0.1147**	**0.0805**	**0.0510**	**0.0364**	**0.0254**	**0.0160**	**0.0114**
$RMSE(\hat{\alpha}_0^{LIH})$	0.2376	0.1896	0.1503	0.1152	0.0932	0.0755	0.0573	0.0458
$RMSE(\hat{\alpha}_E^{GJ})$	0.4539	0.2747	0.1815	0.1166	0.0841	0.0610	0.0411	0.0304
$RMSE(\hat{\alpha}_E^{BR})$	0.1718	0.1242	0.0921	0.0640	0.0488	0.0369	0.0255	0.0194
$RMSE(\hat{\alpha}_E^{IBR})$	**0.1685**	**0.1148**	**0.0805**	**0.0511**	**0.0366**	**0.0255**	**0.0162**	**0.0116**
	$(\alpha, \mu, \sigma) = (1.5, 1.5, 1)$							
$Bias(\hat{\alpha}_0^{LIH})$	–0.3258	–0.2586	–0.1935	–0.1408	–0.1140	–0.0913	–0.0691	–0.0540
$Bias(\hat{\alpha}_E^{GJ})$	0.5120	0.2834	0.1712	0.1025	0.0702	0.0508	0.0362	0.0290
$Bias(\hat{\alpha}_E^{RB})$	-0.2454	–0.1532	–0.1096	-0.0675	–0.0517	-0.0383	–0.0274	–0.0216
$Bias(\hat{\alpha}_E^{IRB})$	**-0.1470**	**–0.0858**	**–0.0628**	**–0.0392**	**–0.0279**	**–0.0186**	**–0.0120**	**–0.0081**
$SD(\hat{\alpha}_0^{LIH})$	0.3728	0.2952	0.2398	0.1864	0.1499	0.1216	0.0921	0.0744
$SD(\hat{\alpha}_E^{GJ})$	0.7303	0.4534	0.3121	0.2062	0.1512	0.1101	0.0734	0.0533
$SD(\hat{\alpha}_E^{BR})$	**0.3017**	0.2259	0.1670	0.1171	0.0879	0.0662	0.0447	0.0331
$SD(\hat{\alpha}_E^{IBR})$	0.3201	**0.2249**	**0.1606**	**0.1026**	**0.0731**	**0.0511**	**0.0322**	**0.0230**
$RMSE(\hat{\alpha}_0^{LIH})$	0.4951	0.3924	0.3081	0.2336	0.1883	0.1521	0.1152	0.0919
$RMSE(\hat{\alpha}_E^{GJ})$	0.8918	0.5347	0.3559	0.2302	0.1667	0.1213	0.0818	0.0606
$RMSE(\hat{\alpha}_E^{BR})$	0.3889	0.2729	0.1997	0.1352	0.1020	0.0764	0.0524	0.0396
$RMSE(\hat{\alpha}_E^{IBR})$	**0.3522**	**0.2407**	**0.1724**	**0.1098**	**0.0783**	**0.0544**	**0.0343**	**0.0244**

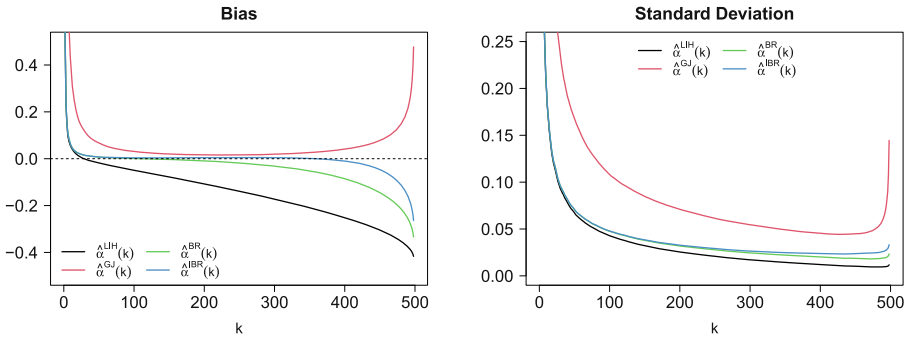

Fig. 2. Simulated bias and standard deviation of the estimators $\hat{\alpha}^{LIH}(k)$, $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$, and $\hat{\alpha}^{IBR}(k)$, as a function of k. Data generated from a log-logistic model with parameters $(\alpha, \mu, \sigma) = (0.5, 1.5, 1)$ and sample size $n = 500$.

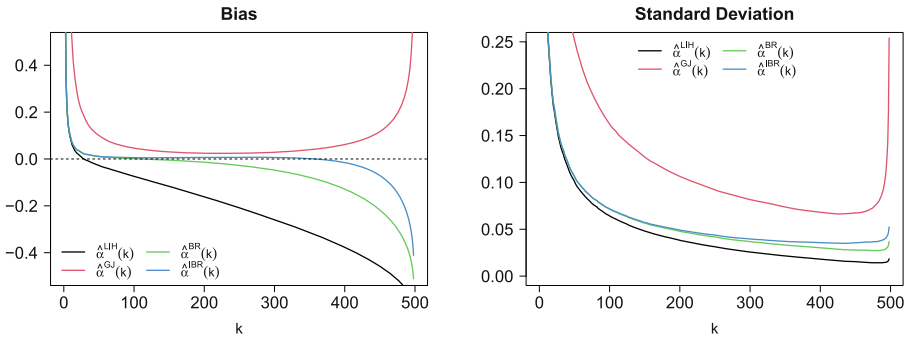

Fig. 3. Simulated bias and standard deviation of the estimators $\hat{\alpha}^{LIH}(k)$, $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$, and $\hat{\alpha}^{IBR}(k)$, as a function of k. Data generated from a log-logistic model with parameters $(\alpha, \mu, \sigma) = (0.75, 1.5, 1)$ and sample size $n = 500$.

Using the maximum likelihood (ML) method, Ahmad et al. [1] computed a shape parameter estimate of $\hat{c} = 0.531$. By the invariance property of ML estimators, the corresponding estimate for $\alpha = c^{-1}$ is $\hat{\alpha}^{ML} = 1.883$. The estimates for the shape parameter α provided by the alternative estimators studied in this paper are

$$\hat{\alpha}_0^{LIH} = 1.389,$$
$$\hat{\alpha}_E^{GJ} = 2.090,$$
$$\hat{\alpha}_E^{RB} = 1.599,$$
$$\hat{\alpha}_E^{IRB} = 1.902.$$

The results indicate that the reduced bias estimates, $\hat{\alpha}_E^{GJ}$, $\hat{\alpha}_E^{RB}$ and $\hat{\alpha}_E^{IRB}$ align more closely with the ML estimate. Moreover, the $\hat{\alpha}_E^{IRB}$ estimate of 1.902 is similar to the ML estimate.

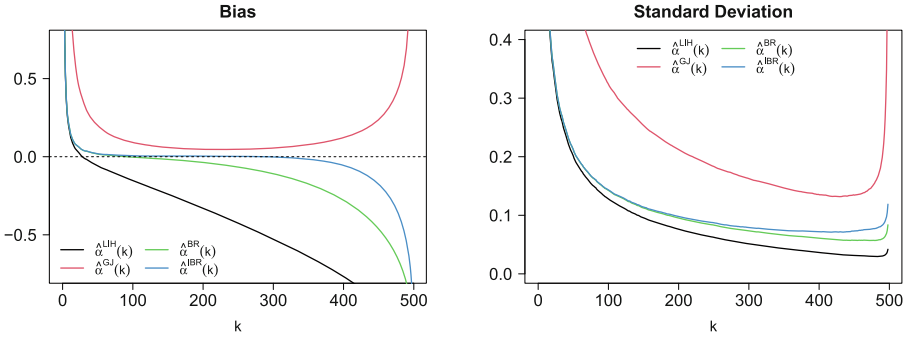

Fig. 4. Simulated bias and standard deviation of the estimators $\hat{\alpha}^{LIH}(k)$, $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$, and $\hat{\alpha}^{IBR}(k)$, as a function of k. Data generated from a log-logistic model with parameters $(\alpha, \mu, \sigma) = (1.5, 1.5, 1)$ and sample size $n = 500$.

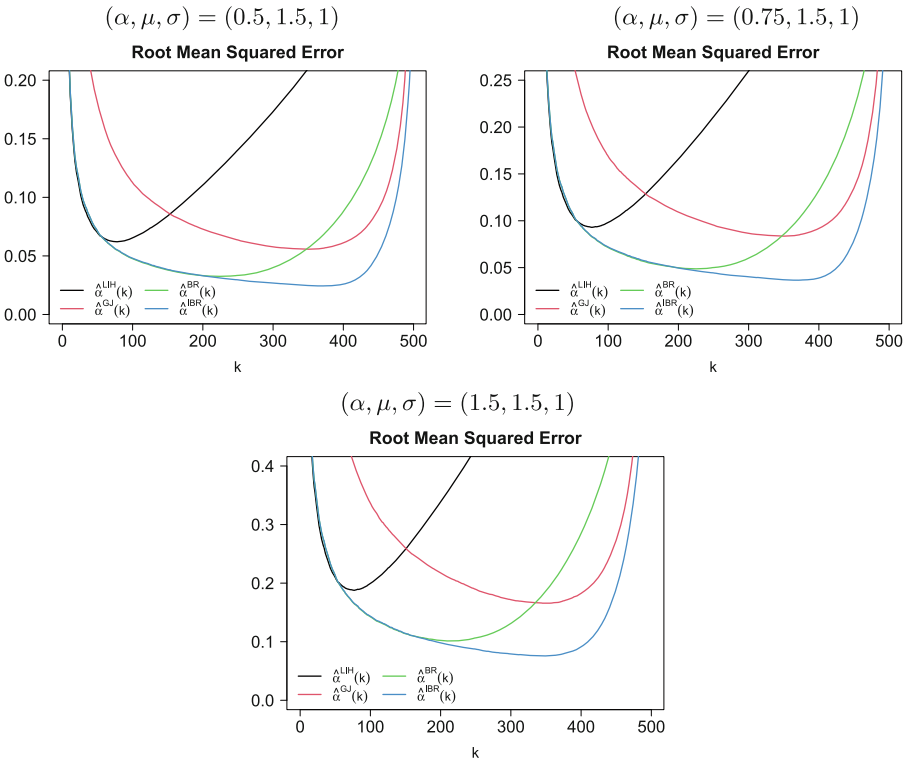

Fig. 5. Simulated RMSE of the estimators $\hat{\alpha}^{LIH}(k)$, $\hat{\alpha}^{GJ}(k)$, $\hat{\alpha}^{BR}(k)$, and $\hat{\alpha}^{IBR}(k)$, as a function of k.

5 Concluding Remarks

In this research, we introduced a modified bias reduced location-invariant estimator for the shape parameter of the three-parameter log-logistic model, which provides an improved bias correction. Our simulation results confirm that this new approach offers superior performance, consistently outperforming existing estimators, including other bias-reduced estimators, in terms of both bias reduction and RMSE. Furthermore, the conceptual framework presented here may be extended to correct higher order bias terms.

Acknowledgments. The authors thank the anonymous reviewers for their comments and suggestions. This study was partially funded by FCT Fundao para a Cincia e a Tecnologia, I.P., under the scope of the projects UID/00297/2025 (https://doi.org/10.54499/UID/00297/2025) and UID/PRR/00297/2025 (https://doi.org/10.54499/UID/PRR/00297/2025) (Center for Mathematics and Applications/NOVA Math).

Disclosure of Interests. The authors have no competing interests.

References

1. Ahmad, M., Sinclair, C., Werritty, A.: Log-logistic flood frequency analysis. J. Hydrol. **98**(3–4), 205–224 (1988)
2. Ahsanullah, M., Alzaatreh, A.: Parameter estimation for the log-logistic distribution based on order statistics. Revstat - Stat. J. **16**(4), 429–443 (2018)
3. Arnold, B.C.: Pareto distributions. Chapman and Hall/CRC, Monographs on Statistics and Applied Probability (2015)
4. Bachem, M., Ergun, L.M., de Vries, C.G.: Covariates hiding in the tails. The Econometrics Journal p. utaf029 (Dec 2025). https://doi.org/10.1093/ectj/utaf029
5. Balakrishnan, N., Malik, H., Puthenpura, S.: Best linear unbiased estimation of location and scale parameters of the log-logistic distribution. Commun. Stat. Theory Methods **16**(12), 3477–3495 (1987)
6. Beirlant, J., Herrmann, K., Teugels, J.: Estimation of the extreme value index. In: Extreme Events in Finance, pp. 97–115. John Wiley & Sons, Inc. (Oct 2016). https://doi.org/10.1002/9781118650318.ch5
7. Beirlant, J., Caeiro, F., Gomes, M.I.: An overview and open research topics in statistics of univariate extremes. REVSTAT – Statistical J. **10**(1), 1–31 (2012). https://doi.org/10.57805/revstat.v10i1.109, https://www.ine.pt/revstat/pdf/rs120101.pdf
8. Brilhante, M.d.F., Pestana, D., Pestana, P., Rocha, M.L.: Measuring the risk of vulnerabilities exploitation. AppliedMath **4**(1), 20–54 (2024). https://doi.org/10.3390/appliedmath4010002, https://www.mdpi.com/2673-9909/4/1/2
9. Caballero-Megido, C., Hillier, J., Wyncoll, D., Bosher, L., Gouldby, B.: Technical note: comparison of methods for threshold selection for extreme sea levels. J. Flood Risk Manag. **11**(2), 127–140 (2017)
10. Cabral, I., Caeiro, F., Gomes, M.I.: Reduced bias hill estimators. In: Simos, T., Kalogiratou, Z., Monovasilis, T. (eds.) AIP Conference Proceedings, vol. 1790, p. 080006 (2016). https://doi.org/10.1063/1.4968687

11. Caeiro, F., Gomes, M.I.: Threshold selection in extreme value analysis. In: Dey, D.K., Yan, J. (eds.) Extreme Value Modeling and Risk Analysis, pp. 69–86. Chapman and Hall/CRC (2016). https://doi.org/10.1201/b19721-5
12. Caeiro, F., Henriques-Rodrigues, L., Gomes, M.I., Cabral, I.: Minimum-variance reduced-bias estimation of the extreme value index: a theoretical and empirical study. Comput. Math. Methods **2**(4), e1101 (2020)
13. Caeiro, F., Mateus, A.: Reduced bias estimation of the shape parameter of the log-logistic distribution. J. Comput. Appl. Math. **436**, 115347 (2024)
14. Fisk, P.R.: The graduation of income distributions. Econometrica **29**(2), 171–185 (1961)
15. Fraga Alves, M.: A location invariant hill-type estimator. Extremes **4**(3), 199–217 (2001)
16. Gomes, M.I., Fraga Alves, M.I., Santos, P.A.: Port hill and moment estimators for heavy-tailed models. Commun. Stat. Simulat. Comput. **37**(7), 1281–1306 (2008)
17. Gomes, M.I., Guillou, A.: Extreme value theory and statistics of univariate extremes: a review. Int. Stat. Rev. **83**(2), 263–292 (2015). https://doi.org/10.1111/insr.12058
18. Gomes, M.I., Henriques-Rodrigues, L.: Comparison at optimal levels of classical tail index estimators: a challenge for reduced-bias estimation? Discussiones Mathematicae Probability Stat. **30**(1), 35–51 (2010). http://eudml.org/doc/277076
19. Gomes, M.I., Martins, M.J., Neves, M.: Alternatives to a semi-parametric estimator of parameters of rare events-the jackknife methodology. Extremes **3**(3), 207–229 (2000)
20. Gomes, M.I., Oliveira, O.: How can non-invariant statistics work in our benefit in the semi-parametric estimation of parameters of rare events. Commun. Stat. Simulat. Comput. **32**(4), 1005–1028 (2003)
21. Hill, B.M.: A simple general approach to inference about the tail of a distribution. Ann. Stat. **3**(5), 1163–1174 (1975). https://doi.org/10.1214/aos/1176343247
22. Hosking, J.R.M.: The four-parameter kappa distribution. IBM J. Res. Dev. **38**(3), 251–258 (1994)
23. Mateus, A., Caeiro, F.: Improved shape parameter estimation for the three-parameter log-logistic distribution. Comput. Math. Methods **2022**, 1–13 (2022)
24. Panaretos, J., Tsourti, Z.: Extreme value index estimators and smoothing alternatives: a critical review. In: Panaretos, J. (ed.) Stochastic Musings: perspectives from the pioneers of the late 20th Century, pp. 141–160. Laurence Erlbaum, New Jersey, USA (2003). https://doi.org/10.4324/9781410609120
25. R Core Team: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria (2025). https://www.R-project.org/
26. Scarrott, C., MacDonald, A.: A review of extreme value threshold estimation and uncertainty quantification. REVSTAT-Stat. J. **10**(1), 33–60 (2012). https://doi.org/10.57805/revstat.v10i1.110, https://revstat.ine.pt/index.php/REVSTAT/article/view/110
27. Shrahili, M., El-Saeed, A.R., Hassan, A.S., Elbatal, I., Elgarhy, M.: Estimation of entropy for log-logistic distribution under progressive type ii censoring. J. Nanomater. **2022**(1), 2739606 (2022)
28. Singh, V., Guo, H., Yu, F.: Parameter estimation for 3-parameter log-logistic distribution (lld3) by pome. Stoch. Hydrol. Hydraul. **7**(3), 163–177 (1993)



A Comparative NLP Study of Semantic and Thematic Changes in AI Research Pre - and Post-generative AI

Inês Sena(✉), João Mendes, and Ana I. Pereira

Research Centre in Digitalization and Intelligent Robotics (CeDRI), Laboratório para a Sustentabilidade e Tecnologia em Regiões de Montanha (SusTEC), Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
{ines.sena,joao.cmendes,apereira}@ipb.pt

Abstract. With the advancement of Artificial Intelligence (AI), particularly with the emergence of Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs), has significantly influenced the scope, accessibility, and dissemination of scientific research. This evolution raises important questions regarding how AI-related scientific discourse has changed over time. This study search lexical, semantic, and thematic changes in a dataset of approximately 50,000 AI-related publications retrieved from the Web of Science (WoS), comparing a pre-GenAI period (2015–2019) with a post-GenAI period (2022–2025). Natural Language Processing (NLP) techniques were applied, including TF-IDF for lexical analysis, Logistic Regression for temporal classification, and BERTopic for topic modeling. The results achieved an accuracy of 89%, revealing a clear separability between the two periods based on textual features. The analyses reveal a transition from predominantly technical and methodological terminology, focused on algorithms and infrastructure, toward more application-oriented and generative themes, including GenAI and LLMs. Furthermore, topic modeling shows the growing prominence of interdisciplinary domains such as education, healthcare, and business applications in the post-GenAI period. Overall, the findings suggest a significant lexical and thematic evolution in AI-related scientific discourse, reflecting broader transformations in research priorities and application domains following the emergence of GenAI technologies.

Keywords: Artificial Intelligence · GenAI · Natural Language Processing · Semantic Analysis · Topic Modeling

1 Introduction

Since 2022, a new Artificial Intelligence (AI) paradigm, namely Generative Artificial Intelligence (GenAI), experienced rapid growth with the emergence of tools such as ChatGPT and Gemini [6]. GenAI enables the generation of original content based on training over large-scale datasets containing text, images, audio, and video [10].

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 272–288, 2027.
https://doi.org/10.1007/978-3-032-30518-3_18

GenAI is primarily driven by Large Language Models (LLMs), whose emergence has considerably transformed scientific communication and research dissemination. These tools now support tasks such as language editing, manuscript structuring, literature synthesis, text generation, and scientific reasoning. While offering new opportunities, they also raise important concerns regarding authorship, accountability, originality, and transparency in the scientific production process [5,8].

Current discussion extend beyond the technical capabilities of these systems. A growing body of literature points out that GenAI challenges traditional notions of scientific authorship, historically associated with human intellectual contribution, responsibility, and accountability. In this context, the relationship between human authors and GenAI should not be interpreted as replacement or opposition, but rather as an emerging form of techno-social collaboration that demands new conceptual and ethical considerations for academic production. [10].

At the same time, several studies show that the growing integration of GenAI into academic workflows provides operational benefits, including improved efficiency, support for non-native English speakers, acceleration of literature reviews, and assistance in research planning and analysis. However, these advantages coexist with persisting concerns regarding plagiarism, superficial analysis, bias, hallucinations, fabricated references, and loss of authorship. This reinforces the need for human supervision and transparent disclosure mechanisms regarding the use of these tools [5,8].

Furthermore, recent studies suggests that GenAI is no longer limited to supporting micro-level writing tasks but is increasingly influencing broader knowledge-production processes, including automated literature synthesis and support for scientific discovery. Nevertheless, the quality of generated outputs remains highly dependent on factors such as input data quality, prompt design, retrieval mechanisms, and human critical evaluation, which currently prevents the replacement of human scientific judgment [7,9].

In this context, it becomes relevant to explore not only how researchers use GenAI, but also how its emergence coincides with lexical, semantic and thematic changes in scientific discourse. This study addresses this challenge through the application of Natural Language Processing (NLP) techniques, which provide effective tools for analyzing large-scale textual corpora and identifying lexical patterns, semantic differences, and topic evolution over time.

This study searchs how scientific literature related to AI has evolved following the emergence of GenAI, through the analysis of a large-scale dataset of 50,000 publications extracted from the Web of Science (WoS) for two distinct periods: a pre-GenAI phase (2015–2019) and a post-GenAI phase (2022–2025). Several NLP techniques were applied, including TF-IDF for lexical analysis, supervised classification using Logistic Regression to distinguish periods, and BERTopic for topic modeling, aiming to identify structural lexical, semantic, and thematic changes in scientific discourse.

The main objective of this work is to study whether AI research has changed from predominantly technical topics toward more applied and gener-

ative approaches. To achieve this objective, the study adresses the following research questions (RQ):

- RQ1: How has the lexical and semantic representation of AI research evolved between the pre-GenAI (2015–2019) and post-GenAI (2022–2025) periods?
- RQ2: To what extent can NLP-based classification techniques h distinguish AI research between these two periods?
- RQ3: What tematic changes can be identified in AI research topics across the two periods?

The remainder of this paper is structured as follows. Section 2 reviews work on GenAI in scientific writing and on the use of NLP approaches to analyze the scientific literature. Section 3 describes the dataset and the data collection process. Section 4 presents the procedure, including pre-processing techniques and NLP approaches applied. Section 5 presents the results, showing the lexical, semantic, and thematic changes. Section 6 answers the research questions, comparing the results with the state-of-the-art findings. Finally, Sect. 7 concludes the article and describes future research directions.

The remainder of this paper is structured as follows. Section 2 reviews previous work on GenAI in academic contexts and the application of NLP techniques for scientific literature analysis. Section 3 describes the dataset and the data collection process. Section 4 presents the methodology, including preprocessing procedures and NLP techniques applied in the study. Section 5 presents the obtained results, showing lexical, semantic, and thematic changes across periods. Section 6 discusses the findings in relation to the research questions and the literature. Finally, Sect. 7 concludes the paper and outlines future research directions.

2 Related Work

This section presents the current context regarding the impact of Generative Artificial Intelligence (GenAI) on scientific communication and academic production, showing the main challenges, limitations, and research gaps identified in the literature. It also reviews the application of Natural Language Processing (NLP) approaches for the analysis of lexical, semantic, and thematic evolution in scientific discourse, providing the theoretical foundation for the methodological approach adopted in this study.

2.1 GenAI and Scientific Communication

Scientific communication practices have been increasingly influenced by the emergence of GenAI tools, as researchers progressively integrate these systems into academic workflows for tasks such as language improvement, literature review, text generation, research planning, and idea structuring [8]. Studies show benefits including improved efficiency, reduced writing effort, increased productivity, and better accessibility for non-native English speakers through translation and language assistance capabilities [8].

Due to the capabilities of Large Language Models (LLMs), GenAI tools are also being adopted in more advanced stages of scientific production, including automated literature synthesis, semantic exploration of large corpus, and support for knowledge discovery processes [9].

Despite these advantages, the growing use of GenAI raises significant ethical and methodological concerns regarding transparency, plagiarism, accountability, and research integrity. The use of AI-generated content challenges traditional notions of human intellectual contribution and scientific responsibility [10]. Since AI systems cannot assume accountability for scientific claims or conclusions, several publishers and academic institutions have explicitly rejected the attribution of authorship to AI systems [5].

Furthermore, the quality and reliability of AI-generated outputs remain strongly dependent on the quality of training data, prompt design, and human supervision [7]. Previous studies report that generated content may present superficial reasoning, fabricated references, hallucinations, limited originality, and embedded biases, reinforcing the importance of critical human evaluation and validation [5].

2.2 NLP for Scientific Literature Analysis

NLP has evolved considerably over recent decades, transitioning from rule-based and statistical approaches to Machine Learning (ML), Deep Learning, and transformer-based architectures [2]. Early techniques such as Word2Vec and GloVe enabled the representation of words through dense semantic vectors capable of capturing lexical similarity, although without contextual awareness [3]. More advanced transformer-based models, including BERT and related architectures, introduced contextual embeddings that substantially improved language understanding across multiple NLP tasks [3].

The emergence of Transformer architectures significantly advanced NLP capabilities by enabling models to capture long-range contextual dependencies and semantic relationships in large textual corpus [3]. These architectures showed improved performance in tasks such as text classification, semantic analysis, information extraction, document clustering, and topic modeling [3].

NLP techniques have increasingly been applied to the analysis of scientific literature, allowing researchers to extract structured information from large-scale unstructured corpus. Applications include information extraction, bibliometric analysis, trend identification, semantic similarity analysis, and topic modeling [1]. Embedding-based topic modeling approaches have also enabled the identification of temporal thematic evolution across scientific domains [4].

Recent studies further reveal the growing role of NLP in supporting automated knowledge discovery from scientific literature, particularly in tasks involving literature review automation, semantic exploration, and trend detection [4]. These capabilities make NLP especially suitable for investigating how scientific discourse evolves in response to technological transformations such as the emergence of GenAI.

2.3 Research Gap

Although studies extensively discuss the applications, opportunities, and limitations of GenAI in academic contexts, comparatively less attention has been given to understanding how the emergence of GenAI is reflected in the evolution of scientific discourse itself. Most previous research focuses either on ethical implications or on the direct use of AI-assisted writing tools, rather than on broader lexical, semantic, and thematic transformations in scientific literature.

Additionally, current literature lacks large-scale empirical analyses combining NLP techniques with longitudinal scientific corpus to search how AI-related research discourse has evolved before and after the emergence of GenAI technologies. In particular, there remains limited evidence regarding whether recent changes observed in AI literature reflect a broader transition from technically oriented research toward more applied, interdisciplinary, and generative-oriented topics.To address this gap, the present study applies multiple NLP techniques to a large corpus of scientific publications extracted from the Web of Science (WoS), comparing pre-GenAI and post-GenAI periods to identify lexical, semantic, and thematic changes in AI-related scientific discourse.

3 Data Characterization

This study is based on a large-scale dataset of scientific publications obtained from the WoS database. The dataset focuses on research related to AI and was divided into two distinct temporal periods to enable a comparative analysis between pre- and post-GenAI research trends. Figure 1 presents the data collection and processing workflow adopted in this study.

The first period, referred to as the pre-GenAI phase, covers the years between 2015 and 2019 and represents the state of AI-related scientific research prior to the widespread adoption of GenAI technologies. The second period, referred to as the post-GenAI phase, covers the years between 2022 and 2025, reflecting the rapid expansion of GenAI and LLMs in scientific and technological contexts.

The years 2020 and 2021 were intentionally excluded from the analysis to create a clearer separation between the two phases and reduce overlap between emerging and consolidated GenAI-related discourse. These years correspond to a transitional period characterized by the increasing popularity of Transformer-based architectures and the emergence of early LLMs, such as chatGPT-3, which had not yet achieved broad integration into scientific communication practices.

For each period, 25,000 publications were extracted from WoS. Due to the platform limitation of exporting a maximum of 1,000 records per extraction, the data collection process was performed in multiple stages, generating several files that were subsequently merged into a unified dataset containing approximately 50,000 publications.

The WoS database provides metadata fields including title, abstract, keywords, authors, year of publication, source title, and document identifiers (e.g., DOI) for each publication. However, for the Natural Language Processing

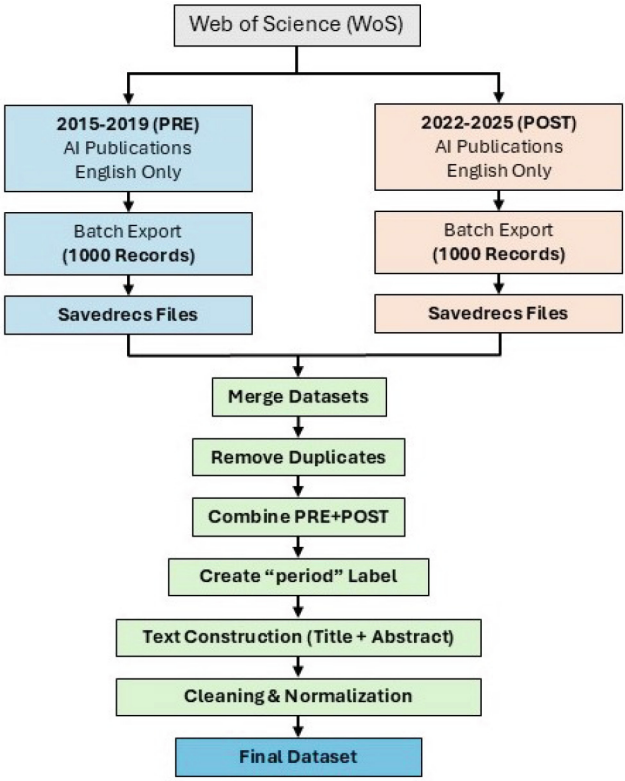

Fig. 1. Data collection and processing workflow.

(NLP) analysis conducted in this study, only textual fields, primarily titles and abstracts, were considered.

Following data collection, the dataset underwent a preprocessing stage involving duplicate removal, handling of missing values, and textual normalization. Additionally, a categorical variable indicating the corresponding temporal period (pre-GenAI or post-GenAI) was assigned to each document to support comparative analysis.

The final dataset contains structured metadata fields, including publication year, title, abstract, keywords, DOI, document type, source title, and source file information, as well as a processed textual field combining titles and abstracts. After preprocessing, the dataset contained 24,781 documents for the pre-GenAI period and 24,865 documents for the post-GenAI period.

The resulting corpus was subsequently prepared for NLP analysis and served as the basis for the lexical analysis, supervised classification, and topic modeling procedures described in the following sections.

4 Methodology

This study uses a multi-stage Natural Language Processing (NLP) pipeline to analyze lexical, semantic, and thematic differences between AI-related scientific publications from the pre-GenAI and post-GenAI periods. The methodology combines traditional statistical text analysis techniques with recent embedding-based approaches for document classification and topic modeling. Figure 2 presents the methodological workflow adopted in this study.

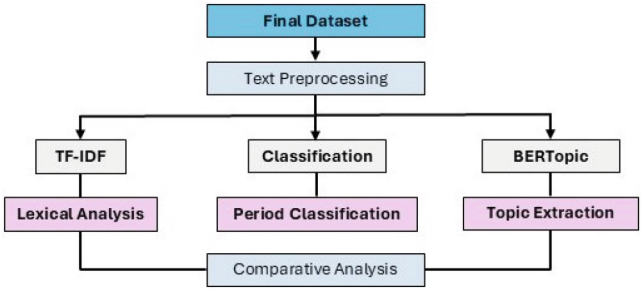

Fig. 2. Methodological pipeline used.

4.1 Text Preprocessing

The textual corpus was constructed by combining the title and abstract of each publication into a single textual field. Before analysis, several preprocessing procedures were applied to improve data consistency and reduce textual noise. These procedures included lowercasing, removal of punctuation and special characters, normalization of whitespace, and elimination of incomplete or duplicated records.

The preprocessing stage is particularly important in NLP tasks, as it reduces lexical variability and improves the quality of feature extraction and semantic representations [2].

4.2 Lexical Analysis Using TF-IDF

To identify the most representative lexical patterns within each temporal period, the Term Frequency–Inverse Document Frequency (TF-IDF) method was applied. TF-IDF is a statistical weighting approach that evaluates the relevance of terms based on their frequency within a document while reducing the importance of highly frequent words across the entire corpus.

TF-IDF is widely used in text mining and information retrieval due to its interpretability and effectiveness in identifying discriminative lexical features [2]. In this study, TF-IDF was employed to compare dominant terms between the pre-GenAI and post-GenAI corpus, enabling the identification of lexical modifications associated with changes in AI-related scientific discourse.

4.3 Supervised Classification

A supervised classification approach was used to evaluate whether scientific publications from the two temporal periods could be automatically distinguished based on their textual characteristics. Each document was labeled according to its corresponding period (pre-GenAI or post-GenAI), and textual features were extracted using TF-IDF vectorization.

The dataset was divided into training and testing subsets using a 70/30 split strategy, where 70% of the documents were used for training and 30% for evaluation. To reduce the risk of overfitting and improve generalization, the Logistic Regression classifier was selected due to its robustness, interpretability, and suitability for sparse and high-dimensional textual data [3]. In addition, regularization mechanisms inherent to Logistic Regression were maintained during training to penalize excessively complex decision boundaries and improve model stability.

The classification model was evaluated using unseen test data through metrics including accuracy, precision, recall, F1-score, and confusion matrix analysis. The confusion matrix was particularly important for identifying classification balance and potential misclassification patterns between the two periods.

Beyond predictive performance, the model coefficients were also analyzed to identify the most discriminative terms associated with each period. This procedure provided additional insights into lexical and semantic differences between pre-GenAI and post-GenAI AI-related scientific publications.

4.4 Topic Modeling with BERTopic

To analyse thematic evolution across the two periods, BERTopic was applied as the primary topic modeling technique. BERTopic combines Transformer-based embeddings with clustering algorithms to generate semantically coherent topics from large textual corpora.

Document embeddings were generated using a pre-trained SentenceTransformer model, which produces dense contextual vector representations capable of capturing semantic relationships beyond isolated lexical features [3]. Compared to traditional bag-of-words representations, contextual embeddings provide improved semantic understanding of complex scientific text.

BERTopic leverages these embeddings together with dimensionality reduction and density-based clustering techniques to identify groups of semantically related documents. This approach is particularly suitable for dynamic and multidisciplinary domains such as AI research, where terminology and thematic structures evolve rapidly over time [4].

The identified topics were subsequently analyzed and compared between periods to observe thematic transitions associated with the emergence and expansion of GenAI-related research.

4.5 Comparative Analysis

To analyze changes in terms, thematic focus, and research direction associated with the emergence of GenAI, results from the two periods were compared with respect to the extracted variables, classification outcomes, and topic distributions.

The combined methodological approach enables the integration of interpretable lexical analysis with more advanced embedding-based semantic representations, providing a comprehensive framework for studying the evolution of AI-related scientific discourse before and after the emergence of GenAI technologies.

5 Results

This section presents the results obtained from applying the proposed NLP pipeline to the AI-related scientific corpus. The findings are organized into three main analytical components: lexical analysis using TF-IDF, supervised classification, and topic modeling using BERTopic.

5.1 Lexical Analysis Using TF-IDF

To analyze lexical differences between the pre-GenAI (2015–2019) and post-GenAI (2022–2025) periods, TF-IDF was applied to identify the most representative terms in each corpus. Table 1 presents the top 15 terms with the highest average TF-IDF values for each period.

The results reveal clear lexical differences between the two periods. In the pre-GenAI period, the most representative terms are strongly associated with technical and methodological aspects of AI research, including neural networks, algorithms, machine learning, and model development. Terms such as *learning*, *neural*, *network*, and *algorithm* indicate a stronger focus on foundational AI methods and computational infrastructure.

In contrast, the post-GenAI period presents a broader and more application-oriented vocabulary. Terms such as *applications*, *analysis*, *review*, and *research* become more prominent, suggesting an expansion of AI research toward interdisciplinary and applied domains. In addition, GenAI-related terminology, including references to LLMs and generative systems, emerged more frequently during this period.

Overall, these findings suggest a lexical transition from technically centered AI research toward a more diversified scientific discourse characterized by broader application contexts and interdisciplinary integration.

5.2 Supervised Classification

To evaluate the separability between the two temporal periods, a supervised classification model based on Logistic Regression and TF-IDF features was trained

Table 1. The top 15 of the most representative terms in the pre-GenAI (2015–2019) and post-GenAI (2022–2025) periods.

Term PRE (2015–2019)	Mean TF-IDF	Term POST (2022–2025)	Mean TF-IDF
learning	0.032	ai	0.060
artificial	0.028	learning	0.024
data	0.027	data	0.023
intelligence	0.027	research	0.022
artificial intelligence	0.024	study	0.020
ai	0.022	review	0.017
model	0.0211	systems	0.017
based	0.0198	models	0.017
neural	0.0197	use	0.016
using	0.0196	intelligence ai	0.016
network	0.0194	analysis	0.016
algorithm	0.0194	applications	0.016
machine	0.0183	using	0.016
paper	0.0176	development	0.015
systems	0.0173	model	0.015

and tested. The dataset was divided using a 70/30 split, where 70% of the documents were used for training and 30% for testing. Table 2 presents the classification results obtained on the test dataset.

Table 2. Classification results for separability between the two periods.

Label	Accuracy	Recall	F1-score	Support
2015–2019 (0)	0.882	0.894	0.888	7392
2022–2025 (1)	0.894	0.881	0.887	7460

The results in Table 2 show an accuracy of 89%, revealing a strong distinction in the terms used between pre- and post-GenAI publications. Furthermore, they show a balance between classes, revealing consistent classification performance, as the accuracy and recall values presented are balanced between the two periods. To provide a more detailed analysis of classification performance. Figure 3 presents the confusion matrix obtained from the test dataset.

The confusion matrix shows that the majority of documents were correctly classified in their corresponding periods, with relatively low misclassification rates between classes. Specifically, most pre-GenAI documents were correctly

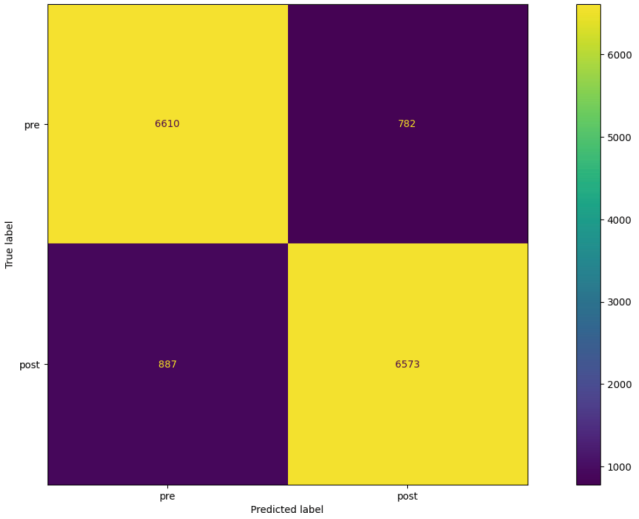

Fig. 3. Confusion matrix of Logistic Regression results to classify terms in pre-GenAI (2015–2019) and post-GenAI (2022–2025).

identified as belonging to the 2015–2019 period, while most post-GenAI publications were correctly assigned to the 2022–2025 period. The balanced distribution of false positives and false negatives suggests that the lexical differences between periods are sufficiently robust to support reliable automatic classification.

A more detailed examination of the classifier coefficients further revealed distinct discriminative vocabularies associated with each period. Publications from the pre-GenAI phase were more strongly associated with terms related to traditional AI methodologies, ML infrastructure, and algorithmic development. In contrast, publications from the post-GenAI phase contained a greater prevalence of terms associated with GenAI, emerging technologies, interdisciplinary applications, and broader societal integration of AI systems.

These findings reinforce the hypothesis that AI-related scientific discourse underwent substantial lexical and semantic changes following the emergence of GenAI technologies.

5.3 Topic Modeling with BERTopic

To analyze thematic evolution across periods, BERTopic was applied to the corpus to identify latent semantic topics based on documents embedding and clustering.

Figure 4 presents the most significant thematic changes between the pre-GenAI (2015–2019) and post-GenAI (2022–2025) periods.

The BERTopic analysis reveals a substantial increase in application-oriented and interdisciplinary topics during the post-GenAI period. Educational themes, including *education*, *students*, and *teaching*, showed some of the largest increases,

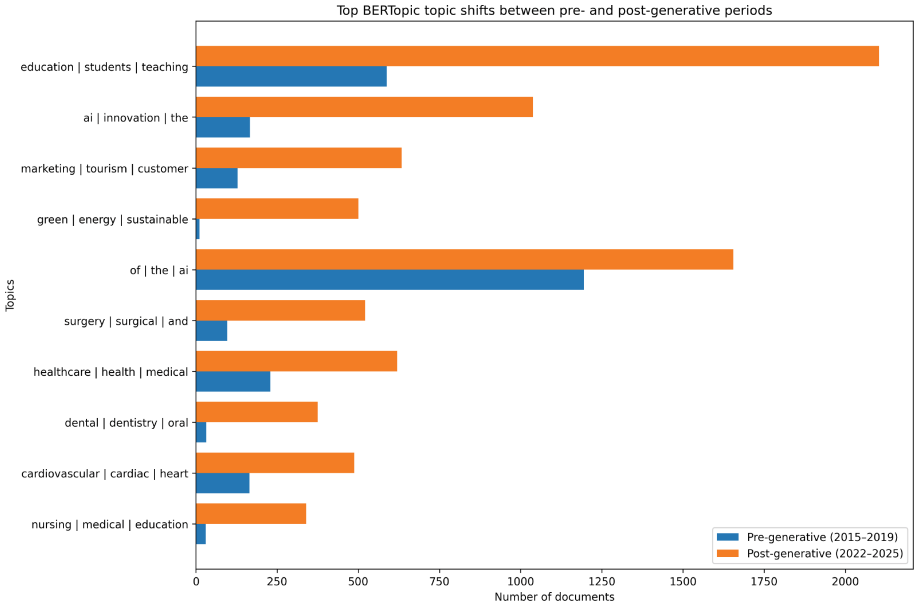

Fig. 4. Comparison of the most significant topic changes between pre-GenAI (2015–2019) and post-GenAI (2022–2025) periods.

suggesting a growing integration of AI technologies within educational environments.

Similarly, topics associated with business, innovation, and marketing expanded considerably, indicating a broader industrial and commercial adoption of AI systems. Medical and healthcare-related topics, including *health*, *surgery*, *dentistry*, and *cardiovascular* applications, also showed strong growth, reflecting the increasing role of AI in clinical and healthcare research.

Environmental and sustainability-related themes, such as *green energy* and *sustainability*, became more prominent in the post-GenAI period, further emphasizing the diversification of AI applications into new scientific and societal domains.

Although foundational AI topics remained present across both periods, their relative prominence was lower compared to the rapid growth of application-oriented themes. Overall, the BERTopic results suggest a thematic transition from predominantly technical and methodological research toward a broader interdisciplinary and application-focused research landscape.

These findings are consistent with the lexical and classification analyses, collectively supporting the substantial structural changes in AI-related scientific discourse following the emergence and expansion of GenAI technologies.

6 Discussion

The results of this study provide consistent evidence of a transformation in AI-related scientific discourse, particularly following the emergence and growth of GenAI. By combining lexical analysis, supervised classification, and topic modeling, it was possible to identify both surface-level linguistic changes and broader structural transformations in research themes and scientific communication patterns across the two analyzed periods.

The outcomes suggest that the evolution of AI technologies is reflected not only in the topics being study but also in the vocabulary, semantic structures, and research orientation of scientific publications. In addition, the results reinforce the idea that the rise of GenAI coincides with a transition from predominantly methodological and infrastructure-oriented research toward broader interdisciplinary and application-driven scientific production.

6.1 RQ1 How Has the Lexical and Semantic Representation of AI Research Evolved Between the Pre-GenAI (2015–2019) and Post-GenAI (2022–2025) Periods?

The TF-IDF analysis revealed substantial lexical differences between the two periods. In the pre-GenAI phase, the dominant vocabulary was strongly associated with technical and methodological aspects of AI research, including terms related to algorithms, neural networks, ML models, and computational systems. This pattern reflects a research landscape primarily focused on the development and optimization of AI techniques and infrastructure.

In contrast, the post-GenAI period presented a broader and more application-oriented vocabulary. Terms associated with healthcare, education, management, analysis, and applications became considerably more prominent, suggesting an expansion of AI research into interdisciplinary and real-world domains. Furthermore, the emergence of terms directly associated with GenAI and LLMs, including references to generative systems and conversational AI tools such as ChatGPT, show the growing influence of these technologies on scientific discourse.

These outcomes are consistent with previous studies describing the expansion of NLP and LLM applications across multiple scientific domains [3,4]. However, while most previous literature focuses on technological capabilities and ethical implications, this study contributes empirical evidence showing how these technological developments are reflected in the language used in scientific publications.

Nevertheless, the results also reveal an important distinction between thematic evolution and stylistic transformation. Although the study identifies significant lexical and semantic changes, these changes are mainly associated with changes in research topics and application domains rather than with deeper stylistic characteristics of scientific writing itself. Since the analysis was limited to titles and abstracts, additional stylistic modifications present in full-text sec-

tions such as methodology, discussion, and conclusions may not have been fully captured.

6.2 RQ2: To What Extent Can NLP-Based Classification Techniques H Distinguish AI Research Between These Two Periods?

The supervised classification results demonstrate that publications from the two periods can be reliably distinguished based solely on textual features. The Logistic Regression classifier achieved approximately 89% accuracy, showing that the lexical and semantic differences between periods are sufficiently strong to support automatic temporal classification.

The confusion matrix further confirmed the robustness of the model, showing balanced classification performance across both classes and relatively low misclassification rates. The similarity between precision, recall, and F1-score values suggests that the classifier generalized consistently and that overfitting effects were minimized through the use of TF-IDF vectorization, regularized Logistic Regression, and the separation of independent training and testing datasets.

These results reinforce previous research demonstrating the effectiveness of NLP techniques in identifying patterns and trends within large scientific corpus [2]. More importantly, the outcomes reveal that the growth of GenAI corresponds to measurable linguistic and semantic transformations in scientific communication.

Therefore, while the classification results clearly demonstrate the existence of temporal linguistic differences, they do not independently establish causality between GenAI-assisted writing and the observed discourse transformations.

6.3 RQ3: What Tematic Changes Can Be Identified in AI Research Topics Across the Two Periods?

The BERTopic analysis provided strong evidence of thematic evolution between the pre-GenAI and post-GenAI periods. While the earlier period was dominated by topics associated with technical AI development and computational methodologies, the later period showed a substantial diversification of research themes.

The most significant growth was observed in application-oriented domains such as education, healthcare, business, sustainability, and interdisciplinary AI adoption. These outcomes indicate that AI research increasingly extends beyond core methodological development toward broader societal, industrial, and scientific applications.

This thematic diversification is aligned with recent literature showing that transformer-based NLP methods and embedding-based topic modeling approaches are capable of identifying long-term research evolution patterns within scientific corpusa [1,4]. The present study extends these results by specifically showing how the emergence of GenAI technologies coincides with a measurable expansion of interdisciplinary AI-related research topics.

At the same time, the results suggest that AI is progressively transitioning from a specialized technical discipline into a general-purpose scientific and technological paradigm integrated across multiple domains.

6.4 Implications and Research Contributions

In general, the results suggest a paradigmatic change in AI research, characterized by a transition from method-centered development to generative and application-oriented approaches. This change is not reflected only in the topics studied but also in the language used by researchers.

From a methodological perspective, this study shows the effectiveness of combining classical NLP approaches with modern embedding-based approaches for the analysis of large scientific corpora. Furthermore, it enhances the growing corpus of literature on the impact of GenAI on scientific writing by supplying quantitative evidence of changes in discourse patterns, rather than focusing only on conceptual or moral discussions.

Overall, the findings suggest the existence of a paradigmatic transformation in AI-related scientific discourse following the emergence of GenAI technologies. The results indicate a transition from predominantly technical and method-centered research toward broader application-oriented and interdisciplinary approaches.

From a methodological perspective, this study demonstrates the effectiveness of combining classical NLP techniques, such as TF-IDF and supervised classification, with modern embedding-based topic modeling approaches for analyzing large-scale scientific corpus. The integration of these methods enabled the identification of lexical, semantic, and thematic changes across distinct temporal periods.

The study also contributes to the growing literature on GenAI and scientific writing by providing quantitative evidence of changes in scientific discourse patterns. While previous research has primarily focused on ethical concerns, authorship issues, and conceptual discussions surrounding GenAI, this work offers empirical evidence of measurable transformations in AI-related scientific communication.

7 Conclusion and Future Works

This study explore the evolution of AI-related scientific research before and after the emergence of GenAI. Using a large-scale dataset of approximately 50,000 publications extracted from the Web of Science (WoS), a combination of NLP techniques, including TF-IDF, supervised classification, and BERTopic, was applied to identify lexical, semantic, and thematic changes in scientific discourse across two distinct periods.

The results provide consistent evidence of a significant transformation in AI research following the emergence of GenAI technologies. The lexical analysis

revealed a transition from predominantly technical and methodological terminology toward a broader and more application-oriented vocabulary. Terms associated with infrastructure, algorithms, neural networks, and model development became progressively accompanied by concepts related to education, healthcare, business, sustainability, and generative systems.

The supervised classification results further showed that the two periods can be reliably distinguished based solely on textual characteristics, confirming the existence of systematic linguistic differences between pre- and post-GenAI scientific publications. In addition, the BERTopic analysis revealed a thematic diversification of AI research, showing a change from method-centered investigations toward interdisciplinary and application-driven domains strongly associated with the use of GenAI and LLMs.

Taken together, these results suggest the paradigmatic transformation in AI-related scientific discourse. The outcomes indicate that the rapid development and widespread use of GenAI technologies are influencing not only the topics being searched but also the way scientific knowledge is communicated and structured within the research community.

From a methodological perspective, this study shows the effectiveness of combining classical NLP techniques with modern embedding-based approaches for the large-scale analysis of scientific literature. The proposed framework provides a replicable methodology for studying temporal changes in research domains and scientific communication patterns.

Despite these contributions, several limitations should be acknowledged. First, the analysis was limited to titles and abstracts, which may not fully capture deeper methodological, theoretical, or stylistic transformations present in complete scientific papers. Second, the exclusion of the 2020–2021 transition period intentionally created a clearer separation between the pre- and post-GenAI phases but may have omitted important intermediate developments associated with the early use of transformer-based models such as chatGPT-3. Third, although the results reveal strong correlations between the emergence of GenAI and changes in scientific discourse, they do not establish direct causal relationships between AI-assisted writing tools and the observed linguistic transformations.

Future research may address these limitations by incorporating full-text analysis, stylistic and syntactic metrics, citation networks, and longitudinal analyses including transitional periods. Additional studies may also integrate other scientific databases and explore more advanced semantic approaches based on LLMs to obtain deeper contextual representations of scientific discourse. Furthermore, future investigations could focus more directly on the influence of GenAI tools on authorship practices, writing styles, scientific originality, and transparency in academic communication.

Overall, this study contributes to a broader understanding of how emerging AI technologies are reshaping scientific research and communication, providing a foundation for future interdisciplinary investigations into the long-term impact of GenAI on scientific production.

Acknowledgments. This work was supported by FCT- Fundação para a Ciência e Tecnologia, I.P. by projects: CeDRI, UID/05757/2025 (DOI: 10.54499/UID/05757/2025) and UID/PRR/05757/2025 (DOI:10.54499/UID/PRR/05757/2025); LA/P/0007/2020 (DOI: 10.54499/LA/P/0007/2020). This work has been supported by project Ted-SOEP - 101179402 and project Di-Hor - 2025-1-IT02-KA220-SCH-000359420.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article. During the preparation of this work, the authors used ChatGPT-5 to improve the paper's writing and clarity. After using this tool/service, the authors reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

1. Natural language processing in construction: A scientometric review. Automation in Construction (2022). https://doi.org/10.1016/j.autcon.2022.104169
2. Natural language processing for low-resource languages: A survey. Natural Language Processing (2024). https://doi.org/10.1017/nlp.2024.33
3. A comprehensive review of bert and its applications in nlp. Artificial Intelligence Review (2025). https://doi.org/10.1007/s10462-025-11162-5
4. Natural language processing for materials science. npj Computational Materials (2025). https://doi.org/10.1038/s41524-025-01554-0
5. Chauhan, C.: The impact of generative artificial intelligence in scientific content synthesis for authors. Am. J. Pathol. **194**(8), 1406–1408 (2024)
6. Dhamani, N., Engler, M.: Introduction to generative AI. Simon and Schuster (2026)
7. Leong, S.X., Loh, N.D.: Writing scientific review articles in the age of generative artificial intelligence (2025)
8. Meishar-Tal, H.: Scientific writing in the age of generative ai: usage patterns and ethical considerations among israeli researchers. Ethics & Behav., 1–16 (2025)
9. Reddy, C.K., Shojaee, P.: Towards scientific discovery with generative ai: progress, opportunities, and challenges. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, pp. 28601–28609 (2025)
10. Watson, S., Brezovec, E., Romic, J.: The role of generative ai in academic and scientific authorship: an autopoietic perspective. AI & Soc. **40**(5), 3225–3235 (2025)



Customer Churn Prediction in the Telecom Industry Via a Clustering-Random Forest-XGBoost Approach: Discovering the Role of Delta in PCA Loadings

Chien-Sing Lee[1,2(✉)] and Sharini Sivadas Nanda[1]

[1] Faculty of Engineering and Technology, Sunway University, Subang Jaya, Malaysia
chiensingl@sunway.edu.my
[2] IEEE, New York City, USA
csleester@gmail.com

Abstract. This study is framed within Technology Acceptance Models 1, 2, 3, in identifying key features which influence churn in the telecommunications industry. Data is obtained from 225 Malaysian survey respondents via random sampling. Mining adapts from Bach et al.'s 3-stage approach. For Phase 1, *k*-means, *k*-medoids and hierarchical agglomerative clustering are utilized to determine the most suitable clustering technique to segment customers. Subsequently, bar graphs form more detailed preliminary inferences and cluster profiling. In Phase 2a, features selected via correlation analysis are complemented with principal component analysis. In Phase 2b, non-linear relationships are assessed using Random Forest and XGBoost. The top 15 identified key features are ranked and compared. Outcomes are compared with findings in Phase I and Phase 2a. Findings highlight the viability of a combination of these techniques, depending on the research objectives, the nature of the dataset, the objectives/ nature of the respective algorithms and the domain investigated. Correlation (covariance matrix) and PCA loadings in Phyton reveal different polarities. Moreover, the delta value of PCA loadings are consistent across all clusters. Chi squared and *p* in CHAID performs comparably with RF and XGBoost. A very humble analogy to quantum states due to polarity, robustness in correlational feature selection reflecting information gain and communication's parity bits to ensure data integrity in interconnected systems due to covariance-PCA loadings' delta as identifiers, may be useful for STEM education, depending on model fit, and to finance, e.g. Buy Now Pay Later.

Keywords: Customer Churn Prediction · Telecom · clustering · PCA loadings · delta · Random Forest · XGBoost · Information Science

1 Introduction

The telecommunications industry is a service-oriented industry that significantly contributes to national income especially in developing countries, such as Malaysia. The Information and Communications Technology (ICT) industry has contributed 23.2%

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 289–307, 2027.
https://doi.org/10.1007/978-3-032-30518-3_19

to Malaysia's GDP in 2021. This is projected to rise to 25.5% by 2025 [1]. Similarly, Mordor Intelligence's Industry Report [2] estimates that the 2025 Malaysian telecom market is valued at USD 9.45 billion. However, it is expected to reach a value of USD 11.54 billion by 2030 at a compound annual growth rate (CAGR) of 4.07%.

Despite the substantial business growth of this industry in recent years, telecommunication firms are in an extremely competitive market because many regions have multiple telecommunications providers offering similar services. This makes it easy for customers to switch providers, if they are dissatisfied with the speed, reliability and affordability of the telecommunication services [3]. Customer Churn, where existing customers choose to unsubscribe from their current telecommunication service provider over a given period [4, 5] is thus a significant challenge.

1.1 Problem Statement

[6] point out that the telecommunications industry has the largest annual churn rate of any major industry (over 30%). Thus, companies often turn to 3 main revenue strategies, i.e., attracting new customers, upselling to the existing customers and increasing customer retention duration [3]. These strategies include aggressive pricing strategies, discounts, promotions, and bundle services to differentiate themselves and attract customers. However, if the cost to acquire new customers is high and the newly acquired customers churn too quickly, this will lead to poor returns on acquisition efforts. Furthermore, customer retention programs are sometimes ineffective, because telecommunication firms sometimes offer a one-size-fits-all retention program to entice them back into resubscribing the service. This fails to address individual needs. Hence, personalized customer retention approaches for churn prediction are critical.

1.2 Research Objectives

Research in this field typically focuses on three key areas, i.e., churn prediction modelling, customer segmentation and model explainability [4, 6, 7]. For churn prediction modelling, we need to ensure that the predictive model adequately considers all relevant factors and corresponding relationships to derive improved churn prediction accuracies and generate rules to aid data-driven decision making.

Given the data collected during this study, the research objectives are to:

a) analyse what the *k*-means, *k*-medoids and hierarchical clustering algorithms will reveal about the clusters;
b) evaluate which features for each cluster are more significant based on correlation;
c) identify which churn prediction features for each cluster are more significant based on Random Forest (RF) and XGBoost;
d) test and compare feature selection performance via correlation with PCA loadings vs. RF/XGBoost's
e) gauge potential for chi-squared automatic interaction detector (CHAID) to enhance RF/XGBoost on model explainability
f) investigate role of delta between PCA loadings
g) possibilities of extending to STEM education

1.3 Research Questions

Our research questions are:

RQ1: what will *k*-means, *k*-medoids and HAC reveal about the dataset's and each cluster's distribution?

RQ2: which features for each cluster are more significant based on correlation analyses and PCA loadings?

RQ3: which features for each cluster are more significant based on RF and XGBoost?

RQ4: how comparable is the Chi-Squared Automatic Interaction Detector (CHAID) with RF, XGBoost?

RQ5: can PCA loadings' delta be analogized for STEM education?

1.4 Research Hypotheses

From the problem statements, hypotheses corresponding to the research objectives are:

- Hypothesis 1: *k*-means will perform better than *k*-medoids for the dataset since it is expected to be imbalanced.
- Hypothesis 2: Tenure and contract type are important predictors not only through intra-cluster correlation analysis but also via principal component analysis (PCA).
- Hypothesis 3: Tenure and contract type will be key predictors across all clusters in RF and XGBoost.
- Hypothesis 4: CHAID's feature selection is comparable to RF/XGBoost, and its correlation complemented by *p* provides a leaner solution
- H5: There are analogical possibilities to quantum computing, which may spark interest, but it needs to be scoped and validated as per scientific inquiry

2 Literature Review

2.1 Significant Factors which Influence Churn

[8, 9] have identified Service Usage Amount, Internet Type, Customer Service Interactions, and Promotional Discounts as the key churn drivers [10, 11]. have pointed to high Monthly Invoices and Billing inaccuracies as large contributors to churn. Contract Type and Tenure also indicate high influence, as month-to-month subscribers and customers with shorter tenure are more likely to churn [12, 13]. Tenure, Contract Type, Monthly Charges, Billing Inaccuracies, Internet Type, Service Usage Amount, Service Type, Customer Service Interactions, Promotional Discounts, Total Services Subscribed are the strongest churn predictors in the telco industry.

These are mainly financial [14] in nature. Hence, an area of interest is the Technology Acceptance Model (TAM) 1 by [15], TAM 2 on social factors [16] and TAM 3 on human factors by [17]. TAM 1, 2, 3 are illustrated in Fig. 2a below.

2.2 Mining Techniques

In terms of mining, clustering, and tree-based machine learning models for clustering/customer segmentation are of interest to this study. These will enable cluster profiling. Next, correlation analysis with Principal Component Analysis (PCA) loadings are reviewed. These enable feature selection, specific to each segment.

2.2.1 Clustering

k-means is considered to be one of the more well-known unsuper-vised clustering techniques, mainly due to its simplicity. However, *k*-medoids and hierarchical agglomerative clustering (HAC) may be more useful when the dataset is more balanced. The precision, recall, and F-measure performance of *k*-means, *k*-medoids and HAC are compared and evaluated using the Silhouette Score (SS) and Davies-Bouldin (DB) Index to determine cluster quality [18–21].

k-means and k-medoids involve initialization, assignment and update. *k*-means splits the dataset into *k* clusters while minimizing the within-cluster sum of squares, which is the total sum of squared distances between each point and its assigned cluster centroid (initialization). The distance between points and centroids are calculated as the Euclidean distance. The goal of *k*-means is to reduce this distance. In the assignment step, each point is assigned to the nearest centroid. With each iteration, the distance to the centroid is updated. The iteration ends when it reaches convergence. *k*-medoids uses minimal average dissimilarity to other data points within a cluster, instead of centroids. Hence, medoids are more robust to outliers (extreme values).

HAC takes a bottom-up approach. Similar clusters are merged based on linkage type, e.g. Ward (minimizes variance within clusters), Average (average distance between points) or Complete (maximum distance between points). This process continues until all points have been agglomerated. To visualize the best number of clusters, a dendrogram is created and the optimal cut-off threshold chosen.

2.2.2 Correlation/Covariance and PCA Loadings

Pearson correlation can enable identification of important churn drivers in the high-risk churn cluster [7, 11]. Correlation, however, has its own limitations as it is not able to capture nonlinear relationships [6]. According to [22], an r of 0.1–0.29 implies weak correlation, 0.3–0.5 moderate, and > 0.5 strong. In Phyton, when PCA is applied on raw data, covariance matrix is used. Chi squared is not used to compute PCA loadings. PC1 loadings refer to that of the variable's, whereas PC2 loadings refer to that of churn.

2.2.3 Tree-Based Machine Learning Models

Decision trees build on information gain or its inverse, entropy. They can manage small tabular datasets effectively and accurately, while using less computational resources compared to deep learning models [23–25]. Rule-extraction provides explainability/identification of optimal paths for personalizing customer retention strategies [26].

For Random Forest, the tree splits until the stop criteria such as max_depth and min_samples_leaf are met. The final RF prediction is obtained by aggregating the prediction of all trees by majority voting. XGBoost optimizes loss function through gradient boosting, i.e., XGBoost builds each new tree by correcting the residual errors from the previous prediction. Chi-Squared Automatic Interaction Detector (CHAID) [27–29] applies chi square and p values.

3 Methodology

3.1 Research Design

The Technology Acceptance Models (TAM) 1, 2, 3 (Fig. 1a) are used to frame the research as the authors regard churn as a rejection of the technology [28].'s 3-stage model is presented in Fig. 1b.

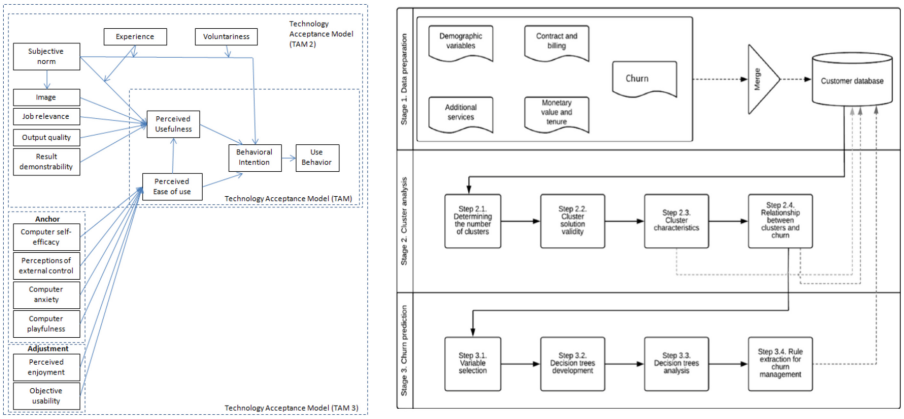

(a) **(b)**

Fig. 1. (a) TAM 1, 2, 3 (b) [28]'s model

Ten hypotheses (H1–H10) are developed based on the significant features identified from literature reviewed in Sect. 2 and mapped to TAM 1, 2, 3 in Table 1. The mapped model (Table 1) is used to develop the survey questions and to develop customer retention strategies.

Table 1. Hypotheses, TAM constructs and explanations

Features/Factors	Mapped TAM Construct(s)	Hypotheses
H1: Promos / Discounts (Q18) [8]	Subjective Norm	If family, friends and colleagues enjoy the promotions or discounts provided, then this subjective norm in terms of social influence increase the user's perceived usefulness of the telecom service.

(*continued*)

Table 1. (*continued*)

Features/Factors	Mapped TAM Construct(s)	Hypotheses
H2: Tenure (Q9) [12, 13]	Result Demonstrability	If tenure is longer, then the benefits and reliability of the service become more evident over time (Result Demonstrability), which in turn increases the perceived usefulness of the telecom service.
H3: Monthly Charges (Q12) [10, 11]		If monthly charges are higher, then the benefits and reliability of the service become more evident over time (Result Demonstrability). This increases perceived usefulness.
H4: Usage (Q15) [8]		If usage is higher, then the benefits and reliability of the service become more evident over time (Result Demonstrability), which in turn increases the perceived usefulness of the telecom service.
H5: Internet Type (Q14) (Ribeiro et al., 2023–4) [8, 9]	Output Quality	If the customer is subscribed to a higher-performing Internet type e.g. fibre optics rather than DSL, then the output quality increases with faster speeds and greater reliability, leading to better user experience, and increased perceived ease of use of the service.
H6: Billing Issues (Q13) (Saleh & Saha, 2023; Mouli et al., 2024) [10, 11]		If billing issues are frequent or unresolved, then output quality as well as customer trust decreases, which in turn decreases the perceived ease of use of the telecom service.

(*continued*)

Table 1. (*continued*)

Features/Factors	Mapped TAM Construct(s)	Hypotheses
H7: Contract Type (Q11) (Chong et al., 2023; Zheng, 2025) [12, 13]	Objective Usability	If the contract type is more flexible (month-to-month), then objective usability is higher based on measurable features such as contract length and increases the perceived ease of use of the service.
H8: Customer Service Interactions (Q17) (Ribeiro et al., 2023) [8]	Computer Self-Efficiency	If there are frequent help requests, then Computer Self-Efficiency is low, and decreases the perceived ease of use of the telecom service.
H9: Service Types/ Bundles (Q16) (Ribeiro et al., 2024) [9]	Perceived Enjoyment	If the service types or bundles are more varied and customizable, then perceived enjoyment increases, which in turn increases the perceived ease of use of the telecom service.
H10: Promotions / Discounts (Q18) (Ribeiro et al., 2023) [8]		If promotions or discounts are offered and are perceived as suitable or relevant to the needs of the customer, then perceived enjoyment increases, and increases the perceived ease of use of the service as the service feels more accessible and rewarding to engage with.

3.2 Survey Design/Data Collection

The Malaysian survey dataset is developed based on the reviewed literature and TAM 1, 2, 3. The Google survey form is distributed via social media from May–June 2025. Hence, sampling is random. The researchers have obtained consent from the survey participants via the Google Form in line with Malaysia's Personal Data Protection Act (PDPA). The collected dataset consists of 225 instances.

3.3 Mining

Random Forest and XGBoost are applied to pre-process feature selection before feeding data to CHAID. In Fig. 1b's step 3.3, the CHAID algorithm [28, 29] serves as the final technique as CHAID is sensitive to irrelevant or noisy features. Moreover, CHAID's

decision rules in Step 3.4 can be used to design interventions. For example, if the rule from Node 6 (Fig. 2) flags new Fiber optics Internet customers whose Tenure is 1 month or less, as having a probability of churning at 88.17%, then the system can recommend customer retention strategies such as providing them with proactive welcome call programs, early discounts, or assign priority customer service support for their first month to reduce dissatisfaction. The telecommunication firm can also investigate whether service issues or price sensitivity is the true cause for customers churning early specifically for Fiber optic users.

Adapting from [28]'s research design (Fig. 1b), in the same domain, Fig. 3 shows this study's 3-stage research design. The differences are a) the dataset, i.e., [28]'s is in Croatia and ours, in Malaysia (Asia), b) the addition of RF and XGB prior to CHAID in our study.

```
/* Node 13 */ (IF tenure <= 1) AND (InternetService = "Fiber optic") AND (MultipleLines = "Yes" AND
MultipleLines = "No phone service") THEN Prediction = 'Yes'; Probability = 51.45%
/* Node 14 */ IF ((tenure <= 1) AND (InternetService = "Fiber optic") AND (MultipleLines = "Yes" OR MultipleLines
= "No phone service") THEN Prediction = 'Yes'; Probability = 78.00%
/* Node 6 */ IF (tenure <= 1) AND (InternetService = "Fiber optic") THEN Prediction = 'Yes'; Probability = 88.17%
/* Node 15 */ IF (tenure <= 5) AND (InternetService = "Fiber optic") AND (MonthlyCharges <= 55.25) THEN
Prediction = 'Yes'; Probability = 53.50%
/* Node 16 */ IF (tenure > 1 AND tenure <= 5) AND (InternetService= "Fiber optic") AND (MonthlyCharges >
55.25) THEN Prediction = 'No'; Probability = 71.93%
/* Node 8 */ IF (tenure > 1 AND tenure <= 5) AND (InternetService = "Fiber optic") THEN Prediction = 'Yes';
Probability = 64.89%
```

Fig. 2. Example of CHAID's rules from [28]

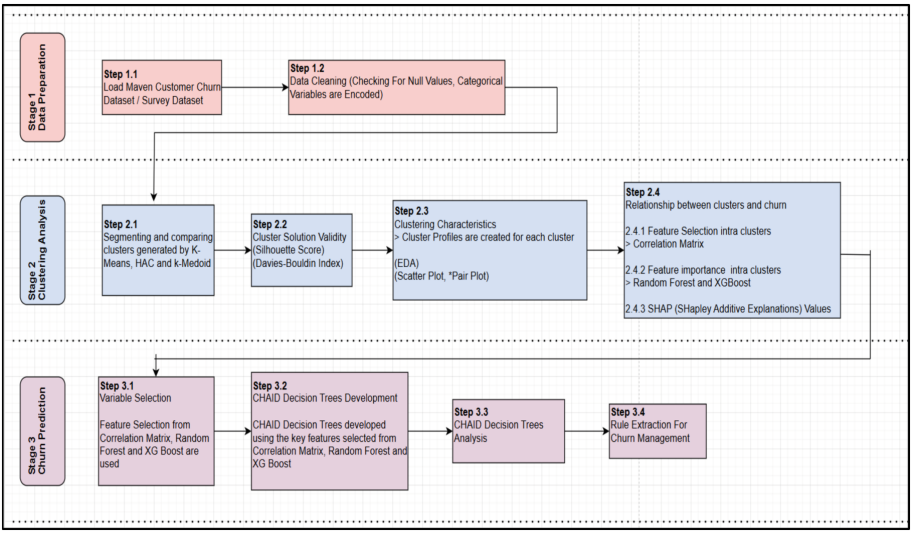

Fig. 3. This study's 3-stage research design adapted from [28]

4 Results

RQ 1: What will *k*-means, *k*-medoids and HAC reveal about the dataset's Malaysian customer segments?

Hypothesis 1: They will reveal slight fuzzy overlaps, due to almost common offerings among providers, differentiated by pricing, capacity/capability of support and provider's incentives.

4.1 Clustering Evaluation

Referring to Table 2, HAC's SS score performs better than *k*-means and *k*-medoids at $k = 2$. The most optimal DB index is also HAC's $k = 2$ at 2.0300. Hence, $k = 2$ has the best cluster quality. The *n* for Cluster 0 is 101, Cluster 1 114 and Cluster 2 10. Hence, only Clusters 0 and 1 are considered. Cluster 2's sample size is too small.

Table 2. SS and DB index of *k*-means, *k*-medoids and HAC across $k = 2$ to $k = 6$

k	SS	DB Index	*k*	SS	DB Index	*k*	SS	DB Index
	k-means			*k*-medoids			HAC	
2	0.0814	3.2377	**2**	**0.0735**	**3.5409**	**2**	**0.2574**	**2.0300**
3	**0.0863**	**2.6510**	3	0.0554	3.5292	3	0.0554	3.3164
4	0.0854	2.6955	4	0.0294	3.5390	4	0.0542	2.7643
5	0.0701	2.8252	5	0.0412	3.1541	5	0.0582	2.7112
6	0.0575	2.8294	6	0.0455	2.9229	6	0.0587	2.6686

4.2 Cluster Profiling

The clusters are profiled based on demographics, service usage, contracts and behaviour. This is to identify the unique aspects of each cluster.

4.2.1 C0 Cluster Profiling

From Table 3, C0's churn rate is 24%, and C1's likelihood of churning is 54%. A possible reason may be the plan type. 84% of C0 customers have family plans, whereas 75% of C1 customers are on individual plans. It is easier for the latter to switch whenever the need arises.

Family plans correlate with 78% of C0's more than 2 years Tenure, and 62% with 2 years Contract. These are binding commitments. Family plans also correlate with most (74%) of C0 customers using fibre optic Internet Type, 44% have Average Monthly GB of 50–100GB and 34% 100-200GB. However, only 51% of C0 customers have Multiple Lines. 68% of C0 customers use mobile data plans.

Due to the Average Monthly GB usage (Table 3), the monthly charges is Gaussian, mainly between RM 101–150 and RM 151–200 (Table 4). Moreover, they are likely to be

mostly well-educated, as most (91%) have no billing issues, and likely to be tech-savvy, as 93% do not need technical support.

Table 3. Key characteristics of C0, and C1 for the Malaysian survey dataset (n = 225)

Features	C0	C1	Plan Type	C0	C1	Contract	C0	C1
n	101	114	Individual	16%	75%	monthly	23%	70%
Predicted churn	24%	54%	Family	84%	25%	1-year	15%	6%
						2-year	62%	24%
Tenure	**C0**	**C1**	**Avg. Monthly Use**	**C0**	**C1**	**Internet type**	**C0**	**C1**
<6 months	0%	1%	< 50 GB	8%	40%	Cable	11%	7%
6–12 months	21%	41%	50–100 GB	44%	40%	DSL	9%	14%
1–2 years	1%	13%	100–200 GB	34%	12%	Fibre	74%	56%
>2 years	78%	45%	Over 200 GB	15%	7%	None	6%	23%

Table 4. Monthly Charge, Multiple Lines, Type of Line, Billing Issues and Tech Support for C0

Monthly Charge	C0	C1	Billing Issues	C0	C1	Tech. Support	C0	C1	Many Lines	C0	C1
< RM 50	2%	24%	No	91%	76%	No	93%	98%	No	49%	77%
RM 50–100	11%	39%	Once	8%	10%	Yes	7%	2%	Yes	51%	23%
RM 101–150	23%	23%	Many	1%	14%						
RM 151–200	22%	11%									
RM 201–250	13%	1%				**Mobile**	**C0**	**C1**	**Land**	**C0**	**C1**
RM 251–300	13%	0%				No	32%	19%	No	78%	93%
>RM 300	17%	2%				Yes	68%	81%	Yes	22%	7%

In terms of promotions, 60% have not received any, possibly because they are mostly satisfied customers (Table 5). Another 30% have received and are still receiving, indicating that some in this group appreciate the discounts more, possibly due to family income. Another 12% however, have not used the discounts received and have allowed the discounts to expire. This implies that possibly 72% of C0 are from M40 or T20.

For this M40 or T20 group, from a cross-selling/market basket analysis (simpler form of [30]'s perspective, the incentive may be in the form of cloud backup, as 97% in C0 have not utilized this and families would need this especially for school-going children's homework. Those in the 18–24 age group will usually use their campus' wifi. For those in the workforce, extending personal cloud backup would be useful. As they get into more senior positions, such personal cloud backup may remain attractive.

Another incentive for those aged 45–54 (23%), and 55–64 (23%), may be beauty/travel-related promotions for female customers (55%). Male customers (45%) may want sports/travel-related promos. Retirees (>65) may want cloud backup to store pictures from travels/gatherings.

Table 5. Possible bases for customer acquisition/retention

Age	C0	C1	Gender	C0	C1	Promo. discount received	C0	C1
18–24	5%	8%	Female	55%	63%	No	60%	54%
25–34	9%	32%	Male	45%	37%	Yes, but expired	12%	16%
35–44	13%	29%				Yes, its ongoing	30%	30%
45–54	23%	12%						
55–64	23%	11%				**Cloud backup**	**C0**	**C1**
> = 65	28%	7%				No	97%	96%
						Yes	3%	4%

4.2.2 C1 Cluster Profiling

For C1, 75% are on individual plan, with 41% having 6–12 months tenure and 45% on more than 2 years' tenure (Table 3). This may imply that the 41% group are renting premises on a yearly basis, whereas the 45% group are renting/staying in premises on a more permanent basis. Based on contract, 70% are on month-to-month, whereas 24% are on 2-year contract. Hence, possibly, 29% out of the 45% are renting their place of stay. More than half (56%) however, have fibre optic Internet Type. This indicates possible use for streaming sports/movies. These are seasonal/on demand. As such, the Average Monthly GB usage is less than 50GB (40%) and for 50–100 GB (40%).

Being on the move explains C1's 81% mobile usage and the lesser charges (less than RM50-RM150) for 86% of C1 (Table 4). Most (76%) compared to C0's 91%, have no Billing Issues, while 14% (compared to C0's 1%) have requested for Technical Support multiple times. Hence, some are not as well-educated or tech-savvy as C0.

In terms of promotions (Table 5), 16% of the promotional discounts have expired, possibly because they may have shifted. Hence, the promotion should enable them to utilize discounts regardless of where they have shifted to. Moreover, most (63%) of C1 are ladies. Hence, from a market basket analysis perspective, they may like discounts for shoes or clothes and groceries and for male customers (37%), travel/clothes.

4.3 Feature Importance Analysis Part I: Correlation (Covariance) Analysis and Principal Component Analysis (PCA)

RQ 2: For each cluster, which features are more significant based on correlation analyses with PCA loadings?

Hypothesis 2: Tenure and Contract Type will be important correlation analysis-PCA predictors.

This section presents the findings of the correlation analysis and Principal Component Analysis (PCA) to guide feature selection for predictive modelling in Random Forest and XGBoost. PC1 represents the feature and PC2 churn. The degree of variance which they represent, implies how much influence PC1 will have on PC2.

In the following Tables, green highlights represents negative correlations (not likely to churn), orange positive correlations (likely to churn). Higher PCA loadings imply higher variance, and thus higher significance.

4.3.1 Correlation (Based on Covariance Matrix) Analysis with PCA Loadings for Cluster 0 (Mostly Customers with Family).

Features which are negatively correlated with Churn (not likely to churn) in descending order (Table 6) are Tenure (strongly negatively correlated), while Average Monthly GB Usage, and Broadband (Internet) are weakly negatively correlated with Churn. Tenure and Number of Referrals are weakly negative, while Unlimited Data is weakly positive (customers may be swayed).

Table 6. Top 10 features' negative vs. positive correlations with Churn and PCA loadings for Cluster 0 (mostly older adults and families)

Feature	Corr.	PC1	PC2	delta	Feature	Corr.	PC1	PC2	delta
Tenure	**-0.885**	**0.100**	**-0.215**	**-0.315**	Broadband	0.239	0.251	0.012	-0.239
Average Monthly GB Usage	-0.277	0.356	0.012	-0.344	Gender	0.141	0.098	-0.040	-0.138
No. of Referrals	-0.116	0.082	-0.255	-0.337	Contract	0.074	0.216	-0.278	**-0.494**
Unlimited Data	-0.086	0.156	0.206	0.050	Monthly Charge	0.068	0.404	-0.201	**-0.605**
Mobile	-0.084	-0.028	0.162	0.190	Age	0.062	0.157	-0.293	**-0.450**

Features which do contribute to the likelihood of Churn are weakly correlated i.e., Broadband (Internet), followed by Gender, Contract, Monthly Charge and Age. However, Age, Contract, and Monthly Charge's PC2 loadings are weakly negative and higher. From Sect. 4.2's findings, bundled offers, e.g. providing discounts for Unlimited Data may increase the Number of Referrals, especially via social media.

4.3.2 Correlation (Based on Covariance Matrix) Analysis with PCA Loadings for Cluster 1 (C1)

The strongest correlation coefficient for C1 is Tenure (Table 7). Contract and Promotional Discounts are weakly and negatively correlated with Churn. The highest positive correlation is Multiple Lines, followed by Cloud Backup. Moreover, though Plan Type's correlation is positive and weakly correlated, it's PC-PC2 loading's delta is the highest, at -0.620. Hence, if the Plan Type fits their lifestyle needs, this can offset other features' influence on the likelihood to Churn. Acquisition/retention strategies can include discounts for multiple lines and cloud backup.

Table 7. Top 10 features most correlated (based on covariance) with Churn and their PCA loadings – Cluster 1 (mostly single and middle/lower income individuals)

Feature	Corr.	PC1	PC2	delta	Feature	Corr.	PC1	PC2	delta
Tenure	**-0.818**	**0.100**	**-0.215**	**-0.315**	Billing Issues	0.163	-0.057	0.202	0.259
Contract	-0.215	0.216	-0.278	**-0.494**	TV / Streaming	0.161	0.286	0.223	-0.063
Promotional Disc. Rec.	-0.178	0.071	0.045	-0.026	Broadband (Internet)	0.158	0.251	0.012	-0.239
Multiple Lines	0.216	0.205	-0.198	**-0.403**	Plan Type	0.155	0.285	-0.335	**-0.620**
Cloud Backup	0.176	0.285	0.379	0.094	Unlimited Data	0.147	0.156	0.206	0.050

4.4 Feature Importance Analysis Part II: Random Forest and XGBoost

RQ 3: Which features for each cluster are more significant based on RF, XGB?

Hypothesis 3: Tenure and Contract type will be significant across all RF, XGB clusters.

Table 8 presents the performance metrics of both models by cluster, followed by the top features contributing to churn prediction in each cluster for both datasets. For this dataset, RF and XGB exhibit similar performance scores for Cluster 0. This implies that the cluster is rather Gaussian and homogeneous. However, RF slightly outperforms XGB for cluster 1, especially for recall (0.106), F1 score (0.062), accuracy (0.059). This implies that this cluster is skewed, possibly due to 63% being female.

Table 8. Model performance metrics by cluster

Cluster 0	RF	XGB	Churn	Cluster 1	RF	XGB	Churn
Accuracy	0.935	0.935		Accuracy	0.912	0.853	
Precision	1.000	1.000		Precision	0.944	0.938	
Recall	0.750	0.750		Recall	0.895	0.789	
F1 Score	0.857	0.857	24%	F1 Score	0.919	0.857	54%

RF = Random Forest, XGB = XGBoost, CR = predicted Churn rate

4.5 Summary of Findings

All 10 hypotheses are accepted as these features are found among the top 16 predictors of churn based on the average rankings by RF and XGB across clusters. Cluster 0 (lowest churn) comprise of customers who prefer bundles of home services and Cluster 1 (highest churn) consists mainly of customers who rely on mobile services with light use of broadband, TV, landline, cloud backup services. From the above, a summarized profile of Clusters 0, 1 and 2 is presented in Table 9 below.

Table 9. Cluster-profiling for Cluster 0, Cluster 1

Features	Cluster 0 (n = 101)	Cluster 1 (n = 114)
Churn status	Low churn rate (CR 24%)	Higher churn rate (CR 54%)
Age	74% adults >35	Gaussian, mainly 25–44
Tenure	78% >2 years	6–12 months (41%), >2 years (45%)
Contract	62% 2-year contracts	70% monthly
Plan Type	84% use family plans	75% use individual plans
Internet Type	74% fibre optic	56% fibre optic, 23% data plan
Multiple Lines	49% single, 51% multiple lines	77% use single line
Avg. Monthly GB Usage	Gaussian 44% 50–100GB, 34% 100-200GB	Skewed left 40% 50–100GB, 40% 100-200GB
Payment Method	46% online banking, 29% credit card	41% online banking, 36% mobile apps
Monthly Charge	Gaussian. 45% RM 101–200	Skewed left. 86% (< RM50–150)
Billing Issues	91% no billing issues	76% no issues, 14% many issues
Promo. Disc. Received	60% no promotion, 30% on-going, 12% expired	54% no promotion, 30% ongoing, 16% expired
Cust. Service Contact Freq.	Contact customer service occasionally	Contact customer service occasionally

5 Comparative Tree-Based Findings: RF, XGB, CHAID

RQ 4: How comparable is CHAID's feature selection to RF/XGBoost?

Hypothesis 4: CHAID's correlation (covariance) complemented by p feature selection is comparable to RF/XGBoost's, for both homogeneous and heterogeneous datasets

RF utilizes majority voting across parallel decision trees. Though rankings differ between RF and XGB, Table 10 and 11 confirm that RF, and XGB concur with regards to each cluster's top 5 and that CHAID is more sensitive towards skewed clusters.

Table 10. RF-XGB ranking across clusters

	Cluster 0		Cluster 1	
Features	**RF Ranking**	**XGB Ranking**	**RF Ranking**	**XGB Ranking**
Tenure	1	1	1	1
Monthly Charge	6	5	5	2
Age	5	3	4	8
Payment Method	3	11	2	6
Number of Referrals	4	9	3	3
Multiple Lines	11	6	9	4
Avg Monthly GB Usage	2	10	7	10
Promotional Discounts Received	10	7	11	9
Contract	9	14	6	5
Gender	7	4	17	13
Broadband (Internet)	8	2	10	15
Plan Type	17	12	14	14
Unlimited Data	12	8	15	11
Billing Issues	14	13	13	17
Internet Type	16	16	8	7
Customer Service Contact Frequency	13	22	12	12
TV / Streaming	15	17	20	18
Mobile	18	15	16	16
Landline	19	18	18	21
Cloud Backup	22	20	19	20
Security Services	21	19	21	19
Technical Support	20	21	22	22

Table 11. XGBoost's top 5 and CHAID's significant vs. RF's top 5 across clusters

RF		**XGBoost vs. RF**		**CHAID vs. RF**	
C0	**C1**	**C0**	**C1**	**C0 (2/3 similar)**	**C1 (3/3 similar)**
Tenure	Tenure	Tenure	Tenure	Tenure	Tenure
Avg. Monthly GB Use	Payment Method	Broadband	Payment Method	No. of Referrals	Payment Method
Payment Method	No. of Referrals	Age	No. of Referrals	Monthly Charge	Monthly Charge
No. of Referrals	Age	Gender	Multiple lines		
Age	Monthly Charge	Monthly Charge	Contract		

Moreover, CHAID's *p*-value is moderated by chi square. For example, for C0, the *chi2* and *p*-value for Payment Method are 3.41 and 0.3324 respectively. CHAID considers it as not significant (Table 9 left). However, for C1, for Payment Method, the *chi2* is 0.96 and *p* value 0.8107 respectively. CHAID regards it as significant. Examples of CHAID-generated rules are presented in Table 12 right.

In summary, the quality of outcomes depends on the quality of the model. RF performs better for more Gaussian and homogeneous datasets, while XGBoost and CHAID are more suited for skewed, imbalanced clusters. Covariance-PCA loadings and its delta may be more stable. We are not certain as the data size is only 225. Further experiments are needed to confirm.

Table 12. Sample CHAID-identified significant features (left) and CHAID rules (right)

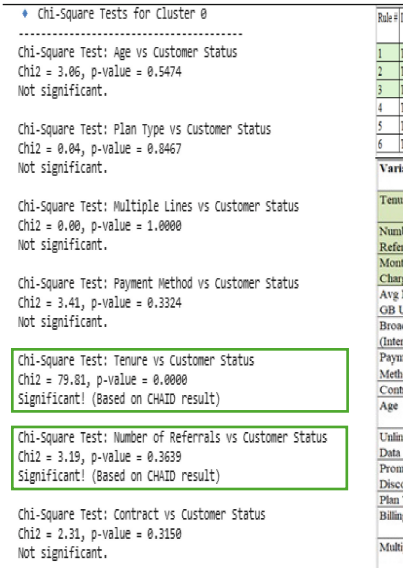

```
♦ Chi-Square Tests for Cluster 0
----------------------------------------
Chi-Square Test: Age vs Customer Status
Chi2 = 3.06, p-value = 0.5474
Not significant.

Chi-Square Test: Plan Type vs Customer Status
Chi2 = 0.04, p-value = 0.8467
Not significant.

Chi-Square Test: Multiple Lines vs Customer Status
Chi2 = 0.00, p-value = 1.0000
Not significant.

Chi-Square Test: Payment Method vs Customer Status
Chi2 = 3.41, p-value = 0.3324
Not significant.

Chi-Square Test: Tenure vs Customer Status
Chi2 = 79.81, p-value = 0.0000
Significant! (Based on CHAID result)

Chi-Square Test: Number of Referrals vs Customer Status
Chi2 = 3.19, p-value = 0.3639
Significant! (Based on CHAID result)

Chi-Square Test: Contract vs Customer Status
Chi2 = 2.31, p-value = 0.3150
Not significant.
```

Rule #	IF Conditions	THEN Churn Probability (%)
1	Tenure ∈ {'1–2 years', '6–12 months'} AND Number of Referrals = 0	100.00%
2	Tenure ∈ {'1–2 years', '6–12 months'} AND Number of Referrals = 1	100.00%
3	Tenure ∈ {'1–2 years', '6–12 months'} AND Number of Referrals = 2	50.00%
4	Tenure = 'More than 2 years' AND Monthly Charge ∈ {'Less than RM50', 'RM 251–300', 'RM 151–200', 'RM 101–150'}	0.00%
5	Tenure = 'More than 2 years' AND Monthly Charge = 'Over RM 300'	12.50%
6	Tenure = 'More than 2 years' AND Monthly Charge = 'RM 50–100'	0.00%

Variable	Global p-value 2 (full)	CHAID Local p-value	CHAID Split Used?	Comments
Tenure	0.0040	4.14E-19	Yes	Root split; highly significant both globally and locally
Number of Referrals	0.0043	0.0149	Yes	Second-level split under short-tenure group; locally significant
Monthly Charge	0.0052	0.0797	Yes (borderline)	Used in long-tenure branch; not significant locally ($p > 0.05$)
Avg Monthly GB Usage	0.0259	Not used	No	Globally modest p-value, but not selected by CHAID for any split
Broadband (Internet)	0.0336	Not used	No	Not selected despite moderate global p-value
Payment Method	0.3324	Not used	No	Not significant globally or in CHAID
Contract	0.3150	Not used	No	Not used in this tree
Age	0.5474	Not used	No	High global p-value, not chosen for splitting
Unlimited Data	0.5791	Not used	No	Not selected
Promotional Discounts	0.4652	Not used	No	Not used despite being a model variable
Plan Type	0.8467	Not used	No	Not significant and not used
Billing Issues	0.8523	Not used	No	Highest global p-value; completely excluded
Multiple Lines	1.0000	Not used	No	Not significant and not split

6 Application to STEM Education

RQ5: Can PCA loadings' delta be analogized for STEM education?

From decades of research in quantum information science, [31] points out that there are three primary branches to quantum information science, i.e., quantum computing, quantum communication, and quantum sensing and measurement [32].'s study on quantum correlations and quantum key distribution highlight this as well. This study has proposed covariance-PCA loadings' deltas as identifier/parity bit in metaverses). There may be very humble analogies [33] to quantum states (polarity), robustness (feature selection) and parity bits (covariance-PCA loadings' delta's consistency across clusters). This would extend the role of PCA in [34], covariance in [35].

7 Contributions and Conclusion

For interconnected systems, data size is huge and require lean, accurate discovery of feature selections and adaptivity to dynamic contexts and changes. We have applied *k*-means, *k*-medoids and HAC to determine the optimal number of clusters. Subsequently, correlation (covariance) analysis with PCA loadings are applied to determine significant features. Findings reveal the strengths of RF, XGBoost, CHAID.

Our very humble contributions are:

a) confirmed [6]'s findings on a large dataset that RF exhibits better performance compared to XGBoost for Gaussian datasets,
b) confirmed CHAID's chi square (covariance)-*p* viable applications to finance, similar to [29], but in the future for Buy Now Pay Later explainability,
c) proposed simplifying very complex concepts for STEM education via [33]'s analogical thinking (e.g. RQ5/Sect. 6). However, the scope-validation measures need to be identified according to progressive scientific inquiry.

In conclusion, Information Science grounds the OSI layers. This paper has built on Information Science but at the application layer. Evolutionary learning are promising. These are however, beyond our capability/ resources.

Acknowledgement. This study was initially Sharini's Masters in Data Science capstone project from March-June 2025, supervised by Lee. Lee prepared the paper, added and structured analyses/ interpretations for RQ5, contributions b), c), inspired by Prof. GL Long's CTIS 2025 quantum computing keynote, after Sharini graduated June 2025. We have obtained informed consent and there is no conflict of interest. Lee thanks Prof. Hean-Teik Chuah, Prof. Hong-Tat Ewe, Prof. Yashwant Prasad Singh for past STEM modeling, OSI, and AI foundations (in Multimedia University), Prof. Ashok K. Goel (one of Lee's past Fulbright Visiting Scholar mentors) on analogy, design, creativity, and Prof. GL Long for not rejecting business model-quantum computing concept similarities.

References

1. CEO Insights Asia Team: Malaysia targets High-Tech status by 2030, aims to lead in ICT. CEO Insights Asia (2024)
2. Mordor Intelligence: Telecommunication Industry in Malaysia Size & Share Analysis - Growth Trends & Forecasts (2025–2030). (2024). https://www.mordor intelligence.com
3. Poudel, S.S., Pokharel, S., Timilsina, M.: Explaining customer churn prediction in telecom industry using tabular machine learning models. Mach. Learn. Appl. **17**, 100567 (2024). https://doi.org/10.1016/j.mlwa.2024.100567
4. Maw, M., Haw, S.C., Ho, C.K.: Customer churn prediction in telecommunication: an analysis on issues, techniques and future trends. In: Conference on Sustainable Utilization and Development in Engineering and Technologies, pp. 148–153. IEEE Press, Malaysia (2019). https://doi.org/10.1109/CSUDET47057.2019.9214592
5. Wagh, S.K., Andhale, A.A., Wagh, K.S., Pansare, J.R., Ambadekar, S.P., Gawande, S.: Customer churn prediction in telecom sector using machine learning techniques. Results Control Optim. **14**, 100342 (2023). https://doi.org/10.1016/j.rico.2023.100342
6. Chang, V., Hall, K., Xu, Q., Amao, F., Ganatra, M., Benson, V.: Prediction of customer churn behavior in the telecommunication industry using machine learning models. Algorithms **17**(6), 231 (2024). https://doi.org/10.3390/a17060231
7. Fatima, G., Khan, S., Aadil, F., Kim, D.H., Atteia, G., Alabdulhafith, M.: An autonomous mixed data oversampling method for AIOT-based churn recognition and personalized recommendations using behavioral segmentation. PeerJ Comput. Sci. **9**, e1756 (2024). https://doi.org/10.7717/peerj-cs.1756
8. Ribeiro, H., Barbosa, B., Moreira, A.C., Rodrigues, R.G.: Determinants of churn in telecommunication services: a systematic literature review. Manag. Rev. Q. **74**(3), 1327–1364 (2023). https://doi.org/10.1007/s11301-023-00335-7

9. Ribeiro, H., Barbosa, B., Moreira, A.C., Rodrigues, R.: Customer experience, loyalty, and churn in bundled telecommunications services. SAGE Open **14**(2) (2024). https://doi.org/10.1177/21582440241245191
10. Saleh, S., Saha, S.: Customer retention and churn prediction in the telecommunication industry: a case study on a Danish university. SN Appl. Sci. **5**, 173 (2023). https://doi.org/10.1007/s42452-023-05389-6
11. Mouli, K.C., et al.: An analysis on classification models for customer churn prediction. Cogent Eng. **11**(1) (2024). https://doi.org/10.1080/23311916.2024.2378877
12. Chong, A.Y.W., Khaw, K.W., Yeong, W.C., Chuah, W.X.: Customer churn prediction of telecom company using machine learning algorithms. J. Soft Comput. Data Min. **4**(2), 1–22 (2023). https://doi.org/10.30880/jscdm.2023.04.02.001
13. Zheng, P.: The impact of contract type and monthly charges on customer churn: comparative analysis across age groups. Adv. Econ. Manag. Polit. Sci. **153**(1), 22–27 (2025). https://doi.org/10.54254/2754-1169/2024.19469
14. Kaplan, R.S., Norton, D.P.: The balanced scorecard-measures that drive performance. Harv. Bus. Rev. (1992)
15. Davis, F.: Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Q. **13**(3), 319–340 (1989). https://doi.org/10.2307/249008
16. Venkatesh, V., Davis, F.D.: A theoretical extension of the technology acceptance model: four longitudinal field studies. Manag. Sci. **46**, 186–204 (2000)
17. Venkatesh, V., Bala, H.: Technology acceptance model 3 and a research agenda on interventions. Decis. Sci. **39**(2), 273–315 (2008)
18. Vieri, J.K., Munandar, T.A., Srisulistiowati, D.B.: Exclusive clustering technique for customer segmentation in national telecommunications companies. Int. J. Inf. Technol. Comput. Sci. Appl. **1**(1), 51–57 (2023)
19. Bilal, S.F., Almazroi, A.A., Bashir, S., Khan, F.H., Almazroi, A.A.: An ensemble-based approach using a combination of clustering and classification algorithms to enhance customer churn prediction in telecom industry. PeerJ Comput. Sci. **8**, e854 (2022)
20. Liu, R., et al.: An intelligent hybrid scheme for customer churn prediction integrating clustering and classification algorithms. Appl. Sci. **12**(18), 9355 (2022). https://doi.org/10.3390/app12189355
21. Bagam, N.N.: machine learning models for customer segmentation in telecom. J. Sustain. Solutions **1, 4**, 101–115 (2024). https://doi.org/10.36676/j.sust.sol.v1.i4.42
22. Cohen, J., Cohen, P., West, S.G., Aiken, L.S.: Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences. Routledge (2013)
23. Grinsztajn, L., Oyallon, E., Varoquaux, G.: Why do tree-based models still outperform deep learning on typical tabular data? Adv. Neural Inf. Proces. Syst. **35**, 507–520 (2022). https://doi.org/10.48550/arXiv.2207.08815
24. Borisov, V., Leemann, T., Sebler, K., Haug, J., Pawelczyk, M., Kasneci, G.: Deep neural networks and tabular data: a survey. IEEE Trans. Neural Networks Learn. Syst. **35**(6), 7499–7519 (2022). https://doi.org/10.1109/tnnls.2022.3229161
25. Rabbani, S.B., Medri, I.V., Samad, M.D.: Attention versus contrastive learning of tabular data: a data-centric benchmarking. Int. J. Data Sci. Anal. **20**, 3069–3091 (2025). https://doi.org/10.48550/arXiv.2401.04266
26. Barsotti, A., et al.: A decade of churn prediction techniques in the TelCo domain: a survey. SN Comput. Sci. **5**, 404 (2024). https://doi.org/10.1007/s42979-024-02722-726
27. Kass, G.V.: An exploratory technique for investigating large quantities of categorical data. Appl. Stat. **29**(2), 119–127 (1980). https://doi.org/10.2307/2986296
28. Bach, M.P., Pivar, J., Jaković, B.: Churn Management in Telecommunications: hybrid approach using cluster analysis and decision trees. J. Risk Financ. Manag. **14**(11), 544 (2021). https://doi.org/10.3390/jrfm14110544

29. Yang, Y., Yi, F., Deng, C., Sun, G.: Performance analysis of the CHAID algorithm for accuracy. Mathematics **11**(11), 2558 (2023). https://doi.org/10.3390/math11112558
30. Rudas, I.J., Pap, E., Fodor, J.: Information aggregation in intelligent systems: an application-oriented approach. Knowl. - Based Syst. **38**, 3–13 (2013). https://doi.org/10.1016/j.knosys.2012.07.02
31. Long, G.L.: Progress in quantum information. National Sci. Rev. **12**, nwaf289 (2025). https://doi.org/10.1093/nsr/nwaf289
32. Franz, T.: Quantum Correlations and Quantum Key Distribution. Universität Hannover, Von der Fakultät fürMathematik und Physik der Gottfried Wilhelm Leibniz (2013)
33. Goel, A.: K.: design, analogy and creativity. IEEE Expert **12, 3**, 62–70 (1997). https://doi.org/10.1109/64.590078
34. Lee, C.S., Singh, Y.P.: Student modelling using principal component analysis of SOM clusters. In: IEEE International Conference on Advanced Learning Technologies, pp. 480–484, Joensuu, Finland (2004). https://doi.org/10.1109/ICALT.2004.1357461
35. Lim, A.H.L., Lee, C.S.: Integrated model-driven business evaluation methodology for strategic planning. Int. J. Bus. Inf. Syst. **3**(4), 333–355 (2008). https://doi.org/10.1504/IJBIS.2008.018037

A Multiclass Model for Financial Distress and Bankruptcy Prediction

Ana Sousa(✉), Ana C. Braga, and Jorge Cunha

ALGORITMI Research Centre, LASI, University of Minho, Guimarães, Portugal
amjorgesousa@gmail.com, {acb,jscunha}@dps.uminho.pt

Abstract. Corporate bankruptcy rarely occurs as an abrupt event. Instead, it is often preceded by a period of gradual financial deterioration. However, most bankruptcy prediction studies rely on binary classification models that distinguish only between healthy and failed firms, thereby overlooking intermediate states of financial instability. This study proposes a multiclass early-warning framework that represents firms' financial condition through three staged categories: healthy, financially distressed, and pre-event, where the pre-event class corresponds to the final observable accounting year before bankruptcy. The empirical analysis uses a large firm-year panel from the Orbis database, focusing on Portuguese manufacturing firms. The raw dataset comprises 37,823 firms and 151,292 firm-year observations from 1997 to 2025. After data preparation, the final analytical dataset contains 140,053 observations. Firm-level financial indicators are combined with lagged macroeconomic variables to account for both firm-specific vulnerabilities and the broader economic environment. A multinomial logistic regression model is estimated to predict firms' financial states while preserving interpretability. The model achieves good out-of-sample performance, with a balanced accuracy of around 0.75 and a macro-averaged F1 score of close to 0.70. The results indicate that the pre-event state can be clearly distinguished from both healthy and distressed firms, whereas the intermediate distress state remains more difficult to identify. Overall, the findings provide evidence consistent with a gradual deterioration process and show that interpretable econometric models can provide useful early-warning signals for firms approaching severe financial vulnerability.

Keywords: bankruptcy prediction · financial distress · multiclass classification · early warning systems · multinomial logistic regression · business failure

1 Introduction

Corporate bankruptcy generates substantial economic costs for firms, creditors, investors, employees, and public institutions. Firm failure may lead to unpaid debt, job losses, disruption of business relationships, and broader economic spillovers [7,13,15,18]. For this reason, the early identification of financially vulnerable firms remains a central objective in financial risk assessment

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 308–326, 2027.
https://doi.org/10.1007/978-3-032-30518-3_20

[31]. Bankruptcy prediction has a long tradition in finance and accounting. Early studies showed that financial ratios help distinguish failing from non-failing firms [1,5,24]. Classical statistical approaches, including discriminant analysis, logistic regression, and hazard models, remain relevant because they provide transparent and economically interpretable frameworks [1,4,24,26,31]. More recent research has increasingly applied machine learning techniques to bankruptcy prediction [31]. Methods such as support vector machines, decision trees, random forests, and gradient boosting can capture complex and non-linear relationships between financial indicators and failure risk [4]. However, higher predictive flexibility often comes at the cost of reduced interpretability, which remains a relevant concern in financial decision-making contexts where stakeholders need to understand the drivers of risk [31]. Despite these developments, three limitations remain important. First, most bankruptcy prediction studies rely on binary classification frameworks that distinguish only between healthy and failed firms [1,24]. However, corporate failure rarely occurs as an abrupt event. Firms often exhibit signs of financial deterioration before bankruptcy, including declining profitability, weaker cash-flow generation, and increasing leverage [1,5,23]. Binary models may therefore overlook intermediate stages of financial instability that are particularly relevant from an early-warning perspective [19,23]. Second, although machine learning models can improve predictive accuracy, their lower transparency may limit their usefulness in settings where interpretability is required. Third, reliable early-warning systems require validation designs that respect the temporal and panel structure of firm-level data. Random cross-validation may mix observations from different time periods or from the same firm across training and testing sets, producing overly optimistic estimates of predictive performance [6,26]. Motivated by these limitations, this study develops a multiclass early-warning framework that represents firms' financial condition using three staged categories: healthy, financial distress, and pre-event, with the latter corresponding to the final observable accounting year before bankruptcy. Rather than estimating dynamic transitions between financial states, the proposed framework provides a temporally ordered representation of financial deterioration. This allows the analysis to examine whether intermediate and pre-event states can be empirically distinguished while preserving the interpretability of the modelling approach. The empirical analysis is based on a large panel dataset of Portuguese manufacturing firms obtained from the Orbis database. The raw dataset contains 37,823 firms and 151,292 firm-year observations for the period 1997–2025. After data preparation, the final analytical dataset includes 140,053 observations. Firm-level financial indicators are combined with lagged macroeconomic variables to account for both firm-specific vulnerabilities and broader economic conditions. A multinomial logistic regression model is estimated as the baseline specification, allowing the probabilities of different financial states to be modelled simultaneously while maintaining economic interpretability. The study contributes to the bankruptcy prediction literature by providing empirical evidence on the usefulness of a multiclass representation of financial distress, combining descriptive, event-time, and predictive evidence, and showing that

an interpretable econometric model with a firm-disjoint validation design can provide useful early-warning signals while reducing information leakage. The remainder of the paper is organised as follows. Section 2 presents the data, modelling framework, and validation design. Section 3 reports the empirical results and discussion. Section 4 concludes.

2 Methods

2.1 Data Source and Sample Construction

Firm-level accounting data were obtained from the Orbis database maintained by Bureau van Dijk. The sample includes Portuguese non-financial firms operating in the manufacturing sector (NACE Rev. 2). Financial institutions were excluded because their balance-sheet structure, regulatory environment, and failure mechanisms differ substantially from those of non-financial firms. The analysis was further restricted to manufacturing firms, following a common approach in bankruptcy prediction studies that focus on industrial firms [1,24]. Firm information was extracted from the *Key Financials and Employees* and *Global Standard Format* (GSF) modules. The GSF format was prioritised because it provides harmonised accounting variables that improve comparability across firms and over time. Because the Orbis interface limits the number of records that can be downloaded at once, the dataset was downloaded in several batches. All raw files were preserved unchanged to ensure traceability and reproducibility of the data preparation process. Data preparation, modelling, and figure generation were implemented in `R` using a fully scripted workflow. Orbis reports financial statements using a relative time structure (*t0*–*t3*), where *t0* corresponds to the most recent accounting year available for each firm. Calendar years were reconstructed using the accounting closing date associated with each reporting period. The dataset was then reshaped into a firm–year panel, in which each observation corresponds to a firm in one fiscal year. The raw panel spans the period 1997–2025, depending on firm reporting availability, and contains 37,823 firms and 151,292 firm-year observations. After constructing financial indicators and applying the data availability filters required for modelling, the final analytical dataset used in the empirical analysis contains 140,053 firm-year observations.

Sample Selection Considerations. Restricting the sample to manufacturing firms is consistent with prior bankruptcy prediction studies focusing on industrial firms [1,24]. This restriction improves the homogeneity of the estimation sample and facilitates a more consistent interpretation of financial ratios across firms [3]. Financial ratios may have different meanings across sectors because firms differ in asset intensity, operating cycles, working-capital requirements, leverage structures, and profitability margins. Therefore, focusing on a single sector reduces cross-sector heterogeneity and strengthens internal comparability. At the same time, this design choice limits external validity, as the results may not generalise directly to other sectors or institutional contexts.

Panel Structure. The dataset is organised as an unbalanced firm–year panel. Each observation corresponds to a firm in a specific accounting year. Firms

enter and exit the database at different points, resulting in a different number of observations per firm.

2.2 Dependent Variable: Financial Condition

The dependent variable captures firms' financial condition using three mutually exclusive states that represent different stages of financial deterioration.

- **Healthy**: firms classified as *Active* in Orbis, indicating normal operating conditions.
- **Distress**: firms classified as *Active (insolvency proceedings)*, indicating firms undergoing formal insolvency or restructuring procedures.
- **Pre-event**: firms classified as *Bankruptcy* in Orbis, for which the outcome is assigned to the last accounting year observed before the bankruptcy event. Therefore, this class represents the final pre-bankruptcy observation, not the bankruptcy event itself.

This classification framework is consistent with the view that financial distress may unfold through different degrees of deterioration before failure [28]. Importantly, the proposed framework does not estimate dynamic transitions between financial states. Instead, it provides a staged and temporally ordered representation of firms' financial condition, allowing the analysis to examine whether intermediate and pre-event states can be empirically distinguished. As is typical in bankruptcy prediction datasets, the distribution of observations across classes is highly imbalanced, with healthy firms representing the vast majority of firm-year observations and financially distressed states occurring relatively infrequently [29]. Figure 1 illustrates the temporal logic used to construct the dependent variable.

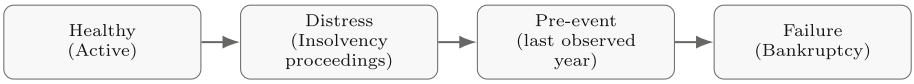

Fig. 1. Conceptual temporal logic used to construct the dependent variable.

Event Definition. The Orbis legal status variable does not always report the exact calendar year of bankruptcy. For this reason, for firms classified as *Bankruptcy*, the pre-event observation is approximated using the last available accounting year before the bankruptcy event. The bankruptcy event itself is not used as a labelled observation, reducing the risk of look-ahead bias arising from post-failure accounting adjustments or incomplete reporting.

2.3 Explanatory Variables

Predictors were selected based on a systematic review of the bankruptcy prediction literature covering 83 empirical studies [27]. The review identified the financial indicators most frequently used in prior research and ensured coverage of the main dimensions of corporate failure risk, including liquidity, profitability, leverage, cash-flow capacity, growth, and operating efficiency. The final set of variables was further refined based on analytical judgement and data availability constraints imposed by the Orbis database, ensuring consistent construction across firms and over time. Explanatory variables are grouped into three categories: (i) firm-level financial indicators, (ii) non-financial firm characteristics, and (iii) macroeconomic conditions. Firm-level variables capture internal financial vulnerabilities, while macroeconomic indicators account for the broader economic environment in which firms operate. Macroeconomic variables were merged by calendar year and included with a one-year lag, so that conditions in year $t-1$ explain firm outcomes in year t. This temporal alignment reduces the risk of look-ahead bias. Ownership-related variables were initially considered but excluded due to severe and systematic missingness for distressed and pre-event firms. In particular, including these variables would have substantially reduced the number of observations in the minority classes and could have introduced additional sample-selection bias. Financial ratios follow standard accounting definitions. Growth variables are computed as annual percentage changes, while changes in net profit and equity correspond to first differences. Table 1 summarises all explanatory variables used in the analysis.

2.4 Macroeconomic Data and Temporal Alignment

Macroeconomic indicators were collected from official public sources to capture the broader economic environment in Portugal. These variables account for aggregate conditions that may affect firms' financial stability, including economic activity, inflation, labour market conditions, financial conditions, and the business cycle [3]. All macroeconomic variables were included with a one-year lag. Since the firm-level panel covers the period 1997–2025, macroeconomic series were assembled for the period 1996–2024. This temporal alignment ensures that macroeconomic conditions observed in year $t-1$ are used to explain firm outcomes in year t, thereby reducing the risk of look-ahead bias.

2.5 Modelling Framework and Validation

Data Preprocessing. All preprocessing steps were performed using the training data only to prevent information leakage from the validation and test sets [2,8,22]. Missing values were imputed using the median calculated from the training sample [11]. Median imputation was selected as a transparent and reproducible strategy suitable for a large-scale financial dataset. However, this approach does not fully account for potentially informative missingness, which is acknowledged as a limitation of the study. Predictors were then standardised

Table 1. Explanatory variables used in the analysis.

Category	Variable
Liquidity	Current ratio
	Quick ratio
	Cash ratio
	Working capital to total assets
	Current assets to total assets
Efficiency	Total asset turnover
	Current asset turnover
	Fixed asset turnover
	Receivables turnover
	Inventory turnover
	Working capital to sales
Leverage and capital structure	Gearing
	Equity ratio
	Short-term leverage
	Interest coverage
	Long-term debt to equity
	Long-term debt to total assets
Profitability	ROA
	EBIT to total assets
	ROE
	Gross margin
	EBIT margin
	Net profit margin
	EBITDA margin
Cash flow	Cash flow to total assets
	Cash flow to current liabilities
	Cash to total assets
Growth	Asset growth
	Sales growth
	Change in net profit
	Change in equity
Firm characteristics	Firm size
	Firm age
	Added value per employee
Macroeconomic variables	GDP growth
	Inflation
	Unemployment rate
	Equity market return
	Interest rate
	Oil price growth

using z-score normalisation based on training means and standard deviations [2,8]. The same transformations were applied to the validation and test sets [2]. Financial ratio datasets often contain extreme values [11]. To reduce the influence of outliers in graphical analyses, the variables used in the event-time trajectories were winsorised at the 1st and 99th percentiles. The values used for model estimation remain unchanged.

Baseline Model. The objective of the study is to estimate the probability that firms belong to different financial states based on firm-level financial indicators and macroeconomic conditions. The baseline specification employs a multinomial logistic regression model. This model estimates the probability of each financial state as a function of the explanatory variables [17]. The healthy state serves as the reference category, meaning that coefficients are interpreted relative to financially stable firms.

Model Interpretability. Multinomial logistic regression was selected because the objective of this study is not only to classify firms, but also to preserve the economic interpretability of the early-warning framework. The model estimates class probabilities simultaneously and allows the financial characteristics associated with different stages of deterioration to be examined directly. This is particularly relevant in financial risk assessment, where stakeholders often need to understand the drivers of risk rather than rely only on predictive accuracy [21]. Although more flexible machine learning algorithms may achieve higher predictive performance in some settings, they often reduce transparency and may be less suitable when interpretability is a central requirement [4,21]. Therefore, multinomial logistic regression provides an appropriate balance between predictive performance, interpretability, and operational usefulness.

Alternative Specifications. Robustness was assessed using several complementary specifications. First, Elastic Net regularisation was applied within the multinomial logistic regression framework to perform data-driven variable selection and control for multicollinearity [32]. The penalty parameter was selected using cross-validation within the training sample. Second, principal component analysis (PCA) was employed as a diagnostic tool to analyse the structure of the financial variables and detect possible collinearity patterns [20]. The resulting components were not used as predictors in the final model. Third, reduced-variable models based on economically motivated financial indicators were estimated to evaluate the trade-off between model parsimony and predictive performance [4,21,32]. These specifications were used to assess the robustness of the interpretable modelling framework rather than to maximise predictive performance through black-box optimisation.

Validation Design. The dataset was partitioned into three mutually exclusive subsets at the firm level: training, validation, and testing [2]. Firm identifiers were used to ensure that observations from the same firm do not appear in more than one subset. This firm-disjoint design prevents information leakage and provides a more realistic evaluation of predictive performance [2,22]. Models were estimated using the training set. Model selection and hyperparameter tuning

were performed on the validation set. The test set was used only for the final evaluation of predictive performance.

Performance Metrics. Model performance was evaluated using overall accuracy, balanced accuracy, and the macro-averaged F1 score. Because the dataset is highly imbalanced, balanced accuracy and macro-F1 provide more informative measures than accuracy alone [16]. Model discrimination was additionally assessed using one-versus-rest ROC-AUC and precision-recall AUC metrics, with PR-based measures being particularly informative under class imbalance [25].

Temporal Instability and Target Drift. The prevalence of financial distress may vary over time as macroeconomic conditions change [10]. Such variation can alter the prediction task across different economic regimes and may lead to target drift [14]. Temporal diagnostics of class prevalence were therefore used to interpret differences between validation and test performance and to assess the stability of early-warning signals over time [6].

3 Results and Discussion

This section presents the empirical results of the multiclass bankruptcy prediction framework. It begins by describing the main characteristics of the sample and the differences across financial states. It then examines the macroeconomic context and event-time patterns in financial indicators prior to bankruptcy. Finally, it reports the data partition, model selection procedure, predictive performance, robustness checks, discrimination ability, and early-warning evidence based on model predictions.

3.1 Descriptive Evidence on Financial Conditions

Table 2 reports descriptive statistics by financial condition. This comparison provides an initial assessment of whether the three proposed financial states differ in terms of liquidity, profitability, leverage, cash-flow capacity, efficiency, growth, and firm characteristics.

Clear differences emerge across the three financial states. Healthy firms generally exhibit stronger liquidity, higher profitability, more stable capital structures, and greater cash-flow capacity than distressed and pre-event firms, supporting the staged classification framework proposed in this study. Liquidity differences are particularly visible in the quick and cash ratios, while profitability separates the groups through ROA, ROE, and margin indicators. Capital structure and debt-sustainability measures also point to greater vulnerability among distressed and pre-event firms, reflected in lower equity ratios, higher short-term leverage, and weaker interest coverage. Cash-flow indicators follow a similar pattern. Some variables, especially working-capital and long-term debt indicators, do not follow a strictly monotonic pattern, suggesting heterogeneity in firms' adjustment strategies and financing structures. Overall, the descriptive evidence is consistent with prior bankruptcy prediction research, which shows that firms approaching

Table 2. Descriptive statistics by financial condition.

Variable	Healthy	Distress	Pre-event
Liquidity			
Current ratio	1.78 (1.04–3.44)	1.08 (0.65–1.72)	0.97 (0.70–1.25)
Quick ratio	1.37 (0.70–2.78)	0.77 (0.39–1.31)	0.58 (0.38–0.82)
Cash ratio	0.31 (0.07–1.04)	0.04 (0.01–0.19)	0.03 (0.01–0.09)
WC/Total assets	0.21 (0.05–0.39)	0.16 (-0.01–0.36)	0.28 (0.14–0.42)
Current assets/Total assets	0.77 (0.55–0.93)	0.73 (0.48–0.91)	0.68 (0.49–0.84)
Profitability			
ROA (%)	2.45 (-0.12–8.31)	-0.54 (-11.54–1.12)	0.07 (-8.09–0.97)
ROE (%)	8.22 (1.19–21.60)	1.00 (-19.04–9.75)	0.98 (-9.10–7.28)
Net profit margin	0.02 (0–0.07)	-0.01 (-0.20–0.01)	0 (-0.08–0.01)
EBIT margin	3.76 (0.56–9.79)	0.46 (-13.24–3.83)	2.73 (-5.53–5.61)
EBIT/Assets	0.04 (0–0.10)	0 (-0.12–0.03)	0.02 (-0.06–0.06)
EBITDA margin	7.47 (2.29–14.93)	2.45 (-9.25–7.67)	6.51 (-1.23–10.74)
Gross profit margin	0.64 (0.48–0.84)	0.68 (0.50–0.94)	0.56 (0.41–0.70)
Leverage/Capital structure			
Equity ratio	0.36 (0.12–0.62)	0.12 (-0.14–0.29)	0.15 (-0.04–0.27)
Short-term leverage	0.36 (0.19–0.64)	0.60 (0.37–0.90)	0.72 (0.50–0.93)
Long-term debt/Assets	0.14 (0.01–0.31)	0.25 (0.10–0.44)	0.07 (0–0.28)
Long-term debt/Equity	0.18 (0–0.77)	0.44 (-0.05–1.70)	0 (0–0.77)
Gearing	34.40 (2.79–112.87)	132.03 (42.99–285.54)	77.83 (2.04–161.68)
Interest coverage	5.77 (1.48–24.96)	0.48 (-7.54–2.42)	1.00 (0.20–1.58)
Cash-flow			
Cash-flow/Total assets	0.06 (0.01–0.13)	0 (-0.10–0.05)	0.03 (-0.04–0.06)
Cash-flow/Current liabilities	0.17 (0.02–0.50)	0.01 (-0.14–0.09)	0.05 (-0.05–0.11)
Cash/Total assets	0.12 (0.03–0.30)	0.03 (0.01–0.11)	0.02 (0.01–0.06)
Efficiency			
Total assets turnover	1.11 (0.67–1.75)	0.86 (0.49–1.52)	1.05 (0.71–1.51)
Current assets turnover	1.65 (1.00–2.63)	1.37 (0.79–2.41)	1.73 (1.20–2.44)
Fixed assets turnover	5.38 (2.19–15.05)	4.35 (1.47–13.81)	4.03 (1.82–9.27)
Inventory turnover	8.71 (3.67–25.62)	6.23 (2.57–20.27)	5.51 (3.12–10.65)
Receivables turnover	4.98 (2.93–9.40)	3.89 (1.99–8.52)	3.43 (2.20–5.60)
WC/Sales	0.18 (0.05–0.38)	0.16 (0–0.43)	0.26 (0.13–0.48)
Growth and variation			
Total assets growth	0.04 (-0.06–0.19)	-0.01 (-0.12–0.15)	0.05 (-0.07–0.21)
Sales growth	0.06 (-0.08–0.25)	-0.04 (-0.23–0.16)	-0.01 (-0.16–0.14)
Δ Net profit	0.32 (-12.73–16.02)	-2.90 (-63.44–19.18)	-2.06 (-82.19–21.12)
Δ Equity	6.16 (-2.66–40.24)	-8.60 (-121.17–7.88)	1.71 (-89.02–25.76)
Firm characteristics			
Firm age	15 (6–27)	17 (7–31)	13 (6–23)
Employees	6 (2–15)	17 (7–38)	40 (18.75–91.25)
Added value per employee	20.63 (13.71–31.81)	13.55 (8.83–19.67)	16.25 (10.95–19.22)

Notes: Median values are reported, with 25th–75th percentiles in parentheses. Growth variables exhibit high dispersion due to firm heterogeneity.

failure tend to exhibit weaker liquidity, lower profitability, higher leverage, and reduced cash-flow capacity [1,5,24,31].

Ownership Structure. The BvD Independence Indicator was initially considered as a proxy for ownership structure. However, it is available for 98.1% of healthy firm-year observations but is almost entirely missing for distressed firms (99.6%) and completely absent for pre-event observations. This structural imbalance makes comparisons across financial states unreliable and may introduce bias; therefore, the variable was excluded from the descriptive analysis and predictive modelling framework.

3.2 Macroeconomic Context

To contextualise firm-level behaviour, macroeconomic indicators representing economic activity, labour market conditions, inflation, financial conditions, and

market performance are analysed for the period 1996–2024. Average GDP growth was 1.51%, with considerable volatility (SD = 3.07), reflecting different phases of the business cycle. The unemployment rate averaged 8.23% and reached a maximum of 16.5% during the European sovereign debt crisis. Inflation averaged 2.2%, also showing variation across the period. Public debt increased substantially over time, reaching 134% of GDP. Equity market returns and oil price growth also displayed high volatility. Figure 2 presents the evolution of real economic conditions in Portugal, measured through GDP growth and the unemployment rate. Two major shocks are particularly evident. The European sovereign debt crisis (2010–2013) was associated with a contraction in economic activity and a sharp increase in unemployment, while the COVID-19 shock in 2020 produced a sudden decline in GDP followed by a rapid recovery.

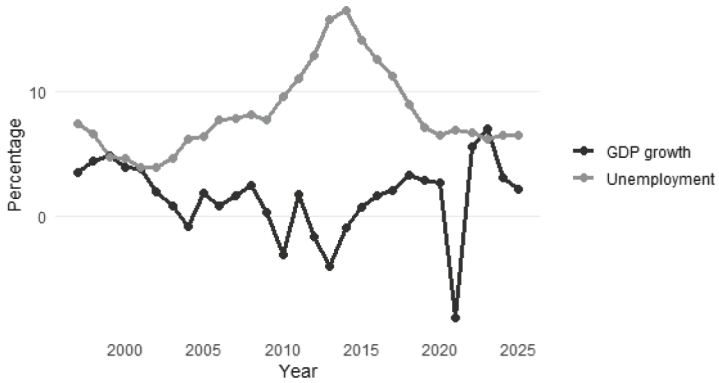

Fig. 2. Real economy in Portugal (1996–2024): GDP growth and unemployment.

Figure 3 summarises financial conditions over the same period, including public debt, equity market returns, and interest rates.

Public debt rose sharply from the early 2000s to the mid-2010s and remained high thereafter. Equity market returns were highly volatile, reflecting changing market expectations and broader financial uncertainty. Interest rates followed a long downward trend until around 2021, followed by an increase associated with post-pandemic monetary tightening. These macroeconomic patterns are relevant because firms' financial conditions are observed under heterogeneous economic regimes. Adverse macroeconomic conditions may increase corporate vulnerability by reducing demand, tightening financing conditions, and weakening firms' ability to service debt [12,30]. However, macroeconomic shocks affect firms unevenly. Their impact depends on firm-specific financial resilience, including liquidity, profitability, leverage, and cash-flow capacity. Therefore, macroeconomic variables are included as lagged controls that complement firm-level indicators rather than replace them as the main source of early-warning information. Having established the macroeconomic context, the next subsection turns

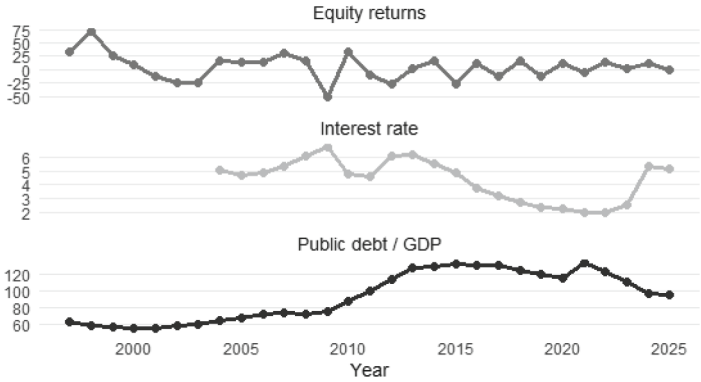

Fig. 3. Financial conditions in Portugal (1996–2024): public debt, equity market returns, and interest rates.

to firm-level trajectories before bankruptcy to examine whether financial deterioration is observable before the event.

3.3 Event-Time Financial Deterioration Prior to Bankruptcy

To examine patterns of financial deterioration before bankruptcy, an event-time analysis was conducted around the bankruptcy event. For bankrupt firms, the reference year (t_0) corresponds to the pre-event observation, defined as the last observed accounting year before bankruptcy. Observations from $t-1$ to $t-3$ represent the three preceding fiscal years. This alignment ensures that all information used in the analysis refers strictly to periods before bankruptcy and therefore avoids look-ahead bias. The notation t_0 is used only within this event-time analysis and should not be confused with the *t0* notation used by Orbis to denote the most recent accounting year available for each firm. Four representative indicators were selected to capture key dimensions of financial distress: liquidity (current ratio), profitability (ROA), capital structure (equity ratio), and operating cash-flow generation (cash flow to total assets). These indicators reflect core dimensions of firms' financial condition that have consistently been identified in the bankruptcy prediction literature as important determinants of corporate failure [1,5,24]. Because financial ratios often contain extreme values, the variables were winsorised at the 1st and 99th percentiles. Median values and interquartile ranges (IQRs) are used to summarise the trajectories. For comparison, a benchmark group of healthy firms was constructed using balanced four-year windows of observations. Each healthy firm contributes a comparable time window aligned to a pseudo-event year, allowing trajectories to be expressed relative to the same event-time structure ($t-3$ to t_0). The number of healthy firms was randomly sampled to approximate the size of the bankrupt sample. Figure 4 shows the median trajectories and interquartile ranges of the selected indicators.



Fig. 4. Event-time trajectories of selected financial indicators for bankrupt and healthy firms. Median values and interquartile ranges are shown. The dashed line marks t_0.

Bankrupt firms exhibit consistently weaker liquidity, with current ratios remaining below those of healthy firms throughout the event-time period. Clear differences are also observed in profitability. Bankrupt firms display persistently lower ROA values, approaching zero before failure, whereas healthy firms maintain positive profitability. Capital structure indicators show a similar pattern: firms approaching bankruptcy have markedly lower equity ratios, indicating weaker capital buffers. In addition, bankrupt firms exhibit lower operating cash flow, while healthy firms maintain higher and more stable cash-flow levels relative to total assets. Overall, the event-time evidence indicates that firms approaching bankruptcy exhibit weaker financial conditions several years before the event. This pattern is consistent with the view that bankruptcy is generally preceded by gradual financial deterioration, reflected in lower liquidity, declining profitability, reduced equity buffers, and weaker operating cash flows, in line with prior findings in the bankruptcy prediction literature [1,5,24].

3.4 Data Partition and Model Selection

The dataset was partitioned into three mutually exclusive subsets at the firm level: training, validation, and testing. Firm identifiers were used to ensure that observations from the same firm were not included in more than one subset. This firm-disjoint design reduces information leakage and provides a more realistic assessment of predictive performance. The training set contains 84,020 firm-year observations from 22,693 firms. The validation set contains 27,989 observations from 7,565 firms, and the test set contains 28,044 observations from 7,565 firms.

Class proportions remain highly stable across the three subsets, as shown in Table 3. In the training sample, 96.66% of observations correspond to healthy firms, 3.09% to distressed firms, and 0.25% to the pre-event condition. The validation and test sets exhibit nearly identical distributions, indicating that the partition preserves the overall class structure of the dataset. This imbalance is

typical in bankruptcy prediction contexts, where failure events are relatively rare and class distributions are highly skewed [29].

Before model estimation, a diagnostic principal component analysis (PCA) was conducted on the training data to examine the structure of the financial variables and identify possible collinearity patterns. The first principal component explains 14.2% of the variance, while the first five components together explain 42.6%. This indicates that the information contained in the predictors is distributed across several dimensions rather than being dominated by a single latent factor. PCA was used only as a diagnostic tool and not for dimensionality reduction in the final predictive model.

Variable selection was performed using Elastic Net regularisation within a multinomial framework. The model was estimated using five-fold cross-validation at the firm level on the training set, evaluating several values of the mixing parameter α. Model selection was based on predictive performance on the validation set, using the macro-F1 score as the main criterion. This metric was prioritised because it gives equal weight to all classes and is therefore more appropriate than overall accuracy in the presence of strong class imbalance.

Table 3. Sample split and class distribution.

Split	Observations	Firms	Healthy (%)	Distress (%)	Pre-event (%)
Train	84,020	22,693	96.66	3.09	0.25
Validation	27,989	7,565	96.70	3.02	0.29
Test	28,044	7,565	96.59	3.15	0.27

As reported in Table 4, the specification with $\alpha = 0.1$ achieved the highest validation macro-F1 and was therefore selected as the final model. The selected model retains 32 predictors capturing liquidity, profitability, leverage, cash flow, efficiency, firm growth, firm characteristics, and macroeconomic conditions. Test metrics are reported for transparency only and were not used for model selection, ensuring an unbiased final evaluation of predictive performance. This result is consistent with the descriptive evidence presented earlier, suggesting that financial deterioration is multidimensional and reflects both firm-level vulnerabilities and broader economic conditions.

3.5 Predictive Performance and Robustness

The test set, which was not used to estimate or select the model, was used to assess predictive performance on unseen data, following standard practice for evaluating generalisation ability [6]. Under the default classification rule, based on the highest predicted probability, the model achieves an overall accuracy of 97.17%, a balanced accuracy of 0.747, and a macro-F1 score of 0.701. These results are reported in Table 5. Because the dataset is highly imbalanced, balanced accuracy and macro-F1 provide a more informative evaluation than overall

Table 4. Elastic Net model selection results.

Specification	α	Pred.	F1 (Val.)	Bal. Acc. (Val.)	F1 (Test)	Bal. Acc. (Test)
alpha_0.1	0.1	32	0.7015	0.7306	0.7009	0.7471
alpha_0.5	0.5	29	0.6994	0.7359	0.7071	0.7482
alpha_1.0	1.0	22	0.6983	0.7338	0.7114	0.7481
alpha_0.9	0.9	23	0.6950	0.7274	0.7095	0.7435

Note: Pred. = predictors retained; Val. = validation; Bal. Acc. = balanced accuracy. Test metrics are reported for completeness only and were not used for model selection.

accuracy alone [9,22,29]. Table 5 also reports class-specific F1 scores, showing substantial differences across financial states.

Table 5. Predictive performance on the test set across alternative specifications.

Specification	Acc.	Bal. Acc.	Macro-F1	Healthy F1	Distress F1	Pre-event F1
Default model	0.9717	0.7471	0.7009	0.9855	0.3859	0.7312
Weighted model	0.8732	0.8826	0.6342	0.9303	0.2780	0.6944
Dual-threshold	0.9715	0.7620	0.7268	0.9855	0.4524	0.7425

Note: Acc. = accuracy; Bal. Acc. = balanced accuracy; F1 = F1-score. Class-specific F1 values are reported for each class.

Healthy firms are identified with a very high F1 score of 0.986, reflecting the model's strong ability to recognise the majority class. The pre-event state also achieves a relatively strong F1 score of 0.731, indicating that the model identifies firms approaching bankruptcy with useful reliability. From an early-warning perspective, this result is particularly important because the pre-event state corresponds to the final observable stage before bankruptcy. In contrast, the distress class is more difficult to detect, with an F1 score of 0.386. This weaker performance should not be interpreted only as model failure. Rather, it suggests that intermediate financial distress is empirically heterogeneous and partially overlaps with both healthy firms and firms approaching the pre-event state [3]. Distressed firms may therefore display financial profiles that are not sharply separated from neighbouring states, making this class the most challenging to identify in a staged early-warning framework. To assess robustness, two additional specifications were examined. First, a weighted multinomial model was estimated using class weights to account for class imbalance. This specification improves balanced accuracy, which increases from 0.747 to 0.883, but this improvement comes at the cost of lower overall accuracy and weaker macro-F1 performance. Specifically, overall accuracy decreases from 0.972 to 0.873, while macro-F1 decreases from 0.701 to 0.634 (Table 5). These results suggest that the weighted model over-corrects for class imbalance by favouring minority-class predictions at the expense of more false positive classifications. Second, a

dual-threshold classification rule was applied to the probabilities produced by the main model. The thresholds were calibrated on the validation set and then applied unchanged to the test set. Under this rule, overall accuracy remains almost unchanged at 0.972, while balanced accuracy increases from 0.747 to 0.762 and macro-F1 rises from 0.701 to 0.727. In particular, the F1 score for distress increases from 0.386 to 0.452, while the F1 score for the pre-event class increases slightly from 0.731 to 0.743. These results indicate that the estimated class probabilities contain useful information about the distress state, but that the default maximum-probability classification rule does not fully exploit this signal.

3.6 Discrimination Performance and Early-Warning Evidence

The model's discrimination ability was further evaluated using one-versus-rest ROC-AUC and precision-recall AUC metrics. In the test set, the ROC-AUC values were 0.895 for healthy firms, 0.886 for distressed firms, and 0.999 for the pre-event class. These results indicate strong discrimination across financial states, particularly for firms approaching the pre-event condition. Because the dataset is highly imbalanced, precision-recall performance provides additional information on minority-class detection and complements ROC-based evaluation [16,25]. The PR-AUC values are 0.994 for healthy firms, 0.432 for distressed firms, and 0.718 for the pre-event state. Together with the class-specific results reported in Table 5, these metrics confirm that the model clearly separates healthy firms from those approaching the pre-event state, while the intermediate distress class remains more difficult to identify, which is a common challenge in imbalanced classification settings [16,29]. The early-warning ability of the model was analysed by aligning observations relative to the year in which the pre-event condition occurs (t_0), corresponding to the final observed accounting year before bankruptcy. This event-time perspective makes it possible to examine whether warning signals become more frequent as firms approach the pre-event year. Table 6 reports the proportion of observations classified as either distress or pre-event in each event-time window, together with the share directly classified as pre-event.

Table 6. Event-time early-warning evidence.

Event time	N	Flagged (distress/pre-event) %	Predicted pre-event %
$t-3$	5	40.0	20.0
$t-2$	13	69.2	15.4
$t-1$	38	76.3	50.0
t_0	75	93.3	82.7

Warning signals become more frequent as firms approach the pre-event year. The share of observations classified as distress or pre-event rises from 40.0% at

$t-3$ to 93.3% at t_0, while the share directly predicted as pre-event increases from 20.0% to 82.7%. Figure 5 illustrates this progression.

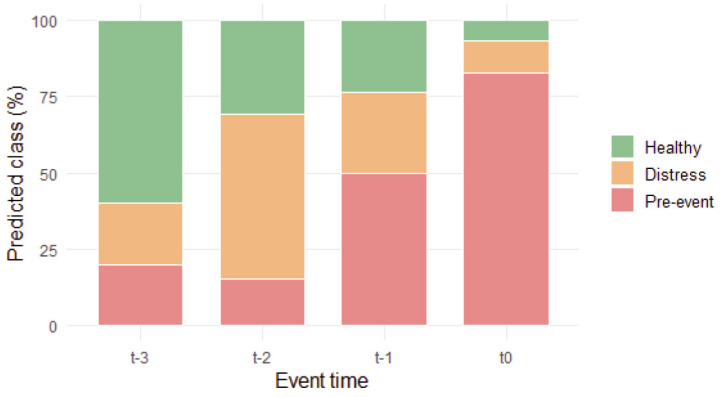

Fig. 5. Distribution of predicted states across event time relative to the pre-event year.

The earliest event-time windows contain relatively few observations and should therefore be interpreted with caution. Even so, as the pre-event year approaches, the share of healthy predictions declines, while distress and especially pre-event classifications become increasingly frequent. This pattern is consistent with the view that financial deterioration tends to emerge gradually rather than appearing only as an abrupt event [3,26]. Overall, the evidence suggests that the framework provides indicative early-warning information by distinguishing firms in more advanced stages of financial deterioration from financially healthy firms.

4 Conclusion

This study develops and evaluates a multiclass early-warning framework that distinguishes between healthy firms, financially distressed firms, and firms in the final observable pre-event state before bankruptcy. Using a large firm-year panel of Portuguese manufacturing firms, the analysis combines descriptive evidence, event-time patterns, and predictive modelling to examine whether different stages of financial deterioration can be empirically distinguished. The findings are consistent with the view that bankruptcy is often preceded by gradual financial deterioration rather than appearing only as an abrupt event. Firms approaching bankruptcy exhibit weaker liquidity and profitability, reduced operating cash flows, and smaller equity buffers. These patterns suggest that financial fragility may become observable before bankruptcy and contain relevant forecasting information. The multinomial logistic model achieves robust predictive performance, with balanced accuracy around 0.75 and a macro-averaged

F1 score close to 0.70 on the test set. The pre-event state is clearly distinguishable from both healthy and distressed firms, whereas the intermediate distress state remains more difficult to classify. This reflects the empirical heterogeneity of distressed firms and their partial overlap with neighbouring financial states. The results also suggest that firm-level financial indicators provide direct early-warning signals, while macroeconomic variables capture broader economic conditions and provide relevant contextual information. This interpretation is consistent with the idea that financial distress is closely associated with firm-specific weaknesses, although the broader economic environment remains important for understanding the conditions under which firms operate. From a practical perspective, the multiclass framework offers advantages over traditional binary bankruptcy models by distinguishing between different degrees of financial vulnerability. This distinction may help stakeholders identify firms at more advanced stages of deterioration and support earlier intervention before bankruptcy occurs. Nevertheless, the results should be interpreted in light of several limitations. Bankruptcy events are rare, leading to strong class imbalance and a limited number of observations in the early pre-event windows. Median imputation provides a transparent and reproducible missing-data strategy, but it does not fully account for potentially informative missingness. Finally, the focus on Portuguese manufacturing firms improves sectoral comparability and supports a more consistent interpretation of financial ratios, but limits external validity, as the results may not generalise directly to other sectors, countries, legal environments, or accounting systems. Future research could compare the proposed multiclass approach with modern machine learning benchmarks and apply it to other sectors and countries. Explicitly dynamic approaches, such as time-to-event models, multi-horizon prediction objectives, transition models, or dynamic model updating, could also be used to model financial deterioration more formally over time. Overall, the findings suggest that combining economically interpretable indicators with a multiclass perspective and a rigorous validation design can provide a realistic and operational early-warning framework for identifying firms at different stages of financial vulnerability.

Acknowledgments. This work has been supported by FCT – Fundação para a Ciência e Tecnologia within the R&D Units Project Scope: UID/00319/2023.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Altman, E.I.: Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. J. Financ. **23**(4), 589–609 (1968). https://doi.org/10.2307/2978933
2. Apicella, A., Donnarumma, F., Isgrò, F., Prevete, R.: Don't push the button: exploring data leakage risks in machine learning pipelines. Artif. Intell. Rev. **58**, 1–58 (2025). https://doi.org/10.1007/s10462-025-11326-3

3. Balcaen, S., Ooghe, H.: 35 years of studies on business failure: an overview of the classic statistical methodologies and their related problems. Br. Account. Rev. **38**(1), 63–93 (2006). https://doi.org/10.1016/j.bar.2005.09.001
4. Barboza, F., Kimura, H., Altman, E.I.: Machine learning models and bankruptcy prediction. Expert Syst. Appl. **83**, 405–417 (2017). https://doi.org/10.1016/j.eswa.2017.04.006
5. Beaver, W.H.: Financial ratios as predictors of failure. J. Account. Res. **4**, 71–111 (1966). https://doi.org/10.2307/2490171
6. Bergmeir, C., Benítez, J.M.: On the use of cross-validation for time series predictor evaluation. Inf. Sci. **191**, 192–213 (2012). https://doi.org/10.1016/j.ins.2011.12.028
7. Bernstein, S., Colonnelli, E., Giroud, X., Iverson, B.: Bankruptcy spillovers. J. Financ. Econ. **133**(3), 608–633 (2019). https://doi.org/10.1016/j.jfineco.2018.09.010
8. Bouke, M.A., Abdullah, A.: An empirical study of pattern leakage impact during data preprocessing on machine learning-based intrusion detection models reliability. Expert Syst. Appl. **230**, 120715 (2023). https://doi.org/10.1016/j.eswa.2023.120715
9. Dasilas, A., Rigani, A.: Machine learning techniques in bankruptcy prediction: a systematic literature review. Expert Syst. Appl. **255**, 124761 (2024). https://doi.org/10.1016/j.eswa.2024.124761
10. Duffie, D., Saita, L., Wang, K.: Multi-period corporate default prediction with stochastic covariates. J. Financ. Econ. **83**(3), 635–665 (2007). https://doi.org/10.1016/j.jfineco.2005.10.011
11. El Madou, K., Marso, S., El Kharrim, M., El Merouani, M.: Evolutions in machine learning technology for financial distress prediction: a comprehensive review and comparative analysis. Expert. Syst. **41**(2), e13485 (2024). https://doi.org/10.1111/exsy.13485
12. European Central Bank: Loss-given-default and macroeconomic conditions. ECB Working Paper Series (2954) (2024). https://doi.org/10.2866/10039
13. Franks, J., Lóránth, G.: A study of bankruptcy costs and the allocation of control. Rev. Finance **18**(3), 961–997 (2014). https://doi.org/10.1093/rof/rft020
14. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., Bouchachia, A.: A survey on concept drift adaptation. ACM Comput. Surv. **46**(4), 1–37 (2014). https://doi.org/10.1145/2523813
15. Graham, J.R., Kim, H., Li, S., Qiu, J.: Employee costs of corporate bankruptcy. J. Financ. **78**(4), 2087–2137 (2023). https://doi.org/10.1111/jofi.13251
16. He, H., Garcia, E.A.: Learning from imbalanced data. IEEE Trans. Knowl. Data Eng. **21**(9), 1263–1284 (2009). https://doi.org/10.1109/TKDE.2008.239
17. Hosmer, D.W., Lemeshow, S., Sturdivant, R.X.: Applied Logistic Regression, 3rd edn. Wiley, Hoboken (2013). https://doi.org/10.1002/9781118548387
18. Jia, H., Li, C., Li, D.: Bankruptcy social cost, government intervention and stock price crash risk. Financ. Res. Lett. **64**, 105449 (2024). https://doi.org/10.1016/j.frl.2024.105449
19. Johnsen, T., Melicher, R.W.: Predicting corporate bankruptcy and financial distress: information value added by multinomial logit models. J. Econ. Bus. **46**(3), 269–286 (1994). https://doi.org/10.1016/0148-6195(94)90038-8
20. Jolliffe, I.T., Cadima, J.: Principal component analysis: a review and recent developments. Phil. Trans. R. Soc. A **374**(2065), 20150202 (2016). https://doi.org/10.1098/rsta.2015.0202

21. Jones, S., Johnstone, D., Wilson, R.: Predicting corporate bankruptcy: an evaluation of alternative statistical frameworks. J. Bus. Finance Account. **44**(1–2), 3–34 (2017). https://doi.org/10.1111/jbfa.12218
22. Kapoor, S., Narayanan, A.: Leakage and the reproducibility crisis in machine-learning-based science. Patterns **4**(9), 100804 (2023). https://doi.org/10.1016/j.patter.2023.100804
23. Lau, A.H.L.: A five-state financial distress prediction model. J. Account. Res. **25**(1), 127–138 (1987). https://doi.org/10.2307/2491262
24. Ohlson, J.A.: Financial ratios and the probabilistic prediction of bankruptcy. J. Account. Res. **18**(1), 109–131 (1980). https://doi.org/10.2307/2490395
25. Saito, T., Rehmsmeier, M.: The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. PLoS ONE **10**(3), e0118432 (2015). https://doi.org/10.1371/journal.pone.0118432
26. Shumway, T.: Forecasting bankruptcy more accurately: a simple hazard model. J. Bus. **74**(1), 101–124 (2001). https://doi.org/10.1086/209665
27. Sousa, A., Braga, A.C., Cunha, J.: Mapping the evolution of bankruptcy prediction models: a systematic literature review (2026). Manuscript submitted to International Review of Financial Analysis
28. Sun, J., Li, H., Huang, Q., He, K.: Predicting financial distress and corporate failure: a review from the state-of-the-art definitions, modeling, sampling, and featuring approaches. Knowl.-Based Syst. **57**, 41–56 (2014). https://doi.org/10.1016/j.knosys.2013.12.006
29. Veganzones, D., Séverin, E.: An investigation of bankruptcy prediction in imbalanced datasets. Decis. Support Syst. **112**, 111–124 (2018). https://doi.org/10.1016/j.dss.2018.06.011
30. Xing, K., Luo, D., Liu, L.: Macroeconomic conditions, corporate default, and default clustering. Econ. Model. **118**, 106079 (2023). https://doi.org/10.1016/j.econmod.2022.106079
31. Zhao, J., Ouenniche, J., De Smedt, J.: Survey, classification and critical analysis of the literature on corporate bankruptcy and financial distress prediction. Mach. Learn. Appl. **15**, 100527 (2024). https://doi.org/10.1016/j.mlwa.2024.100527
32. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. J. Roy. Stat. Soc. B **67**(2), 301–320 (2005). https://doi.org/10.1111/j.1467-9868.2005.00503.x

Simulation-Based Data Generation for On-Demand Bus-as-a-Service Analytics

Vânia Sousa[1,2], Martim Santos[1], Bruno Silva[1], Yanelys Llerena[2], Tiago Sá[3], António C. Vieira[1(✉)], and Maribel Y. Santos[1]

[1] ALGORITMI Research Centre, University of Minho, Campus de Azurém, 4800-058 Guimarães, Portugal
avieira@dsi.uminho.pt

[2] CCG/ZGDV Institute, University of Minho, Campus de Azurém, 4800-058 Guimarães, Portugal

[3] Caetanobus - Fabricação de Carroçarias S.A, Vila Nova de Gaia, Portugal

Abstract. On-demand mobility systems, such as bus services, are gaining importance due to their role in sustainable urban transport. Their flexible and context-specific nature poses challenges for planning, optimization, and performance evaluation, particularly due to the scarcity of related operational data needed to support data-driven decision-making. This work is motivated by a real-world Bus-as-a-Service (BaaS) project, where the lack of available data limited the development of analytical tools. To address this gap, this paper proposes an artifact for the generation of event data (i.e., operational data) related to an on-demand bus service. Thus, the main contributions of this paper include the artifact tailored for on-demand bus services, but also the methodology that was followed for this instantiation. The purpose is that such methodology can be replicated in other cases to generate synthetic data. Without access to real data and without the possibility of gathering data from an existing on-demand bus service, the first step consisted of gathering customers' requirements and analyzing similar business models, resulting in the definition of an Entity-Relationship Diagram that allowed us to identify the data that needed to be generated by the artifact. This artifact simulates the required operational processes, capturing critical elements such as passenger requests, vehicle routing, and other service constraints. The results show that this artifact produces valid simulations for multiple scenarios and enables the generation of structured datasets that meet predefined analytical requirements. Thus, the obtained datasets can be consumed by data-driven analytics.

Keywords: Data analytics · Event data · On-demand bus service · Discrete-event simulation · Synthetic data generation

1 Introduction

Process simulation provides a structured approach to capturing the operational dynamics of complex systems, enabling a comprehensive representation of workflows, decision logic, and business rules [9]. Furthermore, simulation also enables

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 327–343, 2027.
https://doi.org/10.1007/978-3-032-30518-3_21

the analysis of complex and dynamic domains, such as healthcare, manufacturing, or transportation [10,13]. Thus, simulation offers not only an understanding of how workflows and systems function but also serve as a foundation for further analysis.

Apart from the aforementioned benefits, simulation also allows one to generate synthetic data regarding events that occur within the emulated environment, which are the result of the modeled processes' dynamic execution [9,11,13]. This capability can be valuable, particularly for domains where data is scarce, inconsistent, costly to obtain, or nonexistent. In fact, this is even more useful when combined with the capacity of simulation to test what-if scenarios, meaning such artifacts can be used to generate synthetic data, even of hypothetical scenarios.

The motivation for this work stems from a real-world project integrated into a Bus-as-a-Service (BaaS) platform that needs to support analytics for on-demand bus systems. One of the main goals of the project is to allow analytics, even though no data can be used since the platform is still under development.

On-demand bus services, a subset of demand-responsive transport, operate without fixed routes or schedules, dynamically adjusting to passenger requests. These systems aim to combine the efficiency of mass transit with the flexibility of ride-hailing, offering a more sustainable and accessible alternative to private car usage. By harnessing digital technologies such as route optimization algorithms and mobile platforms for booking and coordination, these services support the digital transition in urban mobility. At the same time, they contribute to the green transition by reducing unnecessary mileage, lowering emissions, and encouraging shared travel. Typically, users submit trip requests through an app. Based on aggregated demand, tailored routes are planned that group compatible journeys, enabling buses to operate with higher occupancy and reduced environmental impact compared to traditional fixed-route systems.

Due to the absence of data that supports analytics for this domain, there is a need for an artifact that allows such data to be generated and analyzed. In light of this, the objective of this paper is to propose an artifact designed to simulate the operational dynamics of on-demand bus services, with the main purpose of generating synthetic data pertaining to the simulated events. Furthermore, this paper also proposes the methodology that was followed to establish the requirements of the data that needed to be generated, so that applications in other domains can also instantiate it in their projects. Regarding the main contributions of this paper, they include the generated datasets suitable for downstream analytical applications, including data for multiple scenarios. These datasets include operational variables such as waiting times, demand, service coverage, and resource utilization, all of which are essential for informed decision-making in as on-demand bus service analytics.

This paper is structured as follows. Section 2 analyzes related works in the literature. Section 3 describes the proposed methodological approach that was followed to develop the herein proposed artifact. Section 4 describes the main development activities that were conducted, namely the data requirements elicitation approach, the development of the simulation model, and its validation.

Section 5 explores the generated datasets in descriptive analytics. Section 6 discusses the main conclusions, limitations, and future research directions.

2 Related Work

The scarcity of data in certain domains presents a challenge across multiple disciplines, including Healthcare, Artificial Intelligence (AI), and Social Sciences [4,5,9]. Such a lack of sufficient and high-quality data may severely limit the ability to conduct research, develop solutions, and make informed decisions. To address this challenge, the generation of synthetic data has gained significant attention. Among the most prominent approaches are those based on AI and Machine Learning (ML), which have demonstrated remarkable success in creating datasets that mimic the statistical properties of real-world data [1]. These methods are particularly effective in domains where large volumes of data are available, such as computer vision, natural language processing, and fraud detection. However, they often fall short when there is no data from which to learn patterns, as well as in capturing the dynamic, event-driven nature of many real-world systems [2,13]. These limitations have spurred the exploration of alternative methods, such as simulation, which offer unique advantages for generating synthetic and event-based data in data-scarce domains.

Riess [7] proposed a Python-based simulator that randomizes Business Process Model and Notation-style process models and provides rich temporal parameters, such as durations or inter-arrival times. However, it lacks operational details such as resource pools, capacity constraints, queues, and waiting times. Consequently, while control-flow and timing constraints can be effectively modeled, the tool is less effective in analyzing—and hence generating such data—related to resource utilization, bottlenecks, or service-level phenomena.

Skydanienko et al. [8] addressed the limitations of using real-life event logs to evaluate process mining algorithms, highlighting issues such as data privacy, lack of control over log parameters, and difficulty in isolating specific algorithmic behaviors. The author introduced a tool that generates artificial event logs, supporting constraints over both control flow and data. This tool tackles key challenges in declarative log generation by efficiently producing diverse traces using the Alloy model checker, enabling more robust and scalable algorithm evaluations. While resources may be included as static attributes (e.g., performed by), there is no token-based or queue-based dynamic monitoring. Thus, it cannot model phenomena such as workload or waiting time.

Andrews et al. [2] addressed the challenges of generating high-quality event logs for process mining from relational databases, highlighting common data quality issues and the lack of systematic, automated support for log extraction. It introduces a semi-automated, quality-informed tool that assesses source data, supports mapping to event log attributes, and generates logs suitable for process mining. The work aimed to lay the groundwork for a broader process data quality lifecycle, covering detection, repair, and tracking of quality issues. For this purpose, this work introduced an interactive loop to inject noise until

fitness/precision targets are met. The tool focuses on the semi-automated extraction of event logs from relational databases, guided by data quality dimensions and database constraints, rather than injecting synthetic logging errors. Furthermore, the tool does not simulate process executions, as its purpose is to support systematic event log generation from static data.

Li et al. [6] proposed a tool that learns from real logs using a Generative Adversarial Network with a transformer generator to output synthetic event data. While the training data may contain resource identifiers, the tool is agnostic in terms of considering capacity constraints, such as queue lengths, or resource policies, which are crucial to maintaining quality data in terms of being capable of representing attributes relevant to understanding the dynamic nature of the system, such as queue length or service levels.

As the above analysis of related works shows, despite AI- and ML-based approaches having achieved notable success in replicating the statistical properties of real-world data, they often rely on large training datasets and are often limited in capturing dynamic, event-driven behaviors inherent to many operational systems. Alternative approaches, such as process-oriented simulation, are capable of generating event-based data under controlled and configurable conditions. However, they frequently lack support for operational constructs such as resource contention, queues, and service dynamics, which are essential to accurately model a given system. Furthermore, such tools focus primarily on control flow and data constraints or on extracting logs from static relational sources, without simulating the execution of processes. These limitations highlight a gap in current solutions, i.e., the need for a synthetic data generator that can simulate realistic operational dynamics, including timing phenomena, queuing monitoring, and resource utilization, in order to support data-driven experimentation and analysis in contexts where real data is unavailable or insufficient. The simulation artifact proposed in this paper aims to address this gap. Nevertheless, to identify the required data it is necessary to follow specific steps, which are detailed in the next section.

3 Methodology for Synthetic Data Generation

This section describes the methodology that was adopted to develop the herein proposed artifact for data generation of an on-demand bus system. The main phases that comprise this methodology are illustrated in Fig. 1.

As the figure shows, the first step consisted of gathering the requirements in terms of data that the simulation should be able to generate. For this purpose, two aspects should be stressed. First, there is no data to gather information regarding how the system should operate. Second, as such a system is not yet operating, it is also not possible to gather data by alternative means, e.g., by field observations, interviews, business process models' analysis, or other techniques. Therefore, while traditional requirements elicitation for simulation models comprises some of the mentioned techniques, other means may need to be pondered.

The second step consisted of developing the simulation model. This can be done with any commercial tool (e.g., Simio [12,13]) or by programming the

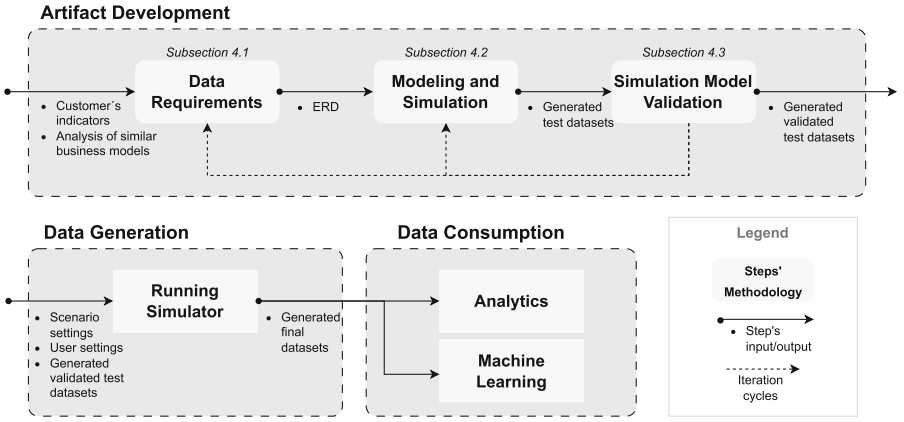

Fig. 1. Methodology followed for the development of the proposed artifact.

simulator (e.g., with Python). The purpose would be to develop models that replicate the behavior of business processes and allow their dynamic simulation, properly accounting for the evolution of resources, queues, and overall system state throughout time. By doing so, it should be possible to generate event-based data related to the events that occur during the simulation.

Having developed the simulation model, it is ready to generate data. However, before considering these data ready to be consumed, the simulation model should be validated, as it could still contain inconsistencies that could be producing data without the required quality. Traditional simulation studies emphasize the importance of simulation model validation, suggesting several approaches, such as statistical analysis, visualization of resulting animation, and comparison of obtained and expected outputs, among others [3,13]. Notwithstanding, due to the importance of obtaining data consistent with the operations of the intended system, alternative means can be pondered, namely, the creation and analysis of visualizations obtained from the generated data. Such visualization could allow one to identify operational inconsistencies, e.g., bus stops without passengers or demand higher than expected. Furthermore, to ensure that the withdrawn conclusions do not apply only to the scenario being tested, it is also important to adapt the validation strategy for multiple scenarios that were previously established. This ensures a higher confidence in the robustness of the simulation model to produce reliable and hence valid synthetic datasets.

Once the simulation model is validated, so are the generated datasets ready to be consumed in the respective layer. On this note, data can be consumed by ML applications, descriptive analytics, or any other analytical approach (e.g., AI, optimization). To validate the proposed methodology and the resulting artifact, synthetically generated datasets will be used for descriptive analytics.

4 Data Generation for On-Demand Bus-as-a-Service Analytics

This section describes the work that was conducted to instantiate the methodology described in the previous section. Concretely, first, the requirements are detailed, which consist of the proposal of the Entity-Relationship Model (ERD) for this system, thus establishing the main concepts that the simulation must contain, as well as the respective data. The next subsection describes how the simulation model was developed. The last subsection covers the validation process that was followed to validate the simulation model.

4.1 Data Requirements

This section describes how the main requirements, in terms of data that needed to be generated, were gathered. As per the methodology illustrated in Fig. 1, first, the requirements in terms of indicators that the project's customers identified as most relevant are gathered. In other words, the generated data should allow such indicators to be analyzed. However, per se, this would not allow one to understand how the on-demand bus system should operate. As such, this approach is complemented by benchmarking other similar business models (e.g., VitrusBus[1], a bus-on-demand service operating in Guimarães, Portugal). Such analysis, along with a careful validation with the project stakeholders, allowed us to identify relevant indicators, as well as specific attributes required to compute such indicators. Table 1 summarizes this mapping, between indicators and attributes, showing how the simulation model will enable the computation of each indicator required by the project's stakeholders.

The decomposition of relevant business indicators into the necessary data attributes, listed in Table 1, allowed us to organize such data in an ERD, as depicted in Fig. 2. As can be seen, six entities were considered, namely `Customer`, `Stop`, `Booking`, `Planned_Route`, `Bus`, and `Trip`.

This also allowed us to extract some relevant insights, such as events that need to be simulated, e.g., customer needs to be able to book trips, or buses need to perform specific trips, which depletes their battery. In fact, a customer can make one or more bookings, each of which includes a requested origin and destination stop, as well as the desired date and time. Each booking is associated with a planned route that aggregates similar requests from different customers. In turn, a planned route can be executed multiple times as separate trips. Each trip represents the operational realization of a planned route and includes actual data, such as execution time and distance traveled. Stops are defined as the physical locations for boarding and alighting, and each booking specifies a bus stop of an origin and another of a destination. The cardinalities between entities reflect these relationships: a one-to-many relationship exists between customers and bookings, bookings and planned routes, and planned routes and trips. Additionally, each booking has two many-to-one relationships with the stop entity: one

[1] https://vitrusambiente.pt/vitrusbus/.

Table 1. Mapping between the indicators given by the project's customer and the identified attributes.

ID	Indicator Description	Relevant Attributes
I1	Total number of bookings	Booking.Booking_ID
I2	Total number of customers	Customer.Customer_ID
I3	Average trip distance	Trip.Distance
I4	Booking trends by date	Booking.Booking_Timestamp
I5	Average trip duration	Trip.Actual_Start_Time, Trip.Actual_End_Time
I6	Main customers (by number of bookings)	Booking.Customer_ID, Customer.Name
I7	Times with highest volume of bookings	Booking.Requested_Time, Booking.Requested_Date
I8	Total active customers	Customer.Customer_ID, Booking.Booking_Timestamp
I9	Total number of trips	Trip.Trip_ID, Bus.Bus_ID
I10	Total battery consumption	Trip_Energy.Consumption_KWh
I11	Energy consumption by route	Trip_Energy.Consumption_KWh, Planned_Route.Route_ID
I12	Route usage map (Origin → Destination)	Planned_Route.Planned_Route_ID, Route_Stop.Stop_ID, Trip.Trip_ID
I13	Bus Utilization and Energy Consumption Analysis	Bus.Bus_ID, Trip.Trip_ID, Trip.Passengers_Onboard, Trip_Energy.Consumption_KWh
I14	Hourly Battery Consumption by Bus	Bus.Bus_ID, Battery_State.Level_Percent, Battery_State.Timestamp
I15	Daily evolution of bus trips	Bus.Bus_ID, Trip.Trip_ID, Trip.Date
I16	Customer segmentation by booking frequency	Customer.Customer_ID, Customer.Name, Booking.Count
I17	Bookings vs travels over time	Booking.Booking_ID, Booking.Booking_Timestamp, Trip.Trip_ID

for the origin and one for the destination stop. It should also be noted that, in addition to the data attributes identified during the requirements gathering phase, additional ones can be generated and calculated as a result of the dynamic execution of the simulation model. Some examples include the requests' waiting times, bus utilization rates, the number of passengers queuing to board the bus per stop, among others.

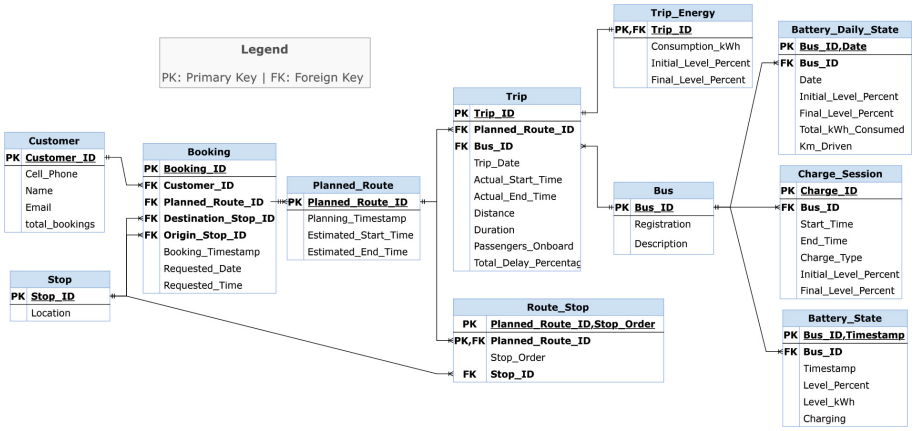

Fig. 2. Relationship Diagram associated with the elicited data requirements for the simulation-based generator.

4.2 Modeling and Simulation

This subsection describes the development of the simulation model that was used to develop the on-demand bus system. The simulation model was implemented in `Python` version 3.12.4[2], namely using the `SimPy` library version 4.1.1[3]. In addition, other libraries were also used, from which the following can be highlighted. `NetworkX`[4] was used to generate the graph that specifies the locations of the bus stops, as well as the distance between them and the different routes that allow buses to travel within the simulation model. In turn, `Faker`[5] was used to generate static contextual data, such as customer names or their contact details. Algorithm 1 provides an overview of the simulation model implemented, by presenting the respective pseudo-code. Initially, the simulator reads a set of parameters from user inputs. While several were implemented, the ones listed in Table 2 can be highlighted.

When the simulation begins, some elements need to be created. First, a list of stops, as well as the respective complete graph. This graph contains a central point to be used as a dockyard from which all buses leave and return after finishing a route. Thereafter, a bus fleet is also created. Finally, other data structures are also created, of which the set of requests can be emphasized, which will serve as a pool for all bus service requests. These requests will act as the main entities of this simulation model.

Afterwards, the main loop in the simulation can initiate, which runs until the simulation clock (specified by the `simtime` variable) reaches the user-defined simulation length parameter (`SIM_LENGTH`). At the start of each new day (every

[2] https://www.python.org/downloads/release/python-3124/.
[3] https://simpy.readthedocs.io/en/4.1.1/contents.html.
[4] https://networkx.org/.
[5] https://faker.readthedocs.io/en/master/.

Table 2. Examples of simulation parameters that can be edited by users.

Symbol	Description
SIM_LENGTH	Duration of a simulation run in minutes
NUM_BUSES	Number of buses used in a simulation run
DEMAND	Expected average number of requests per hour
NUM_STOPS	Number of bus stations. Some stops are more popular
SCALE	Defines the area for the simulation
AT	Maximum time for next bus service in minutes

1440 min), the system plans routes for all pending received requests, with the *plan_routes_for_day(requests)* method. This planning is necessary in order to organize the received requests into future bus routes. This starts by filtering the unprocessed requests scheduled for the current day and sorting them by their scheduled time. Bus routes are then created by combining multiple requests with similar origins, destinations, and compatible schedules. These route groupings are defined by taking into account a few constraints, e.g., not exceeding the defined bus capacity and the maximum expected route duration of 4 h (this estimation can be altered by the user). This process also ensures that no routes are defined with less than two stops. When the planning is finalized, the `logs` function is used to save the required data for later export, in this case into CSV files.

Throughout the simulation, new requests are generated at discrete intervals in time, following an Exponential distribution with an average value related to the expected number of passengers per hour (specified by the `DEMAND` parameter). A randomly selected set of origin and destination bus stops is assigned to these requests, as well as a desired schedule for their bus travels and the associated number of passengers. Finally, such data is also saved with the `logs` function for later export.

When a route's scheduled departure time matches the current simulation time, the route is executed, i.e., the route is assigned to a bus, and it travels to the next bus stop at a randomly assigned speed, following a Uniform distribution with a minimum of 0.83 and a maximum of 1.17 km per minute. In its turn, the travel time is calculated by the *travelDuration(bus.speed, simtime)* method, which considered the speed assigned to the bus, as well as on the current simulation time, thereby allowing to account for the traffic intensity. Moreover, the bus battery is also consumed during each bus travel. To verify if the bus needs to charge, the *charge_check(bus, simtime)* method monitors the battery, as well as the current simulation time, to decide whether the bus should charge, i.e., in periods of higher demand, the battery charging policy changes. In order to allow the simulation clock to advance to the time in which each bus arrives at the intended bus station (next event), the `wait(event)` function mimics such behavior. At each stop, the simulation models (with the `delay()` method) pas-

senger boarding (0.03 min per passenger) and exiting (0.02 min per passenger) the bus before proceeding to the next stop.

Algorithm 1. Pseudocode of the bus simulation implemented in SimPy.

Inputs: SIM_LENGTH, NUM_BUSES, DEMAND, NUM_STOPS, SCALE, AT

Initialization:
    stops ← Create list of stops based on NUM_STOPS
    network ← Create the network of bus stops, considering SCALE
    buses ← Create set of buses based on NUM_BUSES
    requests ← Create empty set of bus requests
    routes ← Create empty set of routes *// including attributes, e.g., battery level*

while simtime < SIM_LENGTH **do**
    if wait(**simtime** % 1440 == 0) **then** *// Start of a new day*
        routes = plan_routes_for_day(requests) *// Planning for pending requests*
        logs(routes) *// Save data logs for later export*
    end if
    if wait(**simtime** == nextRequest) **then** *// Create new request*
        $request.scheduledTime \leftarrow$ **simtime** $+ \mathcal{U}(1440, AT)$
        $request.numPassengers \leftarrow$ if $\mathcal{U}(0,1) \leq 0.5$ then 1 else $\mathcal{U}2, 4$
        $interarrival \leftarrow \mathcal{E}(60/DEMAND)$ *// When to create next request*
        $request.destination \leftarrow$ random bus stop
        $requests.add(request)$
        $logs(request)$ *// Save data logs for later export*
        $nextRequest \leftarrow$ **simtime** $+ interarrival$
    end if
    for each (bus, route) in $(b, r)|b \in buses, r \in routes$ **do**
        if wait(**simtime** == scheduledTime) **then** *// Start new bus route*
            for each stop in route **do**
                $bus.speed \leftarrow \mathcal{U}(0.83, 1.17)$ km per minute
                $bus.travelTime \leftarrow travelDuration(bus.speed,$ **simtime**$)$
                bus.move($stop$) *// Bus travels to the next stop*
                $bus.consumeBattery()$ *// Dynamic consumption of battery*
                $charge_check(bus,$ **simtime**$)$ *// Verifies if the bus need to charge*
                wait(stop) *// Wait for bus to arrive at next stop*
                delay() *// Delay for each passenger boarding or exiting the bus*
                logs(route, bus) *// Save relevant route indicators*
            end for
        end if
    end for *// Go to time of next discrete event*
end while *// End simulation and save data logs*

Finally, relevant indicators, such as waiting time and travel durations, are calculated, and relevant data is saved for later export with the `logs` function. Since SimPy is a process-oriented, discrete-event simulation framework, the next event to simulate is determined based on the simulation time when the event was

scheduled. I.e., at this point, the simulation engine decides whether new routes should be planned, new requests generated, or bus routes modeled.

After completing all stops, the bus returns to the dockyard. The simulation clock advances in discrete time steps until reaching the end of the simulation, when all data events are saved to the respective files. This process captures the logic while handling stochastic demand and travel time variability. Therefore, between different simulation runs, different datasets may be generated. However, studying this stochasticity falls outside the scope of this research. The next subsection covers the validation process that was conducted.

4.3 Simulation Model Validation

This subsection describes the validation of the simulation model, which was carried out to ensure that the model effectively represents the domain, with the ultimate goal of ensuring that the generated data are realistic and useful for analysis. The validation process consisted of analyzing the visualizations obtained from the generated data in order to analyze the coherence of different test scenarios. As per the methodology proposed in Fig. 1, first, a set of scenarios is defined, which guides this validation process. For all scenarios, the following default parameter values were used: maximum antecedence of 1 week (AT) for the scheduling of requests, SCALE of 35 and 20 bus stops (NUM_STOPS); parameters defined in Table 2. In turn, for the defined scenarios, the following parameters are altered:

- **Scenario 1:** The demand parameter is set to 5 and the number of buses is set to 5. Obtained visualizations are depicted in Fig. 3.
- **Scenario 2:** The demand is increased to 20, and the number of buses is set to 10. Obtained visualizations are depicted in Fig. 4.

Both figures confirm that the generated data behaves as expected, exhibiting realistic patterns consistent with the modeled transport process.

The top-left graphs of both figures indicate that the route planning process operates efficiently, generating routes that service multiple requests simultaneously. While, indeed, a planning process was implemented which may considerably affect the operation of the entire simulation, this falls out of the scope of this paper. As expected, the number of requests served per route increases significantly in correlation with higher demand, since the number of buses remain unchanged from one scenario to the other. The top-middle graphs demonstrate that total travel time remains within a defined limit, due to the battery of the buses. As can be seen, under higher demand, the trend is for additional routes being required to satisfy the higher number of passengers with a higher variety of requests. The same can be concluded from the top-right graphs of both figures. In turn, the bottom-left graphs validate that the passenger capacity limit is respected. Furthermore, the increase in demand in scenario 2 shows a realistic increase in passenger load per bus. The bottom-middle graphs, depicting the distribution of distances traveled per bus, reveals a uniform workload across the

fleet. This balanced distribution further suggests that the scheduling algorithm optimizes resource allocation effectively, avoiding the over-utilization of specific vehicles. Finally, the bottom-right graphs illustrate the evolution of requests over time. The graphs differentiate between received, processed and rejected requests, since customers need to schedule their trips, which need to be processed, i.e., allocated to a bus and a route (as per the mentioned planning process).



Fig. 3. Results of scenario 1.

Regarding the rejected requests, this pertains to situations, such as requests that would originate routes consuming too much battery, or simply not enough bus capacity. In some of these situations, customers can simply wait on a queue, however, in some cases they are considered as rejected. The relatively low volume of rejected requests in scenario 1 underscores the effectiveness of the algorithm responsible for allocating the requests to bus routes, and, as expected, under higher demand this number also slightly increases. It is also possible to see that the system achieves a steady state, in what concerns the three variables in question, after roughly one week.

In conclusion, the comparative analysis of these scenarios demonstrates that the simulator is both robust and highly responsive to parameter adjustments. By maintaining operational constraints while achieving a stable steady state, the model proves to be a reliable tool for testing different strategies in on-demand bus transport, providing a realistic foundation for dataset generation and consequent analytics.



Fig. 4. Results of scenario 2.

5 Results

This section proposes to evaluate the effectiveness of the proposed methodology, the resulting artifact, and the datasets it generates by presenting a dashboard to allow descriptive analytics over the generated datasets. The dashboard allows the visualization of the indicators that had been previously identified during the requirements elicitation phase, since such indicators were deemed as relevant for the project's customers. As such, the dashboard was implemented as a proof-of-concept to demonstrate that the synthetic data can effectively support analytical tasks, particularly descriptive analytics. For this purpose, this section displays two different dashboards: one that integrates indicators related to customers and their respective requests (Fig. 5) and another that integrated indicators related to routes and buses (Fig. 6).

The developed dashboard comprises different visualizations, each corresponding to an indicator elicited during the requirements phase. The first visualization displays the total number of bookings (I1), providing an overview of the volume of service usage during the simulation period. The second shows the number of unique users (I2), reflecting the diversity of the user base. The third visualization presents the average distance per trip (I3), offering insights into travel behavior and typical route lengths. The fourth visualization highlights booking trends over time (I4), allowing the identification of usage peaks and seasonal patterns. The fifth graph focuses on the average duration of planned routes (I5), which supports the analysis of operational efficiency. The sixth chart lists the top customers by number of bookings (I6), helping to identify frequent users or institutional clients. Finally, the seventh visualization shows the distribution of bookings across time slots (I7), enabling the detection of peak hours and

demand concentration. In turn, Fig. 7 and Fig. 8 shows that it is also possible to drill down on a route that a user may select. Overall, the results provide empirical evidence that the adopted methodology leads to the generation of usable, structured data aligned with analytical needs.

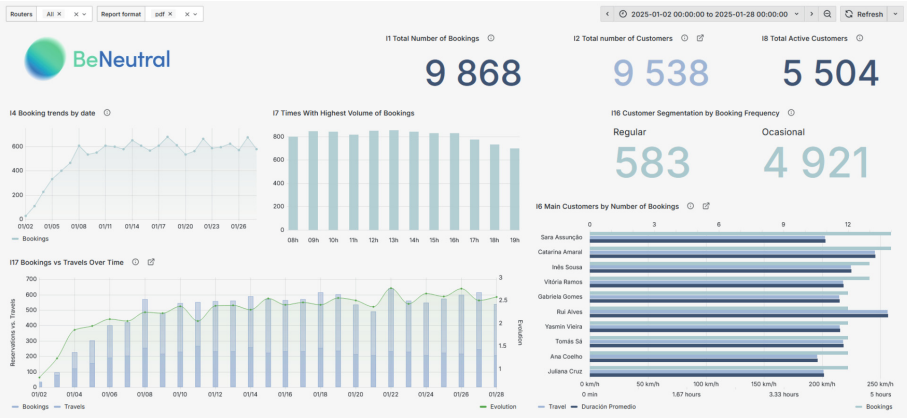

Fig. 5. Dashboard integrating indicators related to customers and their requests.

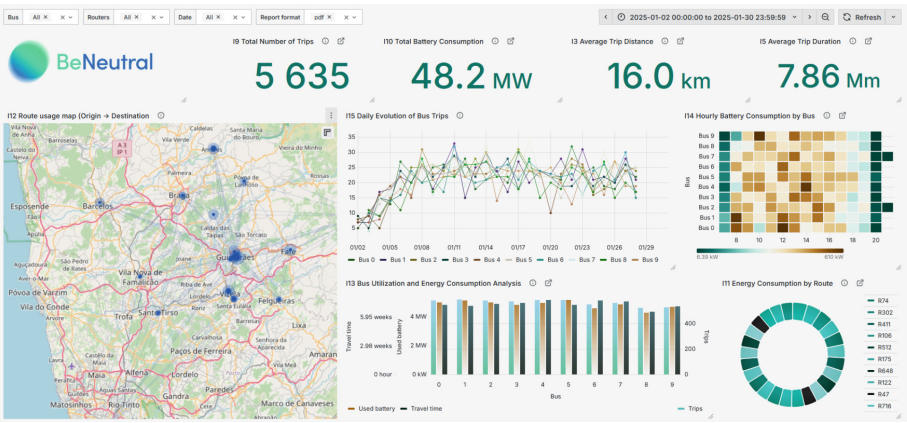

Fig. 6. Dashboard integrating indicators related to routes and buses.

The first dashboard (cf. Fig. 5) reveals a predominantly occasional user base, with regular customers representing a small fraction of the total. Booking trends show an initial ramp-up consistent with simulation cold-start effects, stabilising into a steady daily demand pattern thereafter. The hourly distribution suggests

a mild midday concentration without strongly pronounced peaks. The bookings versus travels chart highlights a persistent gap between reservations and completed trips that gradually narrows over time, reflecting the route planning mechanism adapting to accumulated demand.

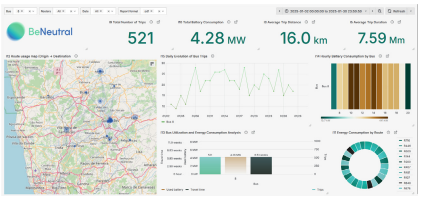

Fig. 7. Dashboard of Fig. 5 filtered by bus.

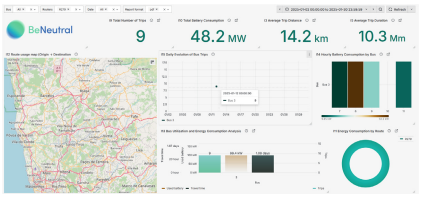

Fig. 8. Dashboard of Fig. 5 filtered by route.

The second dashboard (cf. Fig. 6, Fig. 7 and Fig. 8) shows a well-balanced fleet utilisation, with trip counts and energy consumption distributed evenly across buses throughout the simulation period. An initial stabilisation phase is visible in the first days of operation, after which each bus reaches a consistent activity level. The hourly battery consumption heatmap confirms that energy usage is concentrated within operational hours, with some buses exhibiting higher midday consumption. At the route level, a small subset of routes accounts for higher energy consumption, pointing to differences in route length or demand intensity. The filtered dashboards demonstrate the system's drill-down capability at both bus and route level. Individual bus metrics align closely with fleet averages, confirming consistent behavior across the fleet without systematic assignment bias. Route-level filtering exposes the heterogeneity of the network, with less frequent routes showing distinct distance and duration profiles compared to the overall average.

6 Conclusions

This work addressed the pressing challenge of data scarcity in emerging mobility systems, particularly in the context of on-demand bus services. In the absence of real-world operational data - critical for supporting analytics and decision-making - as well as a real operating system from which data could be gathered, this paper proposed and implemented an artifact capable of generating high-quality synthetic data. The artifact is based on process modeling and simulation, meaning that business processes are modeled and can be simulated, thus considering essential components and dynamics of on-demand bus systems.

The main contributions of this work comprise not only the proposed artifact but also the systematic methodology that guided its construction. Grounded on traditional simulation development methodologies, it extends such methodologies by considering essential steps when aiming to generate high-quality data,

especially without the possibility of gathering data. Namely, the data-driven requirements elicitation process fostered the project stakeholders, which, along with an analysis of similar business models, allowed us to identify attributes that the data should contain. The methodology also reckons that the validation of the simulation model should consider the generated datasets, as well as the assessment of the respective visualizations, which allow for a thorough examination of the operational details that occur throughout the simulation.

One of the advantages of the approach proposed and implemented in this paper lies in its independence from historical data. Unlike methods such as Machine Learning, or generative models (e.g., Large Language Models), which rely on large volumes of data, simulation enables the explicit modeling of processes based on operational rules and process logic. This is valuable in scenarios where data is unavailable or limited, as is the case in early-stage or conceptual systems. Moreover, simulation allows the configuration and evaluation of "what-if" scenarios.

In terms of future work, several main avenues can be pursued. First, key decision processes occurring during simulation runtime, such as charging scheduling, route planning, and fleet dispatching, could be refined through more sophisticated optimization strategies, replacing the current rule-based approaches with adaptive or learning-based mechanisms. Second, the integration of sustainability-oriented indicators, such as carbon emissions and energy efficiency metrics, would broaden the analytical scope of the framework beyond operational performance. Third, the data generated by the simulation should be exploited by machine learning applications, enabling the training of predictive models for demand forecasting, anomaly detection, or route optimization. Fourth, while the simulation produces reliable data, it still warrants the investigation of more flexible generative approaches, since any modification requires explicit rule re-implementation. Finally, the methodology proposed in this paper should also be tested in other cases that require synthetic data that can be generated by simulating process models, e.g., logistics or healthcare.

Acknowledgements. This work has been supported by the European Union under the NextGenerationEU, through a grant of the Portugal's Recovery and Resilience Plan (RRP) Partnership Agreement, within the scope of the project BE.NEUTRAL – Agenda da Mobilidade para a neutralidade carbónica das cidades (Project ref. nr. 35 - C644874240-00000016).

References

1. An, G., Cockrell, C.: Generating synthetic multidimensional molecular time series data for machine learning: considerations. Front. Syst. Biol. **3**, 1188009 (2023)
2. Andrews, R., van Dun, C.G., Wynn, M.T., Kratsch, W., Röglinger, M., ter Hofstede, A.H.: Quality-informed semi-automated event log generation for process mining. Decis. Support Syst. **132**, 113265 (2020)
3. Banks, J.: Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice. Wiley (1998)

4. Borovits, N., Bardelloni, G., Hashemi, H., Tulsiani, M., Tamburri, D.A., van den Heuvel, W.J.: Addressing data scarcity with synthetic data: a secure and GDPR-compliant cloud-based platform. ACM Trans. Softw. Eng. Methodol. (2025)
5. Chen, Y., Kempton, D.J., Ahmadzadeh, A., Wen, J., Ji, A., Angryk, R.A.: CGAN-based synthetic multivariate time-series generation: a solution to data scarcity in solar flare forecasting. Neural Comput. Appl. **34**(16), 13339–13353 (2022)
6. Li, K., Yang, S., Sullivan, T.M., Burd, R.S., Marsic, I.: ProcessGAN: generating privacy-preserving time-aware process data with conditional generative adversarial nets. ACM Trans. Knowl. Discov. Data **18**(9), 1–31 (2024)
7. Riess, M.: SynBPS: a parametric simulation framework for the generation of event-log data. SIMULATION **100**(8), 849–870 (2024)
8. Skydanienko, V., Di Francescomarino, C., Ghidini, C., Maggi, F.M.: A tool for generating event logs from multi-perspective declare models. In: Proceedings of the Dissertation Award, Demonstration, and Industrial Track at BPM 2018, vol. 2196, pp. 111–115 (2018)
9. Sousa, V., Vieira, C., Lavalle, A., Vieira, A., Santos, M.Y.: Integrating business and process analytics for enhanced data-driven decision-making. In: 43rd International Conference on Conceptual Modeling: ER Forum, Special Topics, Posters and Demos, ER-Companion 2024 (2024)
10. Taleb, M., Khalid, R., Ramli, R., Nawawi, M.K.M.: An integrated approach of discrete event simulation and a non-radial super efficiency data envelopment analysis for performance evaluation of an emergency department. Expert Syst. Appl. **220**, 119653 (2023)
11. Vieira, A., Dias, L., Pereira, G., Oliveira, J., Carvalho, M., Martins, P.: Automatic simulation models generation of warehouses with milk runs and pickers. In: 28th European Modeling and Simulation Symposium, EMSS, pp. 231–241 (2016)
12. Vieira, A., Dias, L., Pereira, G., Oliveira, J.: Comparison of Simio and arena simulation tools. In: 12th Annual Industrial Simulation Conference, ISC, pp. 5–13 (2014)
13. Vieira, A.A., Figueira, J.R., Fragoso, R.: A multi-objective simulation-based decision support tool for wine supply chain design and risk management under sustainability goals. Expert Syst. Appl. **232**, 120757 (2023)

Survival Analysis of HIV Seroconversion Using Interval-Censored Data: Nonparametric and Regression Approaches

Rui Alves[1], Filipa Pinto[1], Aurélio Sidumo[1], Carla Moreira[2], Luís Meira-Machado[1(✉)], Paula Meireles[3], Miguel Rocha[3,4], and Maria João Novais[3]

[1] Centre of Mathematics, University of Minho, Campus de Azurém, 4800-058 Guimarães, Portugal
lmachado@math.uminho.pt
[2] RISE-Health, Faculty of Medicine, University of Porto, Porto, Portugal
[3] EPIUnit ITR, Instituto de Saúde Pública da Universidade do Porto, Rua das Taipas 135, 4050-600 Porto, Portugal
[4] GAT – Grupo de Ativistas em Tratamentos, Lisbon, Portugal

Abstract. In prospective HIV cohorts, the exact time of seroconversion is rarely observed. Individuals are tested at irregular follow-up visits, so infection is only known to occur between the last negative and the first positive test. This leads to interval-censored event times, while non-seroconverters are subject to right censoring. Standard survival methods that assume exact or right-censored event times may therefore produce biased estimates.

In this paper, we analyse time to HIV seroconversion using methods specifically designed for interval-censored data. We first describe the Lisbon MSM cohort and its censoring structure. We then estimate the survival function using Turnbull's nonparametric maximum likelihood estimator and compare it with Kaplan–Meier estimators based on left, right and midpoint imputation. Group comparisons are conducted using tests adapted to interval-censored data. Finally, we fit parametric accelerated failure time models and proportional hazards models that explicitly account for interval censoring.

The results show that ad hoc imputation may lead to different survival estimates and that education level is the only covariate showing a statistically significant association with time to seroconversion in this application. In particular, higher educational attainment was associated with longer estimated times to HIV seroconversion under the AFT framework. These findings highlight the importance of using appropriate survival methods in HIV incidence studies based on periodic testing.

Keywords: HIV seroconversion · Interval-censored data · Survival analysis · Turnbull estimator · Imputation methods

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 344–357, 2027.
https://doi.org/10.1007/978-3-032-30518-3_22

1 Introduction

Accurate estimation of the time to HIV seroconversion is a key component in the epidemiological analysis of prospective cohorts and in the assessment of risk factors associated with infection. In many observational studies, individuals are tested for HIV at discrete and often irregular time points. As a result, seroconversion is not observed exactly, but is only known to have occurred between the last negative and the first positive HIV test. This gives rise to interval-censored event times, while individuals who remain HIV-negative throughout follow-up are right-censored.

In practice, survival analyses in HIV research often rely on methods developed for exact or right-censored event times, such as the Kaplan–Meier estimator and the Cox proportional hazards model. These methods are sometimes combined with ad hoc imputation of the event time, for example by taking the midpoint between the last negative and the first positive test. Although simple to implement, such approaches ignore the uncertainty about the true infection time and may affect the estimation of survival probabilities and covariate effects, particularly when the intervals between visits are wide or highly variable.

The statistical literature on interval-censored data provides a rigorous framework to address this problem. The nonparametric maximum likelihood estimator of the survival function for interval-censored data was introduced by Turnbull [1] and remains the standard reference in this setting. Extensions of classical inference procedures, including comparison of survival curves and regression modelling, have also been developed, covering both proportional hazards models [2,3] and accelerated failure time models [4,5]. Detailed discussions of these methods can be found in [6,7].

From an epidemiological perspective, the Lisbon cohort of men who have sex with men (MSM), established at a community-based voluntary counselling and testing centre, provides a relevant example of a prospective study based on repeated and irregular HIV testing [8,9]. This setting is particularly suited to illustrating the practical consequences of analysing interval-censored event times using either appropriate likelihood-based methods or simpler imputation-based alternatives.

The main objectives of this paper are as follows. First, we describe the data structure and the construction of the censoring intervals in the Lisbon MSM cohort. Second, we estimate the survival function using Turnbull's estimator and compare it with Kaplan–Meier estimators obtained after left, right and midpoint imputation. Third, we compare survival curves across levels of selected categorical covariates using methods adapted to interval-censored data. Finally, we assess the effect of covariates through regression models that explicitly account for interval censoring.

2 Data and Censoring Structure

2.1 Study Cohort and Analytical Sample

The data arise from the Lisbon cohort of men who have sex with men (MSM), a prospective observational study conducted at a community-based voluntary HIV counselling and testing centre in Lisbon, Portugal [8,9]. The cohort includes individuals who were initially HIV-negative and were subsequently followed through repeated, irregularly spaced testing visits.

The initial dataset comprised 11,389 individuals, among whom 521 had at least one positive HIV test result. Of these, 346 individuals were already HIV-positive at their first visit and were therefore excluded from the survival analysis, since the time to seroconversion could not be determined. The remaining 175 individuals had a positive HIV test after baseline.

Among the 11,043 individuals who were HIV-negative at the first visit, 5,433 had only a baseline observation and no subsequent HIV test. These individuals were excluded from the survival analysis, since no follow-up interval could be defined and they would otherwise correspond to observations of the form $(0, \infty)$, providing no information on the timing of seroconversion beyond baseline negativity.

After excluding these individuals, 5,610 participants had at least two testing visits. Eligibility criteria were then applied, restricting the analysis to participants aged 18 years or older, of male gender, with a valid testing record, and HIV-negative at the first visit. The final analytical sample included 5,106 participants, among whom 171 experienced HIV seroconversion during follow-up.

Figure 1 summarizes the selection process of the analytical sample.

2.2 Censoring Structure

Let T_i denote the unobserved time to HIV seroconversion for individual i, measured from the date of the first negative test. For individuals who seroconverted during follow-up, T_i is not observed exactly. Instead, it is only known to belong to an interval $(L_i, R_i]$, where L_i is the time of the last negative test and R_i is the time of the first positive test, with $0 \leq L_i < R_i < \infty$.

For individuals who remained HIV-negative throughout the observation period, the event time is right-censored. In this case, L_i is the time of the last negative test and $R_i = \infty$, so that $T_i > L_i$. Thus, the observed information for each participant can be represented by the interval $(L_i, R_i]$, with R_i possibly infinite.

Participants with only a baseline negative test and no subsequent follow-up were not included in the survival analysis, since no censoring interval could be defined beyond the baseline observation.

The visit process was irregular, reflecting the voluntary nature of testing and follow-up. Consequently, the lengths of the censoring intervals varied across participants. This feature is central to the analysis, since replacing each interval by a single imputed time ignores part of the uncertainty about the true seroconversion time.

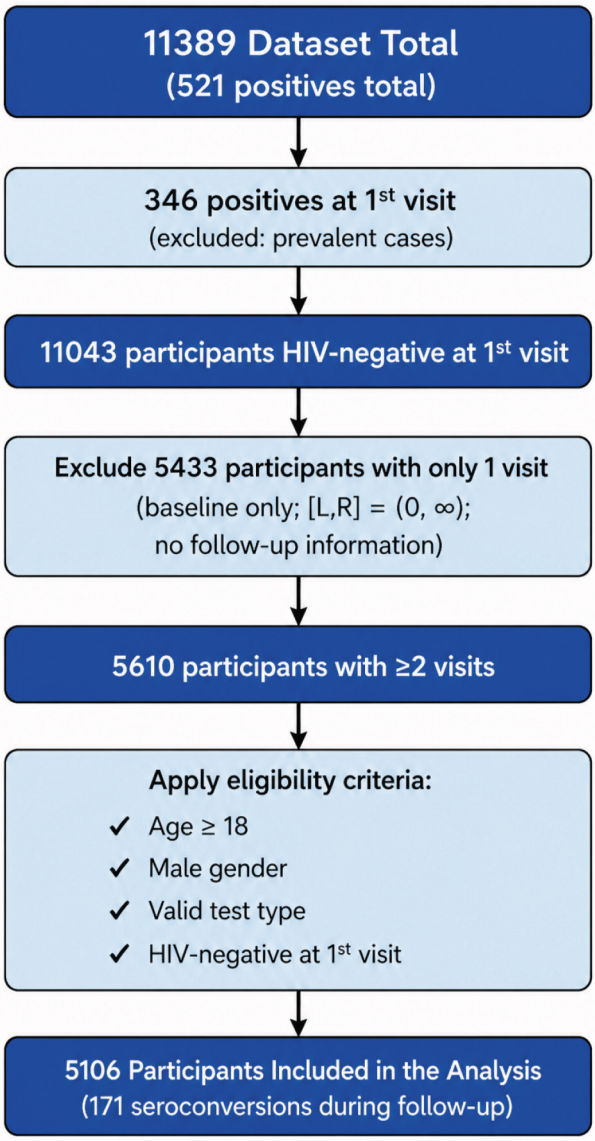

Fig. 1. Flowchart of the selection of the analytical sample for the analysis of interval-censored HIV seroconversion times.

The analysis assumes that the observation process, namely the timing of HIV testing visits, is independent of the underlying time to seroconversion, conditional on the observed covariates. This is a standard working assumption in analyses of interval-censored data collected through periodic observation, but its plausibility is revisited in the Discussion.

2.3 Covariates and Missing Data

The available covariates include demographic and behavioural variables collected through questionnaires at baseline and, when appropriate, during follow-up. These include age, sex assigned at birth, gender identity, previous HIV testing, substance use, sex with casual partners, condom use during anal sex, education level and country of birth. The dataset also includes the dates of the different HIV testing visits, which were used to construct the intervals $(L_i, R_i]$.

Table 1 summarizes the main sociodemographic and behavioural characteristics of the analytical sample. For each variable, the table reports the number and percentage of participants in each observed category, together with the number and percentage of seroconversion events. Missing or unanswered responses are shown explicitly whenever present, since the corresponding denominators differ across variables. Percentages for the categories are calculated using the total analytical sample as denominator, unless otherwise stated.

Some variables contain a substantial proportion of missing or unanswered responses, mainly because not all questions were available or completed for all participants. This is particularly evident for variables such as sex assigned at birth, substance use and country of birth. For this reason, the descriptive table should be interpreted as a summary of the available questionnaire information rather than as a complete profile for all variables.

Although the inclusion criteria restricted the analysis to individuals of male sex, the variable "sex assigned at birth" was not available for all participants in the questionnaire data. Similarly, gender identity was recorded as "man" for almost all participants with available information, but a small number of missing responses remained. These missing values explain the apparent discrepancy between the eligibility criteria and the totals displayed in Table 1.

3 Methods

3.1 Nonparametric Estimation

Let T denote the time to HIV seroconversion, with survival function $S(t) = P(T > t)$. In the presence of interval censoring, each observation is of the form $(L_i, R_i]$, indicating that the event occurred between two observation times.

The nonparametric maximum likelihood estimator (NPMLE) of S in this setting is given by Turnbull's estimator [1]. This estimator generalizes the Kaplan–Meier estimator to interval-censored data and is obtained by maximizing the likelihood

$$\prod_{i=1}^{n} \left[F(R_i) - F(L_i) \right], \tag{1}$$

for interval-censored observations, with contribution $S(L_i)$ for right-censored individuals, where $F(t) = 1 - S(t)$.

The resulting estimator is a step function with jumps at a subset of the interval endpoints and is typically computed using the self-consistency algorithm.

Table 1. Sociodemographic and behavioural characteristics of the analytical sample

Variable	n (%)	n Events (%)
Sex assigned at birth		
Male	603 (11.8)	5 (0.8)
Missing/unanswered	4503 (88.2)	166 (3.7)
Gender identity		
Man	5071 (99.3)	170 (3.4)
Missing/unanswered	35 (0.7)	1 (2.9)
Previous HIV test		
Yes	4148 (81.2)	138 (3.3)
No	889 (17.4)	33 (3.7)
Missing/unanswered	45 (0.9)	0 (0.0)
Substance use		
No	1709 (33.5)	107 (6.3)
Yes	23 (0.5)	1 (4.3)
Does not know	176 (3.4)	19 (10.8)
Missing/unanswered	3193 (62.5)	43 (1.3)
Sex with casual partners		
Yes	4358 (85.4)	153 (3.5)
No	686 (13.4)	16 (2.3)
Missing/unanswered	4 (0.1)	0 (0.0)
Condom use (anal sex)		
Always	737 (14.4)	27 (3.7)
Almost always	681 (13.3)	18 (2.6)
Never	728 (14.3)	28 (3.8)
Rarely	449 (8.8)	15 (3.3)
Sometimes	439 (8.6)	13 (3.0)
Missing/unanswered	2064 (40.4)	69 (3.3)
Education level		
Bachelor's degree	2028 (39.7)	66 (3.3)
Secondary school	1424 (27.9)	66 (4.6)
Master's/PhD	1071 (21.0)	19 (1.8)
Vocational training	337 (6.6)	15 (4.5)
$\leq$ Basic school	165 (3.2)	4 (2.4)
Missing/unanswered	81 (1.6)	1 (1.2)
Country of birth		
Portugal	1494 (29.3)	95 (6.4)
Brazil	187 (3.7)	15 (8.0)
Other European countries	127 (2.5)	7 (5.5)
African countries	71 (1.4)	8 (11.3)
Other countries	37 (0.7)	0 (0.0)
Missing/unanswered	3190 (62.5)	46 (1.4)

3.2 Kaplan–Meier Estimators with Imputation

For comparison purposes, we also consider Kaplan–Meier estimators obtained after imputing a single event time within each censoring interval. Three imputation strategies are used: left imputation ($T_i = L_i$), right imputation ($T_i = R_i$) and midpoint imputation ($T_i = (L_i + R_i)/2$).

These approaches transform interval-censored observations into pseudo-exact event times, allowing the use of standard Kaplan–Meier estimation. However, they ignore the uncertainty about the true event time and may introduce bias in the estimation of survival probabilities.

3.3 Comparison of Survival Curves

To assess differences in survival distributions across levels of categorical covariates, we employ log-rank type tests adapted to interval-censored data. These tests are based on score statistics derived from the proportional hazards model or on permutation procedures using the NPMLE of the survival function [3,6].

The null hypothesis corresponds to equality of survival functions across groups. These methods extend the classical log-rank test by incorporating interval-censored likelihood contributions instead of exact event times.

3.4 Regression Models

To evaluate the effect of covariates on the time to HIV seroconversion, we considered both accelerated failure time (AFT) and proportional hazards (PH) models adapted to interval-censored data, with particular emphasis on the parametric AFT approach.

The main regression analysis was based on parametric accelerated failure time (AFT) models fitted using the `ic_par()` function from the `icenReg` package. A semiparametric proportional hazards model for interval-censored data was also considered using the `ic_sp()` function as a complementary analysis.

Under the AFT parametrization used by `icenReg`, the survival function is expressed as

$$S(t \mid \mathbf{x}_i) = S_0\left(t \exp(\mathbf{x}_i^\top \boldsymbol{\beta})\right), \tag{2}$$

where S_0 is the baseline survival function and $\mathbf{x}_i$ is the vector of covariates. With this parametrization, negative coefficients correspond to longer survival times, whereas positive coefficients correspond to shorter survival times. Equivalently, the time ratio associated with a covariate is given by $\exp(-\beta)$.

Likelihood-based inference is performed by integrating the corresponding density over the interval $(L_i, R_i]$, thus explicitly accounting for interval and right censoring.

In the PH framework, the hazard function is specified as

$$\lambda(t \mid \mathbf{x}_i) = \lambda_0(t) \exp(\mathbf{x}_i^\top \boldsymbol{\beta}) \tag{3}$$

where $\lambda_0(t)$ is an unspecified baseline hazard. Estimation is carried out using likelihood-based methods adapted to interval-censored data [2,3].

3.5 Implementation Details

All analyses were performed in the R statistical environment [10]. Nonparametric estimation and hypothesis testing for interval-censored data were carried out using the `interval` package, while regression models were implemented using the `icenReg` package [11]. Standard survival methods were obtained using the `survival` package.

Several parametric distributions were explored for the AFT models, including log-normal, log-logistic, Weibull and exponential distributions. The log-normal specification provided the best fit according to the Akaike Information Criterion (AIC) and was therefore retained for the main regression analysis.

To avoid degenerate intervals when $L_i = 0$, a small positive constant ($\varepsilon = 10^{-4}$) was added when necessary to ensure valid interval limits. For categorical variables, categories corresponding to missing responses were excluded from regression analyses and appropriate reference categories were defined.

4 Results

4.1 Descriptive Characteristics

The main sociodemographic and behavioural characteristics of the analytical sample are summarized in Table 1. The sample includes 5,106 participants, among whom 171 HIV seroconversion events were observed during follow-up.

Overall, most participants reported previous HIV testing and a high proportion indicated having sex with casual partners. Condom use during anal sex showed considerable variability across individuals. Education levels were heterogeneous, with a large proportion of participants reporting secondary or higher education. Most participants were born in Portugal, although a non-negligible proportion originated from other countries.

As discussed in Sect. 2, several variables include missing or unanswered responses, which should be taken into account when interpreting the distributions presented in Table 1.

4.2 Nonparametric Survival Estimates

Figure 2 presents the estimated survival function for time to HIV seroconversion obtained using Turnbull's estimator. The survival probability remains close to one over most of the follow-up period, reflecting the relatively low incidence of seroconversion in the cohort. A gradual decline is observed as follow-up time increases, corresponding to the occurrence of new infections.

For comparison purposes, Fig. 3 displays the Kaplan–Meier survival curves obtained under the three imputation strategies in a single graph. Systematic differences are observed across the three approaches. As expected, left imputation produces lower survival estimates, reflecting the fact that events are assigned to the earliest possible time within each interval.

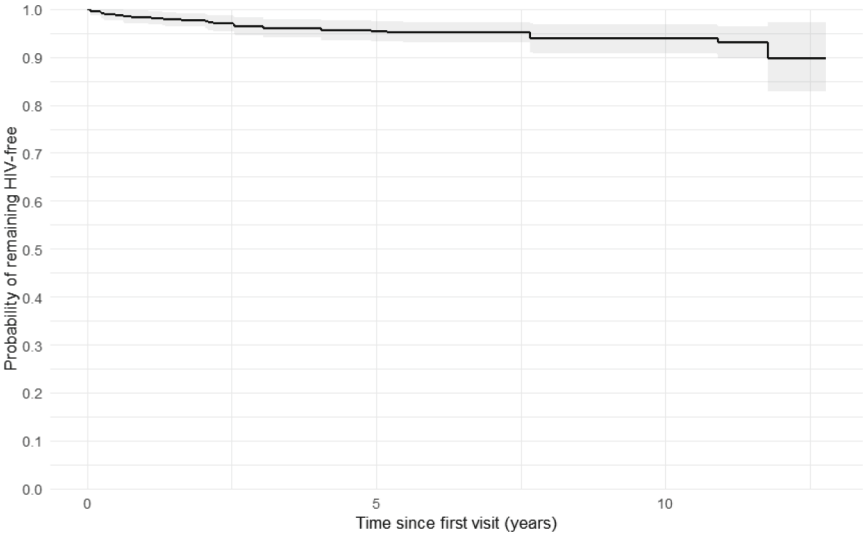

Fig. 2. Estimated survival function for time to HIV seroconversion using Turnbull's estimator.

In this dataset, differences between the imputation strategies reflect the distribution of censoring intervals. When the last negative test occurs at baseline, left imputation assigns the event time close to zero, whereas midpoint and right imputation place the event later within the interval.

In contrast, midpoint and right imputation yield very similar survival estimates. This reflects the fact that, for a large proportion of individuals, the censoring intervals are relatively short or dominated by the timing of the first positive test, so that the midpoint and right endpoints are close. For clarity, the figure is displayed using a restricted vertical scale (0.75–1), which allows these small differences to be more easily visualised.

These differences illustrate the impact of replacing each censoring interval by a single imputed value. In the present data, the effect is not uniform across imputation strategies: left imputation leads to a noticeable downward shift of the survival curve, while midpoint and right imputation produce very similar estimates. This pattern reflects the distribution of censoring intervals in the cohort and highlights how imputation-based approaches may distort the timing of events.

4.3 Comparison Across Covariates

Differences in survival distributions across levels of categorical covariates were assessed using log-rank type tests adapted to interval-censored data.

Among the variables considered, education level was the only covariate for which statistically significant differences were detected (test statistic = 16.37,

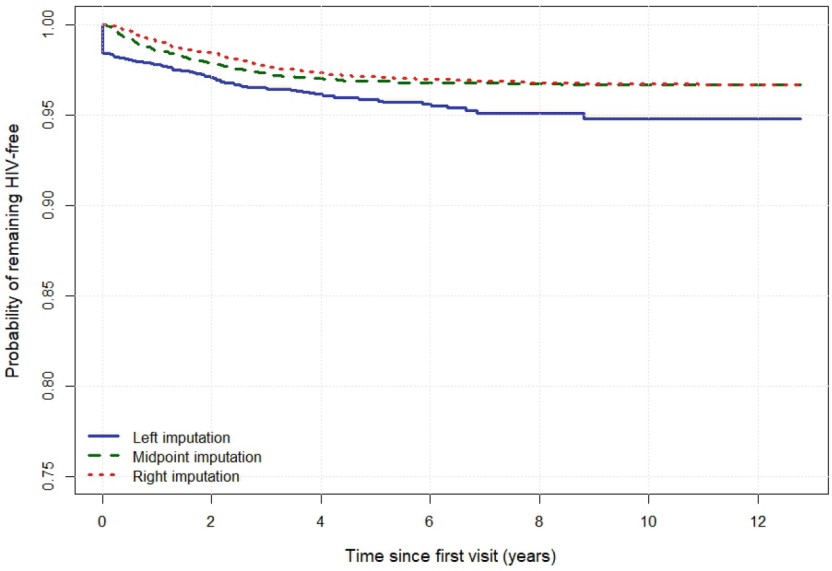

Fig. 3. Kaplan–Meier survival curves obtained under left, midpoint and right imputation. The vertical axis is restricted to the interval (0.75–1) to improve visualisation of the differences.

$p = 0.0026$). No significant differences were observed for condom use, sex with casual partners, substance use or country of birth.

Figures 4 and 5 show the estimated survival curves by education level obtained using Turnbull's nonparametric estimator. The first figure presents the full probability scale, while the second provides a zoomed view to better highlight the differences between groups. Individuals with higher educational attainment tend to present slightly higher probabilities of remaining HIV-negative over time, whereas lower education levels are associated with a somewhat faster decline in the survival function.

Although the absolute differences between curves are relatively small, the statistical test suggests that they are unlikely to be due to random variation.

4.4 Regression Analysis

The results of the univariable log-normal accelerated failure time (AFT) models are presented in Table 2. Most covariates did not show statistically significant associations with time to HIV seroconversion.

Age, condom use, substance use, previous HIV testing and country of birth were not significantly associated with time to HIV seroconversion.

In contrast, education level showed a statistically significant association with time to HIV seroconversion. Under the parametrization used in the AFT model, time ratios greater than one correspond to longer times to seroconversion. Individuals with a bachelor's degree and those with a master's or doctoral degree

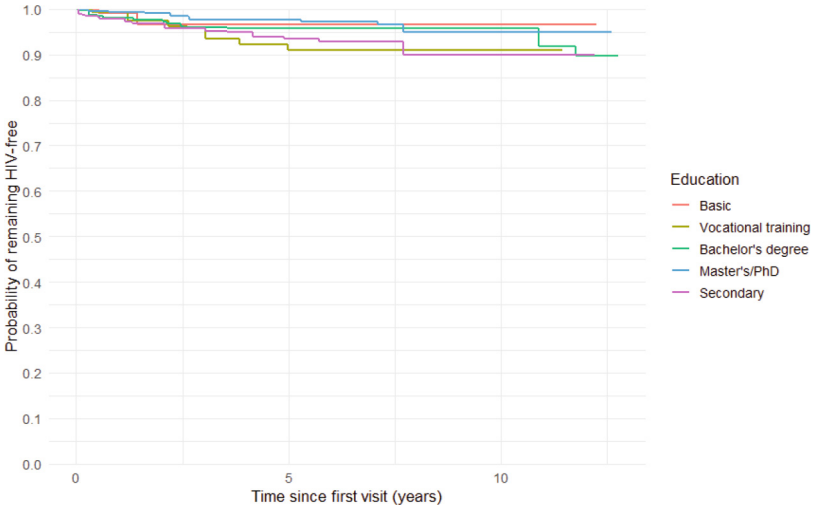

Fig. 4. Estimated survival curves for time to HIV seroconversion according to education level on the full probability scale (0–1).

presented longer estimated times to HIV seroconversion compared with individuals with secondary education.

In particular, participants with a master's or doctoral degree showed an estimated time ratio of 2.56 ($p = 0.0003$), suggesting substantially longer times to HIV seroconversion relative to the reference group. Similarly, individuals with a bachelor's degree presented a time ratio of 1.48 ($p = 0.0246$).

Among the parametric models considered, the log-normal distribution provided the best fit according to the AIC criterion, although differences with other distributions were relatively small. The corresponding hazard ratios obtained from the proportional hazards model were consistent with the direction of the effects observed under the AFT specification.

These findings are consistent with the results obtained from the comparison of survival curves.

5 Discussion

This study examined time to HIV seroconversion in a prospective cohort of men who have sex with men, explicitly accounting for the interval-censored nature of the data arising from periodic and irregular HIV testing. In this setting, the exact infection time is not observed and is only known to lie between consecutive visits, making standard survival methods based on exact or right-censored times potentially inadequate.

The comparison between Turnbull's nonparametric maximum likelihood estimator and Kaplan–Meier estimators based on imputed event times highlights the impact of ignoring interval censoring. Although the differences between

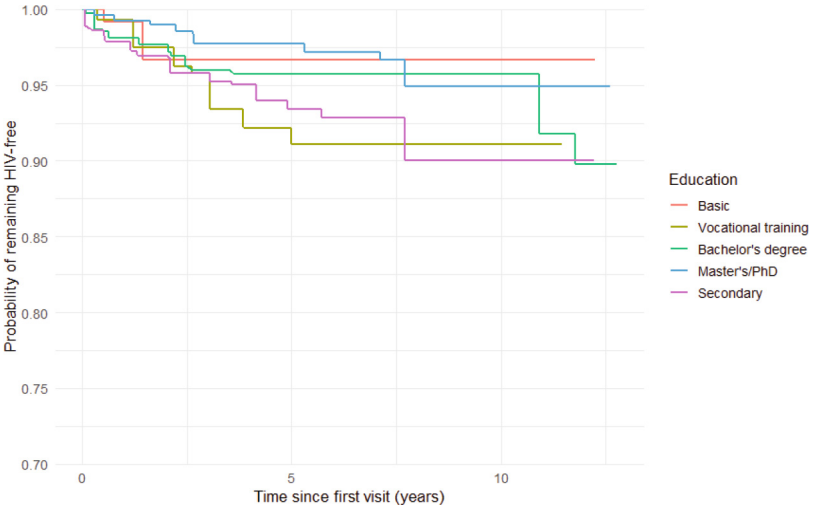

Fig. 5. Estimated survival curves for time to HIV seroconversion according to education level with a restricted vertical scale (0.7–1) to highlight differences between groups.

approaches are moderate, they are systematic: left imputation tends to underestimate survival probabilities, while right imputation may overestimate them. These results illustrate how simplifying assumptions about event timing may affect survival estimates.

The comparison of survival curves across covariates indicated that education level was the only variable showing statistically significant differences. Individuals with higher educational attainment tended to present higher probabilities of remaining HIV-negative over time. This finding was supported by the AFT regression analysis, where higher education levels were associated with longer estimated times to HIV seroconversion.

This association should, however, be interpreted with caution. Education level may act as a proxy for other factors not fully captured in the available data, such as risk perception, access to prevention strategies, testing behaviour, or healthcare access. As a result, the observed relationship does not necessarily imply a direct causal effect.

Participants with only a baseline HIV-negative test and no subsequent follow-up were excluded from the survival analysis. Although this is necessary because no follow-up interval can be defined for these individuals, it may affect the generalizability of the results if participants who did not return for testing differ systematically from those with repeated visits.

Several limitations should be considered. First, the analysis assumes that the timing of HIV testing visits is independent of the underlying time to seroconversion, conditional on observed covariates. In a voluntary testing setting, this assumption may be violated, since individuals may seek testing in response to

Table 2. Univariable log-normal AFT models for time to HIV seroconversion

Variable	Category	Estimate	TR	p-value
Age	—	-0.0086	1.009	0.3335
Condom use	Sometimes	−0.0875	1.091	0.7962
	Never	0.1294	0.879	0.6324
	Almost always	−0.2493	1.283	0.4136
	Rarely	−0.0252	1.026	0.9377
Casual partners	Yes	0.3803	0.684	0.1478
Substance use	Yes	−0.2925	1.340	0.7695
Previous HIV test	Yes	−0.0825	1.086	0.6703
Education level	≤ Primary education	−0.5906	1.805	0.2496
	Vocational training	−0.0083	1.008	0.9767
	Bachelor's degree	−0.3906	1.478	**0.0246**
	Master's/PhD	−0.9410	2.563	**0.0003**
Country of birth	Brazil	0.4297	0.651	0.1220
	Other European countries	−0.0122	1.012	0.9751
	African countries	0.6223	0.537	0.0919
	Other countries	−6.1730	479.6	0.2488

TR = time ratio = exp(−Estimate)
Reference categories: condom use = Always; casual partners = No; substance use = No; previous HIV test = No; education level = Secondary education;
country of birth = Portugal. TR > 1 indicates longer times to seroconversion.

perceived exposure or symptoms. Second, some variables contain a substantial proportion of missing or unanswered responses, which may affect the interpretation of descriptive and regression results. The small adjustment applied to zero left limits ($\varepsilon = 10^{-4}$) was only a technical device to avoid degenerate interval limits. Given its very small magnitude relative to the follow-up times, it is not expected to affect the substantive results. Finally, the relatively low number of seroconversion events limits the power to detect associations for some covariates.

Despite these limitations, the results illustrate the importance of using statistical methods that properly account for interval censoring in the analysis of HIV incidence data.

6 Conclusion

This work illustrates the practical relevance of interval-censored survival methods in the analysis of HIV seroconversion data collected through irregular testing schedules. Methods such as Turnbull's estimator and regression models for interval-censored data provide a principled alternative to approaches based on ad hoc imputation of event times.

In the Lisbon MSM cohort, education level was the only covariate showing a statistically significant association with time to seroconversion. In particular, individuals with higher educational attainment presented longer estimated times to HIV seroconversion under the AFT modelling framework, while most behavioural variables considered did not show statistically significant effects. More broadly, the results support the use of statistical methods that explicitly account for interval censoring in epidemiological studies where event times are only observed between visits.

Acknowledgements. This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the UID/00013: Centro de Matemática da Universidade do Minho (CMAT/UM) Program Contract, and the project reference 2023.14897.PEX (DOI: 10.54499/2023.14897.PEX).

References

1. Turnbull, B.W.: The empirical distribution for interval-censored data. J. Am. Stat. Assoc. **71**, 253–264 (1976)
2. Finkelstein, D.M.: A proportional hazards model for interval-censored failure time data. Biometrics **42**, 845–854 (1986)
3. Pan, W.: A multiple imputation approach to Cox regression with interval-censored data. Biometrics **56**, 199–203 (2000)
4. Lindsey, J.C., Ryan, L.M.: Tutorial in biostatistics: methods for interval-censored data. Stat. Med. **17**, 219–238 (1998)
5. Bogaerts, K., Komarek, A., Lesaffre, E.: Survival analysis with interval-censored data: a practical approach with examples in R. CRC Press, SAS and BUGS (2017)
6. Sun, J.: The Statistical Analysis of Interval-Censored Failure Time Data. Springer (2006)
7. Kleinbaum, D.G., Klein, M.: Survival Analysis: A Self-Learning Text, 3rd edn. Springer (2012)
8. Meireles, P., Lucas, R., Martins, A., et al.: The Lisbon cohort of men who have sex with men. BMJ Open **5**, e007220 (2015)
9. Meireles, P., Lucas, R., Carvalho, C., et al.: Incident risk factors as predictors of HIV seroconversion in the Lisbon cohort of men who have sex with men: first results, 2011–2014. Eurosurveillance **20**(14), 21091 (2015)
10. R Core Team: R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria (2024). https://www.R-project.org
11. Anderson-Bergman, C.: icenReg: regression models for interval censored data in R. J. Stat. Softw. **81**(12), 1–23 (2017)



Experimental Limits of Machine Learning on Near-Random Data: A Case Study of the Romanian Lottery System

Alexandru Dinu, Sara-Gabriela Modan, and Mădălin Frunzete(✉)

Faculty of Electronics, Telecommunications and Information Technology, National University of Science and Technology POLITEHNICA Bucharest, Bucharest, Romania
{alexandru.dinu2408,madalin.frunzete}@upb.ro, sara.gabriela.modan@stud.etti.upb.ro

Abstract. Lottery has attracted people's attention for a long time and the fascination and curiosity on how to beat the system are present today as they were at the beginning. This paper investigates if modern machine learning methods can be useful in extracting information and make predictions based on the Romanian 6/49 lottery historical data (more than 30 years of data and 2510 draws). First of all, we performed a statistical analysis based on the draw sequences and show that it is not statistically different compared to a uniform random process in terms of distribution, independence of data, and informational entropy. Secondly, based on the above-mentioned learnings, we evaluate four statistical and machine learning-based models: a frequency baseline, a Markov model, a multilayer perceptron, and a random forest classifier model. The training and testing data was carefully separated to avoid any leakage. The performance of these models is assessed against the expectation/hypothesis that they will deliver a performance close to random guessing. All models show almost identical hit rates, without any statistically significant difference compared to chance. The findings provide empirical evidence that when the data has information entropy closed to maximum (uniform distribution), no machine learning, as complex as it may be, is not able to achieve any systematic and predictable gains in the long run.

Keywords: Pseudo-random sequences · Lottery prediction · Machine learning limitations · Shannon entropy · Markov chains · Random Forest · Neural networks · Information theory

1 Introduction

Can artificial intelligence/machine learning methods help predict the lottery results? This is not a new question but still a very vivid and interesting one, not only for the general public but also for the researchers in the academia. Practically speaking, there is an entire (small) industry which focuses entirely on

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 358–371, 2027.
https://doi.org/10.1007/978-3-032-30518-3_23

finding a good-enough algorithm to "guess" next weeks lottery numbers (many of them, if not all). From a theoretical point of view, information theory can provide the answer: for a random process, no method can perform better than chance, no matter how elaborate and intricate that method is. The huge gap between the theoretical and the practical perspective is not usually closed empirically and that could be a reason why some people still struggle and work tirelessly to create a lottery-winning algorithm. Most published work either reports minor predictive gains without verifying in the first place that the data is truly random, or the alternative is that the scientific work dismisses the problem without correlating the rejection with the statistical properties of the specific dataset which is studied [1,2].

Romania's lottery system, Loto 6/49, is run by Loteria Romana and has been doing so continuously since 1993. Each draw selects 6 numbers without replacement from $\{1, 2, \ldots, 49\}$, producing a combinatorial space of $\binom{49}{6} = 13\,983\,816$ possible results/sequences. Of course, based on the physical safeguard mechanisms and procedures which are in place, one would expect that the draws follow a uniform distribution. Nonetheless, this assumption should be validated based on data and should not be taken for granted.

Previous work has investigated strategies based on frequency analysis [1], neural network predictors [2], first-order Markov models [3,4], and more recently LSTM-based approaches [5]. The results obtained across these studies are inconsistent, and sometimes positive findings are attributed to overfitting of test data which is not correctly separated in time [6]. Rigorous randomness test batteries such as the NIST Statistical Test Suite [7] and the Diehard tests [8] have been applied to lottery data before and generally support the hypothesis of randomness. However, these analyses are very rarely combined with a systematic machine learning evaluation applied on the same dataset.

This paper addresses that gap. We first verify, through multiple complementary statistical tests, that the Romania Loto 6/49 draw sequence is statistically indistinguishable from a uniform random process. We evaluate next four predictive models—a frequency baseline, a Markov chain, a Multilayer Perceptron, and a Random Forest classifier—against the theoretical random expectation, using strict temporal train/test separation to prevent data leakage. The central hypothesis is that no model will outperform random guessing, and we test this claim empirically rather than simply assuming it.

The current paper is organised as follows:

- Section 2 presents the statistical analysis of the lottery historical data
- Section 3 describes the experimental setup, machine learning models used for comparison, and results
- Section 4 discusses the broader implications and
- Section 5 concludes the paper and mentions the future investigation paths.

2 Statistical Analysis

The authors of this papers applied before statistical independence and uniformity testing via runs tests, autocorrelation analysis, and chi-square goodness-of-fit across diverse sets of sequential data, including chaotic dynamical systems and natural language corpora [9–12]. This provides a solid basis for the tests applied in this study.

The analyzed real dataset includes all official Romanian lottery system, Loto 6/49, winning sequences from August 8, 1993 to January 4, 2026, as published by Loteria Romana. Draw frequency increased from 1–2 per week in the early period to 2–3 per week after 2009. Years with fewer than 40 draws (1993 and 2026) are excluded from annual analyses, leaving 32 complete calendar years (1994–2025). After removing five draws which did not contain a complete sequence of 6 values, the cleaned dataset contains $n = 2510$ complete draws and 15060 (2510×6) individual number observations.

We establish the pseudo-random properties of the draw sequence before evaluating any predictive model, verifying three complementary properties: uniform distribution, near-maximum Shannon entropy, and serial independence of draws.

Uniform Distribution.
Figure 1 shows the observed experimental frequency of each number from 1 to 49, across all 2,510 draws, with counts ranging from 271 (number 31) to 354 (number 5) and a mean of 307.35. Each draw selects 6 numbers out of 49, every number has a theoretical probability of 6/49 of appearing in any given sequence. To verify that all 49 numbers are equally likely to be drawn—that is, that no number is systematically favoured or avoided—we apply the chi-square goodness-of-fit test:

$$\chi^2 = \sum_{i=1}^{49} \frac{(O_i - E_i)^2}{E_i} \tag{1}$$

where $E_i = N \cdot 6/49 = 307.35$. The result ($\chi^2 = 48.99$, $df = 48$, $p = 0.433$) provides no evidence against uniformity.

As an additional check, we look at the sum of the six drawn numbers for each draw. For example, if the numbers 3, 15, 22, 31, 40 and 47 are drawn, the sum is 158. Repeating this for all 2,510 draws gives us 2,510 such sums. The average of these sums, observed in the data, is 147.95, and they spread around this average with a standard deviation of 32.52—meaning most sums fall roughly between 115 and 180.

If the lottery system would be truly fair and all numbers are equally likely to be drawn, probability theory predicts that this average should be close to 150 and the spread close to 32.5. Our observed values of 147.95 and 32.52 match these predictions almost exactly, providing further confirmation that the draw process is uniform.

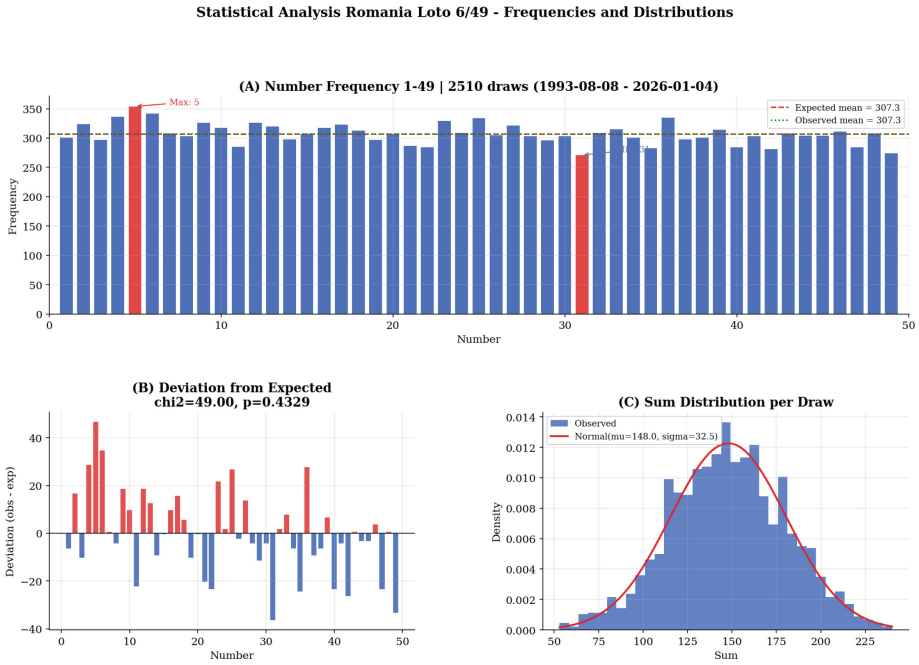

Fig. 1. (A) Observed frequency of each number (1–49) across 2510 draws (1994–2025). Red dashed line: expected frequency under uniformity (307.35); observed mean coincides at 307.35. Numbers 5 (maximum, 354 appearances) and 31 (minimum, 271 appearances) are annotated. (B) Deviation of each number' s frequency from the expected value of 307.35. (C) Distribution of draw sums with best-fitting normal curve.(Color figure online)

2.1 Shannon Entropy

Shannon entropy quantifies predictability directly. We compute the entropy based on 2:

$$H = -\sum_{i=1}^{49} \hat{p}_i \log_2 \hat{p}_i \tag{2}$$

and obtain $H = 5.6124$ bits compared to a theoretical maximum $H_{\max} = \log_2 49 = 5.6147$ bits, which gives an entropy ratio of 99.958%. Figure 2(D) confirms this is stable across all years, reducing the likelihood of temporal drift.

Independence of Draws.
A further important property to verify is that the draws are independent of one another—that is, the numbers drawn in one draw contain no information about the numbers that will appear in future draws. We test this in two complementary ways.

First, we apply the Wald-Wolfowitz runs test [13]. We take all 15,060 extracted numbers and convert them to a sequence of 0 s and 1 s: if a num-

ber is greater than 25 (the median) we write 1, otherwise 0. We then count the number of uninterrupted streaks of the same symbol—this concept is called *runs*. For example, in the sequence 1, 1, 0, 1, 0, 0 there are four runs: (1,1), (0), (1), (0,0). A run can be as short as a single value. If the sequence is truly random, the number of runs should be neither too high nor too low—too many runs would mean the sequence alternates too often, while too few would mean values cluster together more than expected. We observe $R = 7495$ runs against the expected value $\bar{R} = 7517.91$, computed from the counts of zeros and 1 s as $\bar{R} = 2n_0n_1/(n_0+n_1)+1$. This is a negligible difference ($Z = -0.374$, $p = 0.708$), consistent with a random sequence.

Second, we reshape all draws into a single sequence of 15,060 individual numbers in chronological order and measure whether any number tends to be correlated with the numbers appearing 1, 2, or up to 15 positions later in this sequence. This assessment is made by using the serial autocorrelation coefficient:

$$r_\ell = \frac{\sum_{j=1}^{n-\ell}(x_j - \bar{x})(x_{j+\ell} - \bar{x})}{\sum_{j=1}^{n}(x_j - \bar{x})^2} \tag{3}$$

All 15 coefficients depicted in Fig. 2 (A) are approximately zero—the largest is $|r_1| = 0.0085$, well within the 95% confidence bounds of $\pm 2/\sqrt{n} = \pm 0.0163$. This confirms that knowing which numbers appeared in past draws gives no useful information about future draws.

For each of the 49 numbers, we measure how many draws pass between two consecutive appearances—this is called the *gap* for that number. For example, if number 7 appears at draws 5, 12 and 31, its gaps are 7 and 19. Collecting all such gaps across all 49 numbers gives the gap sequence.

If the lottery is truly random and each number has an equal probability of $p = 6/49$ of being drawn at any given draw, then the waiting time between consecutive appearances of any number should follow a geometric distribution—that is, the gaps should decrease exponentially in frequency: short gaps are common, long gaps are rare. The observed gap sequence has mean $\bar{x}_g = 48.81$ and standard deviation $s_g = 45.32$, consistent with the theoretical expectation of $\lambda^{-1} = 49/6 \approx 8.17$ draws between appearances – Fig. 2 (C).

Next, a Kolmogorov-Smirnov test compares the observed gap distribution with the fitted exponential and yields $KS = 0.625$, $p < 0.001$, indicating a small but statistically detectable deviation. However, this deviation is fully expected and carries no predictive implications: since exactly 6 *distinct* numbers are drawn simultaneously at each draw, the same number cannot appear twice in the same draw. This introduces a slight regularity in the gap structure that causes a small deviation from a pure exponential—it is a mathematical consequence of the sampling mechanism and not a sign that the lottery system is predictable.

As a final check, we examine whether any two numbers tend to appear together in the same draw more often than would be expected by chance. For each of the $\binom{49}{2} = 1176$ possible pairs of numbers, we count how many times both numbers were drawn in the same draw—this is the co-occurrence count for

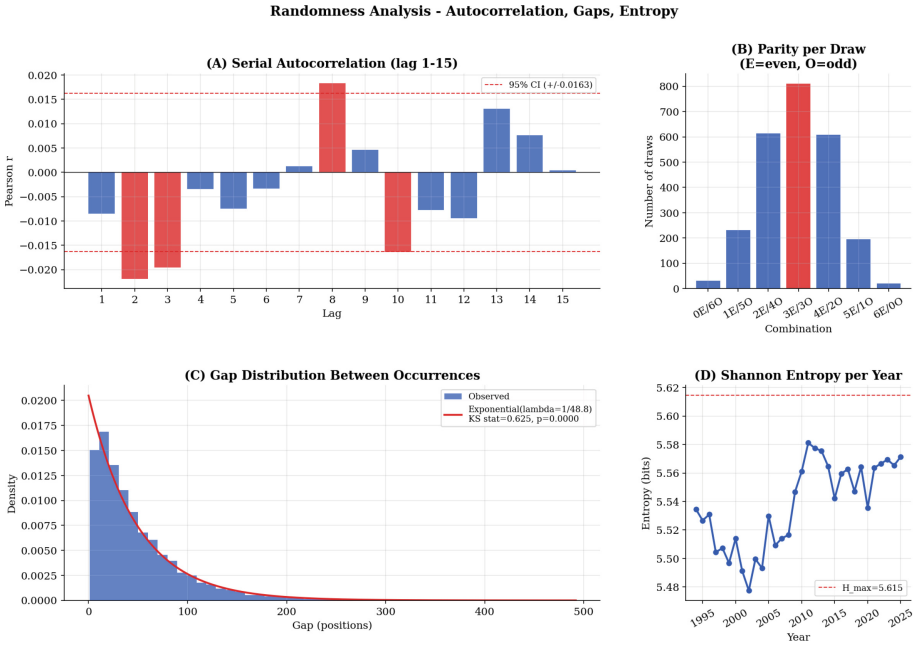

Fig. 2. Randomness analysis of the Romania Loto 6/49 draw sequence. (A) Serial autocorrelation for lags 1–15; dashed lines indicate 95% confidence bounds. (B) Parity distribution per draw; the most frequent combination is 3 even/3 odd (32.3%), consistent with combinatorial expectations. (C) Empirical gap distribution with best-fitting exponential curve. (D) Shannon entropy per year (1994–2025); dashed line indicates $H_{\max} = 5.615$ bits.

that pair. For example, if numbers 7 and 23 appeared together in 19 draws, their co-occurrence count is 19.

If the lottery would be fair and all numbers are independent of one another, each pair should appear together with equal probability at every draw. At each draw, exactly $\binom{6}{2} = 15$ pairs are formed among the 6 drawn numbers, out of a total of $\binom{49}{2} = 1176$ possible pairs. Over $N = 2510$ draws, the expected co-occurrence count for any specific pair is therefore:

$$E[C_{ij}] = N \cdot \frac{\binom{6}{2}}{\binom{49}{2}} = 2510 \cdot \frac{15}{1176} \approx 18.37 \tag{4}$$

The resulting 49×49 co-occurrence matrix is shown in Fig. 3. All observed counts fall close to 18.37, with no pair appearing together substantially more or less often than expected. No systematic pattern is visible—confirming that no number is statistically linked to any other number across the draw history.

Table 1 summarises the key statistical findings.

Figure 4 shows the maximum gap per number across all years; no systematic temporal pattern is visible, confirming stationarity.

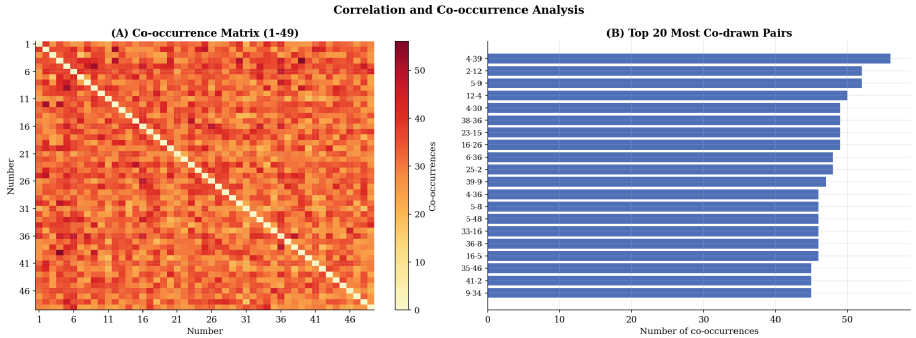

Fig. 3. Co-occurrence analysis. (A) Full 49×49 co-occurrence matrix; colour intensity indicates draw count per pair. No systematic pattern is visible. (B) Top 20 most frequently co-occurring pairs; counts are consistent with sampling variability around the expected value of approximately 18.4.

Table 1. Summary of statistical tests on Romania Loto 6/49 draw data (2,510 draws, 1994–2025).

Test	Null Hypothesis	Statistic	p-value	Decision
Chi-square	Uniform distribution	$\chi^2 = 48.99$	0.433	Not rejected
Shannon entropy	–	$H = 5.6124$ bits	–	99.96% of $H_{\max}$
Wald-Wolfowitz	Serial independence	$Z = -0.374$	0.708	Not rejected
Autocorr. ($\ell = 1$)	No linear dependence	$r = -0.0085$	–	Within 95% CI

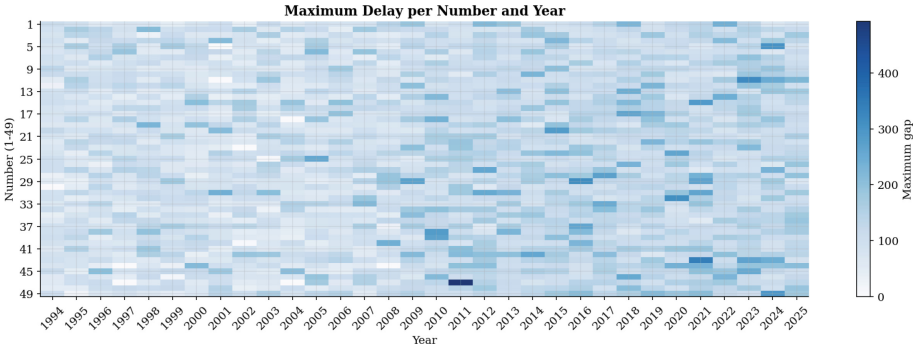

Fig. 4. Maximum gap per number and per year (1994–2025). Colour intensity represents the longest consecutive absence of each number within a given year. No systematic temporal pattern is visible, confirming stationarity of the process.

3 Experimental Setup and Results

Now that we established that the draw sequence is statistically the same as a uniform random process, we turn our attention to the central question of this

paper: can any machine learning model extract useful predictive structure from this data? In this section we describe how we set up the prediction task, which models we evaluate, and what results they achieve.

The main idea is very simple: we give each model access to the history of past draws and ask it to predict the 6 numbers that will be drawn next. We then compare how well each model performs against the simplest possible benchmark—picking 6 numbers completely at random. If a model is truly learning something useful from the data, it should predict correct numbers more often than pure random guessing. If it cannot, that is direct evidence that the data contains no learnable structure.

Each draw $d_t = \{n_1, \ldots, n_6\} \subset \{1, \ldots, 49\}$ is represented as a binary array $\mathbf{v}_t \in \{0,1\}^{49}$, where $v_{t,i} = 1$ if number i was drawn. For the ML models, a sliding window of $W = 10$ consecutive draws is used to construct the feature vectors $\mathbf{x}_t \in \mathbb{R}^{544}$, consisting of: the flattened one-hot encodings of draws $d_{t-W}, \ldots, d_{t-1}$ (dimension 490); five aggregate statistics over the window—normalised mean, standard deviation, sum, fraction of even numbers, and fraction of numbers ≤ 24 (dimension 5); and the normalised cumulative frequency vector up to time t (dimension 49). The target is $\mathbf{y}_t = \mathbf{v}_t$.

All models were implemented in Python using `scikit-learn` and `TensorFlow/Keras`. A strict temporal train/test split was used:: 80% for training (draws 1–2008, approximately 1994–2021) and 20% for testing (draws 2009–2510, approximately 2021–2025). Performance is measured by the hit count

$$h_t = |\hat{D}_t \cap D_t| \tag{5}$$

where $\hat{D}_t$ is the set of 6 predicted numbers and D_t the drawn set. The theoretical expectation under random prediction is

$$E[h] = 6 \cdot \frac{6}{49} = \frac{36}{49} \approx 0.7347 \tag{6}$$

derived from the hypergeometric distribution. This is the primary benchmark; real predictive power requires a mean hit rate significantly above this value.

3.1 Models

We evaluate five models of increasing complexity, ranging from non-parametric baselines to ensemble methods, to test whether predictive failure is model-specific or a general consequence of the data properties.

Model 0 – Random Baseline [1]. The random baseline selects 6 numbers uniformly at random without replacement, making no use of historical data. It converges in expectation to $E[h] \approx 0.7347$ and serves as the reference against which all other models are compared.

Model 1 – Historical Frequency [1]. This model predicts the 6 numbers with the highest cumulative frequency in all preceding draws:

$$\hat{D}_t = \underset{S \subset \{1,\ldots,49\},\, |S|=6}{\arg\max} \sum_{i \in S} F_{t-1}(i) \tag{7}$$

where $F_{t-1}(i)$ is the cumulative count of number i up to draw $t-1$. This embodies the common "hot numbers" heuristic. By the test period, cumulative frequencies are nearly uniform, so the model essentially selects a near-fixed set of slightly above-average numbers.

Model 2 – First-order Markov Chain [3,4]. This model estimates a 49×49 transition matrix $\mathbf{P}$ from training data, where P_{ij} is the empirical probability that number j appears given that i appeared in the preceding draw. Each candidate is scored as

$$s_j = 0.7 \cdot \sum_{i \in D_{t-1}} P_{ij} + 0.3 \cdot \hat{p}_j \tag{8}$$

with the second term regularising against sparse transitions using the empirical marginal frequency. Under true draw independence, $\mathbf{P}$ converges to a matrix with identical rows equal to $\mathbf{1}/49$, reducing the model to the setup described by model 1.

Model 3 – Multilayer Perceptron (MLP) [2]. The MLP maps $\mathbf{x}_t \in \mathbb{R}^{544}$ to the multi-label target $\mathbf{y}_t \in \{0,1\}^{49}$:

$$f_\theta(\mathbf{x}_t) = W_3 \cdot \mathrm{ReLU}(W_2 \cdot \mathrm{ReLU}(W_1 \mathbf{x}_t + \mathbf{b}_1) + \mathbf{b}_2) + \mathbf{b}_3 \tag{9}$$

The architecture uses three hidden layers (256, 128, 64 units, ReLU), Adam with L_2 regularisation ($\alpha = 10^{-4}$), and early stopping on a 10% validation split (patience 15 epochs). It is the most expressive model in the evaluation, and therefore it represents the most aggressive test of the hypothesis that no learnable structure exists.

Model 4 – Random Forest. The Random Forest model consists of 100 decision trees trained on bootstrap samples (max depth 10, min samples per leaf 5), implemented as a multi-output classifier with one binary tree per number. The predicted set consists of the 6 numbers with the highest predicted probability across all 49 classifiers:

$$\hat{D}_t = \underset{S \subset \{1,\ldots,49\},\, |S|=6}{\arg\max} \sum_{i \in S} \hat{P}(v_{t,i} = 1 \mid \mathbf{x}_t) \tag{10}$$

The ensemble structure provides inherent regularisation and makes the Random Forest a reliable detector of tabular structure when it exists [14].

3.2 Results

Table 2 summarises the quantitative results for all models based on the 500-draw test set. All models achieve mean hit rates within ± 0.09 deviation from the theoretical random expectation, and no model achieves more than 4 correct predictions on a single draw.

The Markov model records the highest mean hit rate and the only occurrence of 4 correct predictions (0.4% of draws). However, this can be due to the regularisation term in Eq. (8) creating a slight selection bias toward more frequent

Table 2. Model performance on the 500-draw test set. $\bar{h}$: mean hits per draw; s: standard deviation; Δ: difference from random baseline.

Model	$\bar{h}$	s	Δ	$P(h \geq 2)$	$P(h \geq 3)$
Theoretical random	0.7347	–	–	–	–
Random baseline	0.6880	0.742	0.000	13.6%	1.6%
Hist. frequency	0.7500	0.769	+0.062	16.2%	2.0%
Markov (order 1)	0.7740	0.771	+0.086	15.8%	2.0%
MLP neural net.	0.7080	0.789	+0.020	14.8%	3.0%
Random Forest	0.7360	0.742	+0.048	14.2%	1.8%

numbers. The MLP remains slightly below the random baseline—high-capacity models tend to concentrate on historically frequent numbers rather than sampling uniformly, which affects diversity. The Random Forest produces almost exactly at the theoretical expectation. This is arguably the most informative result: Random Forests are known to extract structure from tabular data when it exists, so failing to beat random here suggests there is little to no structure to exploit.

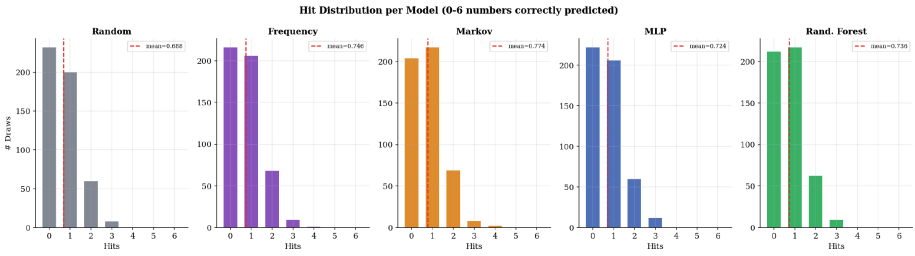

Fig. 5. Hit count distribution per model (500-draw test set).

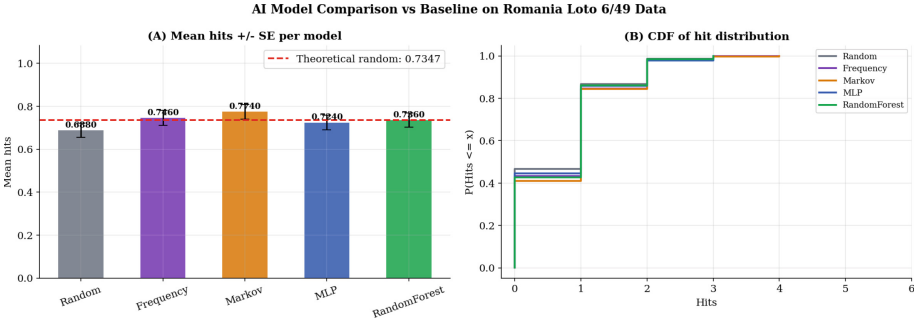

Fig. 6. Mean hit rates with standard error bars (A) and empirical CDFs (B).

Figure 5 shows the distribution of hit counts per model. Figure 6 compares mean hit rates and empirical CDFs across models. Figure 7 tracks rolling mean performance over the test period. Figure 8 breaks down results by hit category. Taken together, the four figures confirm that all models perform at the level of random guessing, with no sustained or structurally distinct above-random behaviour.

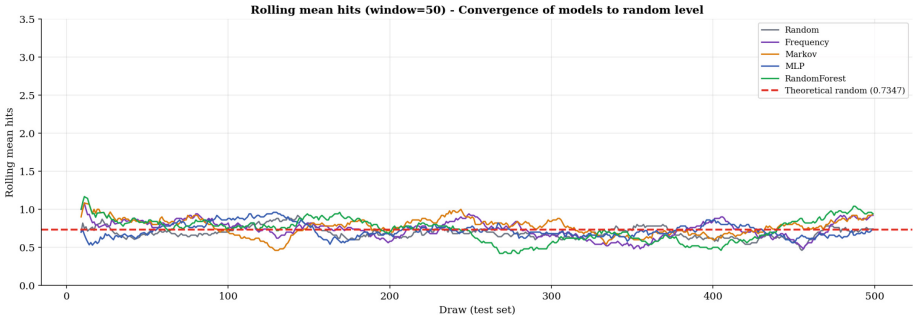

Fig. 7. Rolling mean hit rate (window = 50 draws) over the test period.

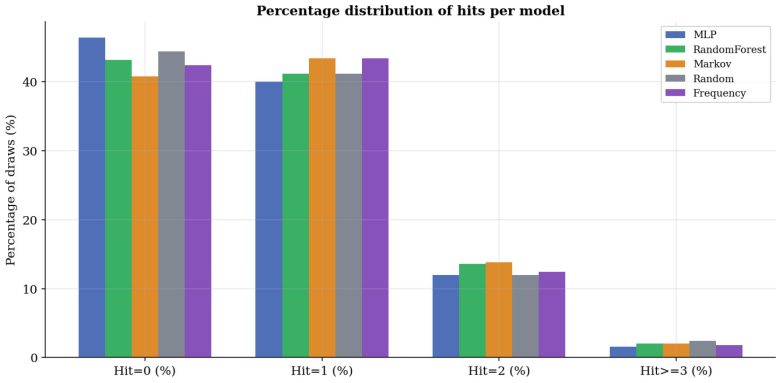

Fig. 8. Percentage of draws per hit category (0, 1, 2, $\geq$3) for all models.

4 Discussion

The lack of ability of all models tested in this paper to produce a good prediction follows directly from information theory basic results. Shannon's source coding theorem establishes that the minimum number of bits required to represent a

random variable is equal to its entropy [15]. A sequence with entropy near the theoretical maximum is essentially incompressible: no representation can encode it in fewer bits than the entropy value. As a direct consequence, any predictive model is searching for a compressed representation that retains information about future draws. However, when no such compression exists, the model is reduced to memorising the noise in the data. This explains why more complex models do not necessarily produce better predictions—the limiting factor is the absence of information in the data, not the model capacity or complexity.

An interesting feature of the results is that the most complex model (MLP) performs slightly worse than the simpler Markov chain. On pseudo-random data, a high-capacity model memorises various correlations in the training set that do not necessarily show general trends. Simpler models have fewer degrees of freedom to overfit and degrade to the random baseline. This has a practical implication beyond the lottery context: on pseudo-random data, model selection should always favour simplicity.

Generally speaking, these results add to the existing reasearch evidence which states that ML performance is limited by the statistical properties of the data, not just by the architecture and training procedure. In fields where data generation approximates a maximum-entropy process—financial data, quantum random noise generators, physical noise sources—the theoretical upper limit on predictability approaches random guessing regardless of the model used. Claims of predictive success on such data should be presented together with rigorous randomness testing, and performance should be compared with expectations based on solid theoretical grounds.

5 Conclusions

In the present paper, we have experimentally shown that machine learning algorithms cannot consistently perform better than random guessing. This conclusion is based on the historical Romanian lotter data, Loto 6/49, using 2,510 draws across 32 complete years (1994–2025). The draw sequence passes several complementary statistical tests for uniformity and statistical independence. In addition, the Shannon entropy for the data set is at 99.96% of its theoretical maximum. Four ML models—a historical frequency baseline, a first-order Markov chain, an MLP neural network, and a Random Forest classifier—were tested against the above-mentioned data set. They achieve mean hit rates consistent with the theoretical random expectation and no model manages to maintain consistent performance better than the average.

The results are not a failure of the specific algorithms put at test. They are a direct consequence of the information-theoretic properties of the data analyzed, and establish a practical upper bound on AI capability for this class of experiments. Future work could extend the analysis to lottery systems with different properties, or investigate whether alternative feature representations change the conclusions for more complex sequence models. The dataset and code can be made available to any interested researcher to support reproducibility of the results.

Acknowledgements. The authors acknowledge the use of publicly available historical data from Loteria Romana. This work was supported by UEFISCDI and National University of Science and Technology Politehnica Bucharest through projects PN-IV-P7-7.1-PTE-2024–0505 and UNSTPB-industry collaboration project 34808/19.09.2025.

Disclosure of Interests. The authors declare no competing interests. AI-assisted tools were used for language editing only; all scientific content and conclusions are the authors' own.

References

1. Haigh, J.: Taking Chances: Winning with Probability. Oxford University Press, Oxford (2003). https://doi.org/10.1093/oso/9780198526636.001.0001
2. Zhang, S., Liu, S., Xiong, J.: Why lottery ticket wins? a theoretical perspective of sample complexity on pruned neural networks. In: Advances in Neural Information Processing Systems (NeurIPS 2021), vol. 34 (2021)
3. Norris, J.R.: Markov Chains. Cambridge University Press, Cambridge (2012). https://doi.org/10.1017/CBO9780511810633
4. Farrell, L., Hartley, R., Lanot, G., Walker, I.: The demand for lotto: the role of conscious selection. J. Bus. Econom. Stat. **18**(2), 228–241 (2000). https://doi.org/10.2307/1392560
5. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. **9**(8), 1735–1780 (1997). https://doi.org/10.1162/neco.1997.9.8.1735
6. Wang, T.V., van Loon, R.J.D.P., van den Assem, M.J., van Dolder, D.: Number preferences in lotteries. Judgm. Decis. Mak. **11**(3), 243–259 (2016). https://doi.org/10.1017/S1930297500003089
7. Bassham, L., et al.: A statistical test suite for random and pseudorandom number generators for cryptographic applications. In: NIST Special Publication 800–22. National Institute of Standards and Technology, Gaithersburg, MD (2010). https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=906762
8. Marsaglia, G.: The Marsaglia Random Number CDROM Including the Diehard Battery of Tests of Randomness. Florida State University, Tallahassee, FL (1995)
9. Dinu, A.: From chaos to security: a comparative study of Lorenz and Rössler systems in cryptography. Cryptography **9**(3), 58 (2025). https://doi.org/10.3390/cryptography9030058
10. Dinu, A., Vlad, A., Hanu, B., Mitrea, A.: The statistical independence for words in printed Romanian language. In: 2020 13th International Conference on Communications (COMM), pp. 319–324, Bucharest, Romania (2020). https://doi.org/10.1109/COMM48946.2020.9142045
11. Dinu, A.: Singularity, observability, and independence: unveiling Lorenz's cryptographic potential. Mathematics **12**(18), 2798 (2024). https://doi.org/10.3390/math12182798
12. Dinu, A., Vlad, A.: Romanian printed language, statistical independence and the type II statistical error. In: 2021 International Conference on Speech Technology and Human-Computer Dialogue (SpeD), pp. 120–125 (2021). https://doi.org/10.1109/SpeD53181.2021.9587397
13. Wald, A., Wolfowitz, J.: On a test whether two samples are from the same population. Ann. Math. Stat. **11**(2), 147–162 (1940). https://doi.org/10.1214/aoms/1177731909

14. Breiman, L.: Random forests. Mach. Learn. **45**(1), 5–32 (2001). https://doi.org/10.1023/A:1010933404324
15. Cover, T.M., Thomas, J.A.: Elements of Information Theory, 2nd edn. Wiley (2006). https://doi.org/10.1002/047174882X

The Generalized Jackknife and the Extremal Index Estimation

Dora Prata Gomes[1(✉)] and M. Manuela Neves[2,3]

[1] NOVA School of Science and Technology (FCT NOVA) and Center for Mathematics and Applications (CMA), FCT NOVA, Caparica, Portugal
dsrp@fct.unl.pt
[2] Instituto Superior de Agronomia, Universidade de Lisboa, Lisbon, Portugal
manela@isa.ulisboa.pt
[3] Faculdade de Ciências, CEAUL – Centro de Estatística e Aplicações, Universidade de Lisboa, Lisbon, Portugal

Abstract. Extreme observations often occur in clusters over time, reflecting persistence in the underlying stochastic process and inducing serial dependence among exceedances. This phenomenon has important implications for risk assessment. Within the framework of Extreme Value Theory, the extremal index is a key parameter for quantifying the degree of clustering of extreme events. In this work, we investigate bias-reduced estimation of the extremal index using the Generalized Jackknife methodology, applied to two widely used estimators: the upcrossing estimator and the Ferro-Segers estimator. A stability-oriented algorithm for extreme value parameter estimation, previously shown to exhibit strong empirical performance, is adapted to guide the selection of both the threshold level k and the tuning parameter inherent to the Generalized Jackknife scheme. This combined approach aims to enhance robustness and reduce sensitivity to threshold choice. An extensive Monte Carlo simulation study is currently in progress to evaluate the finite-sample performance of the proposed methodology under a range of dependence structures. Particular attention is given to bias reduction and variability. Preliminary results suggest that the Generalized Jackknife approach provides a substantial reduction in bias while maintaining competitive overall performance. An application to environmental data illustrates our procedure.

Keywords: Adaptive algorithm · Bias reduction · Extreme value theory · Extremal index · Generalized Jackknife

This work is funded by national funds through the FCT – Fundação para a Ciência e a Tecnologia, I.P., under the scope of the projects UID/00297/2025 (https://doi.org/10.54499/UID/00297/2025) and UID/PRR/00297/2025 (https://doi.org/10.54499/UID/PRR/00297/2025) (Center for Mathematics and Applications – NOVA Math).

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 372–384, 2027.
https://doi.org/10.1007/978-3-032-30518-3_24

1 Preliminaries on Extreme Value Theory

The statistical study of rare and extreme observations plays a central role in many scientific and applied disciplines. In numerous contexts, the main interest lies in unusually large realizations of a stochastic process, such as extreme precipitation, wind speeds, ocean wave heights, financial losses, or wildfire intensities. In other situations, the focus is on extremely small values, for example in the analysis of drought events, structural fatigue, or material degradation. In all these cases, understanding the probabilistic mechanisms governing extreme outcomes is crucial for risk assessment and mitigation.

Extreme events appear in a wide range of areas, including environmental sciences, finance and insurance, telecommunications, engineering reliability, survival analysis and sports science. A key challenge in the statistical analysis of extremes is the intrinsic scarcity of observations in the tail of a distribution. As a consequence, classical statistical methods often perform inadequately in this setting, which motivates the use of specialized probabilistic models and inferential techniques.

Extreme Value Theory (EVT) provides the mathematical framework for describing the stochastic behaviour of extreme observations. Over the last decades, EVT has developed into a well-established area of probability and statistics, supported by solid theoretical foundations and numerous practical applications. Comprehensive accounts of the theory and its applications can be found in [2,3].

Within the EVT framework, particular attention is devoted to quantities that characterize the tail behaviour of the underlying distribution, since these determine the magnitude and frequency of extreme events. Because the full distribution is typically unknown in practice, inference often relies on semiparametric procedures based on asymptotic results that describe the distributional behaviour in the tail region.

Let $\{X_i\}_{i\geq 1}$ denote a sequence of independent and identically distributed (iid) random variables with cumulative distribution function (cdf) F, and define the associated sequence of partial maxima by $M_n := \max(X_1, \ldots, X_n)$. A fundamental result in classical EVT, originating from the pioneering works of Fréchet, Fisher and Tippett and later formalized by Gnedenko and de Haan, establishes that if normalizing sequences exist such that the distribution of the normalized maxima converges to a non-degenerate limit, then the limiting distribution necessarily belongs to the class of extreme value distributions. The corresponding cdf is given by

$$\mathrm{EV}_\xi(\mathrm{x}) = \begin{cases} \exp[-(1+\xi x)^{-1/\xi}], & 1+\xi x > 0 \text{ if } \xi \neq 0 \\ \exp[-\exp(-x)], & x \in R \qquad \text{if } \xi = 0. \end{cases} \tag{1}$$

This formulation has motivated numerous extensions and refinements in the literature, see [8,15].

The parameter ξ in (1), commonly referred to as the *extreme value index*, characterizes the tail heaviness of the distribution and therefore determines the severity of extreme events.

In many real-world processes, extreme observations tend to occur in clusters rather than appearing as isolated events. Such temporal dependence may substantially increase the impact of extreme episodes, since prolonged periods of large values can intensify the underlying phenomenon (see, for example, [4]). In EVT, this clustering behaviour is summarized by the *extremal index* (EI), typically denoted by θ. This parameter takes values in the interval $[0, 1]$, with smaller values indicating stronger clustering of extremes. The limiting case $\theta = 0$ corresponds to a degenerate situation (see, e.g., [17]) and will not be considered in this work.

Although originally introduced within the context of EVT, the EI has also proved useful in several other areas, particularly in the study of dynamical systems (see, for example, [7,18]).

The remainder of the paper is organized as follows. Section 2 presents the theoretical framework for extremes in dependent sequences, introduces the extremal index (EI) as a measure of extremal dependence, reviews existing estimation methods, and describes the Generalized Jackknife estimator considered in this study, together with an automatic stability-based procedure for selecting the tuning parameter δ and the threshold k. Section 3 reports the results of a Monte Carlo simulation study assessing the finite-sample performance of the estimators under temporal dependence. In Sect. 4 we discuss an application to a set of data studied in [2,5] to illustrate our procedure. Finally, a brief conclusion is presented in Sect. 5.

2 Classical and Bias-Reduced Estimators of the Extremal Index

The *extremal index* is a fundamental quantity in the study of extremes for dependent stochastic sequences, as it measures how temporal dependence affects exceedances over high thresholds. In particular, it summarizes how the dependence structure of a stationary process influences both the occurrence rate and the clustering of extreme observations.

Let $\{u_n\}$ be a sequence of increasing thresholds such that, for some constant $\tau > 0$,

$$F(u_n) = 1 - \frac{\tau}{n} + o\left(\frac{1}{n}\right), \qquad n \to \infty, \tag{2}$$

where F denotes the marginal cdf of the process. Within this setting, the EI can be formally defined as follows.

Definition 1 ([17]). *Let $\{X_n\}_{n\geq 1}$ be a strictly stationary sequence with marginal cdf F, and define the partial maxima by*

$$M_n := \max(X_1, \ldots, X_n).$$

The process is said to admit an EI, $\theta \in (0, 1]$ if, for each $\tau > 0$, there exists a sequence of thresholds $\{u_n(\tau)\}_{n\geq 1}$ satisfying

$$n \Pr\{X_1 > u_n(\tau)\} \longrightarrow \tau \quad \text{and} \quad \Pr\{\mathrm{M_n} \leq \mathrm{u_n}(\tau)\} \longrightarrow \mathrm{e}^{-\theta\tau}, \quad \mathrm{n} \to \infty. \tag{3}$$

Several interpretations of the EI have been proposed in the literature. A concise overview of these perspectives is given in [18]. One of the most widely cited interpretations, introduced by [17], links the EI to the clustering of exceedances. Under this viewpoint, θ can be understood as the reciprocal of the limiting mean size of clusters formed by exceedances over high thresholds.

Considerable attention has been devoted to statistical inference for the EI, and a variety of estimation methods have been developed. Early and commonly used approaches include block and runs estimators [16,19,23–25]. These techniques rely on identifying clusters of exceedances through predefined schemes and typically involve tuning parameters whose choice may affect the resulting estimates.

Beyond cluster-based procedures, alternative estimators exploit other probabilistic characterizations of the EI derived from the behaviour of exceedances. Two widely used examples are the upcrossings estimator and the intervals estimator proposed in [6], which are briefly reviewed below.

2.1 The Upcrossings Estimator

Definition 1 also provides the basis for the so-called upcrossings approach to the estimation of the EI, [9–11,19]. The corresponding estimator is defined as

$$\widehat{\theta}^{\mathrm{UC}}(u_n) = \frac{\sum_{i=1}^{n-1} I(X_i \le u_n < X_{i+1})}{\sum_{i=1}^{n} I(X_i > u_n)}. \tag{4}$$

A practical implementation of this estimator was proposed by Gomes *et al.* [12], who suggested selecting the threshold deterministically using order statistics, namely $u = X_{n-k:n}$. This choice leads to the estimator

$$\widehat{\theta}^{\mathrm{UC}}(k) = \frac{1}{k} \sum_{i=1}^{n-1} I(X_i \le X_{n-k:n} < X_{i+1}). \tag{5}$$

For a wide class of dependence structures, such those we are going to consider later in our simulation study, the bias of $\widehat{\theta}^{\mathrm{UC}}(k)$ admits the asymptotic expansion

$$\mathrm{Bias}\left[\widehat{\theta}^{\mathrm{UC}}(k)\right] = \varphi_1(\theta)\frac{k}{n} + \varphi_2(\theta)\frac{1}{k} + o\left(\frac{k}{n}\right) + o\left(\frac{1}{k}\right), \tag{6}$$

as $n \to \infty$, with $k \to \infty$ and $k = o(n)$.

This bias structure naturally motivates the use of the Generalized Jackknife methodology introduced by [14]. Exploiting the expansion in (6), Gomes *et al.* [12] proposed a second–order Generalized Jackknife estimator based on $\widehat{\theta}^{\mathrm{UC}}$ evaluated at three different levels, namely k, $[k/2]+1$ and $[k/4]+1$. The resulting estimator takes the form

$$5\,\widehat{\theta}^{\mathrm{UC}}([k/2]+1) - 2\Big(\widehat{\theta}^{\mathrm{UC}}([k/4]+1) + \widehat{\theta}^{\mathrm{UC}}(k)\Big). \tag{7}$$

Although this estimator substantially reduces bias, it does not improve the mean squared error relative to the original estimator when the threshold level is chosen optimally. To address this limitation, the authors subsequently considered a more general set of evaluation levels, namely k, $[\delta k]+1$ and $[\delta^2 k]+1$, where the parameter δ satisfies $0<\delta<1$.

This idea leads to a family of estimators, denoted by $\widehat{\Theta}^{\mathrm{GJ}(\delta)}$, constructed from linear combinations of $\widehat{\Theta}^{\mathrm{UC}}$ evaluated at the three levels k, $[\delta k]+1$ and $[\delta^2 k]+1$. In particular, Gomes *et al.* [12] introduced the estimator

$$\widehat{\Theta}^{\mathrm{GJ.UC}(\delta)} = \frac{(\delta^2+1)\,\widehat{\Theta}^{\mathrm{UC}}([\delta k]+1) - \delta\Big(\widehat{\Theta}^{\mathrm{UC}}([\delta^2 k]+1) + \widehat{\Theta}^{\mathrm{UC}}(k)\Big)}{(1-\delta)^2}, \tag{8}$$

which is designed to remove the two leading bias components. Consequently, the estimator is asymptotically unbiased. Furthermore, under suitable regularity conditions, both the upcrossings estimator $\widehat{\Theta}^{\mathrm{UC}}$ and the Generalized Jackknife estimator $\widehat{\Theta}^{\mathrm{GJ.UC}(\delta)}$ are consistent and asymptotically normal for $\theta<1$, see [12,19].

2.2 The Intervals Estimator

An alternative approach relies on the temporal spacing between exceedances of a high threshold. The clustering behaviour of extremes is reflected in the distribution of the time intervals separating consecutive exceedances.

Let $S_1<S_2<\cdots<S_N$ denote the ordered times at which exceedances above a high threshold u occur. The corresponding inter-exceedance times are defined as

$$T_j = S_{j+1}-S_j, \qquad j=1,\ldots,N-1.$$

Ferro and Segers [6] showed that the asymptotic distribution of these intervals contains information about EI. Based on this property, they proposed the *intervals estimator*, defined as

$$\widehat{\Theta}^{\mathrm{FS}} = \frac{2\left(\sum_{j=1}^{N-1}(T_j-1)\right)^2}{(N-1)\sum_{j=1}^{N-1}(T_j-1)(T_j-2)}. \tag{9}$$

Unlike runs-based estimators, this approach avoids explicit declustering and instead relies solely on the spacing between exceedances. Under suitable regularity conditions, the estimator is consistent and asymptotically normal.

The bias structure observed for the UC estimator suggests that similar bias–reduction techniques may also be beneficial for other EI estimators. In particular, the Generalized Jackknife methodology could be naturally extended to the intervals estimator of Ferro and Segers, $\widehat{\Theta}^{\mathrm{FS}}$. This idea is explored in the present work, where a bias-reduced versions of the $\widehat{\Theta}^{\mathrm{FS}}$ estimator are constructed as

$$\widehat{\Theta}^{\mathrm{GJ.FS}(\delta)} = \frac{(\delta^2+1)\,\widehat{\Theta}^{\mathrm{FS}}([\delta k]+1) - \delta\Big(\widehat{\Theta}^{\mathrm{FS}}([\delta^2 k]+1) + \widehat{\Theta}^{\mathrm{FS}}(k)\Big)}{(1-\delta)^2}, \tag{10}$$

and their finite-sample performance are also investigated in a simulation study.

2.3 Computational Algorithm for Parameter Selection

The practical implementation of the bias–reduced estimator $\widehat{\Theta}^{GJ.\bullet(\delta)}$, where $\bullet$ denotes an EI estimator, requires the specification of two parameters: the number of upper order statistics k and the tuning parameter δ. The performance of the estimator may depend substantially on these choices, making their selection an important aspect of the estimation procedure.

To address this issue, we adopt a stability–based strategy inspired by the algorithmic approach proposed by [13] and later refined by [20]. This methodology exploits regions of stability in the sequence of estimates as a function of k, allowing the identification of suitable parameter values. Similar ideas have also been considered in the context of extremal index estimation; see [22].

The procedure described below provides a data–driven method for selecting both the tuning parameter δ and the threshold level k.

1. For a given sample $(x_1, \ldots, x_n)$, compute the estimator $\widehat{\Theta}^{\mathrm{GJ}.\bullet}(k)$, as those defined in defined in (8) and (10), for all $k = 1, \ldots, n-1$ and for a grid of values $\delta \in (0, 1)$.
2. For each k, round the estimates to zero decimal places and define

$$a_k^{\widehat{\Theta}^{GJ.\bullet(\delta)}}(0) = \mathrm{round}\Big(\widehat{\Theta}^{GJ.\bullet(\delta)}(k), 0\Big).$$

3. For each δ, identify consecutive values of k for which $a_k^{\widehat{\Theta}^{GJ.\bullet(\delta)}}(0)$ remains constant. Let $k_{\min}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}$ and $k_{\max}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}$ denote the endpoints of the longest such sequence, with length

$$l_{\widehat{\Theta}^{GJ.\bullet(\delta)}} = k_{\max}^{\widehat{\Theta}^{GJ.\bullet(\delta)}} - k_{\min}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}.$$

4. Determine the smallest integer $j_0 \geq 0$ such that the rounded values

$$a_k^{\widehat{\Theta}^{GJ.\bullet(\delta)}}(j) = \mathrm{round}\Big(\widehat{\Theta}^{GJ.\bullet(\delta)}(k), j\Big)$$

   are not all identical when $j = j_0$.
5. Using the sequence $a_k^{\widehat{\Theta}^{GJ.\bullet(\delta)}}(j_0)$, identify the longest run of equal values and denote its endpoints by $k_{\min}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}$ and $k_{\max}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}$.

6. For $k_{\min}^{\widehat{\Theta}^{GJ.\bullet(\delta)}} \leq k \leq k_{\max}^{\widehat{\Theta}^{GJ.\bullet(\delta)}}$, recompute the estimates rounded to $j_0 + 2$ decimal places and determine their modal value. Let $\mathcal{K}_{\widehat{\Theta}^{GJ.\bullet(\delta)}}$ denote the set of indices attaining this value.

7. Define

$$\hat{k}_{\widehat{\Theta}^{GJ.\bullet(\delta)}} = \max \mathcal{K}_{\widehat{\Theta}^{GJ.\bullet(\delta)}}.$$

8. Choose δ_0 as the value of δ that maximizes $l_{\widehat{\Theta}^{GJ.\bullet(\delta)}}$, for all values of δ considered.

9. The final estimate of the extremal index is given by

$$\widehat{\Theta}^{GJ.\bullet(\delta_0)}\left(\hat{k}_{\widehat{\Theta}^{GJ.\bullet(\delta_0)}}\right).$$

The algorithm therefore identifies stability regions in the sequence of estimates and uses this information to determine suitable values for both the tuning parameter δ and the threshold level k.

3 Monte Carlo Study of Extremal Index Estimators Under Temporal Dependence

Monte Carlo simulation experiments are a standard tool for evaluating the performance of statistical estimators in finite samples. Such experiments allow full control over the data-generating mechanism and the true parameter values, making it possible to analyse the behaviour of competing estimators under well-defined stochastic settings. All simulations were conducted using the *R* software environment. In this section we examine the empirical performance of the EI estimators considered in this work. In particular, the study includes both the bias-reduced estimators based on the Generalized Jackknife methodology. The comparison is carried out through Monte Carlo experiments under dependent stochastic models with known extremal properties.

The simulated data are generated from first-order max-autoregressive processes, commonly referred to as ARMAX models. These processes are frequently used in extreme value analysis because they possess analytically tractable extremal characteristics while allowing the strength of extremal dependence to be controlled through a single parameter.

Max-Autoregressive Process I (ARMAX I) Let $\{Z_i\}_{i\geq 1}$ be independent random variables with unit Fréchet distribution. For a parameter $0 < \theta \leq 1$, define the stochastic process $\{X_i\}_{i\geq 1}$ recursively as

$$X_1 = Z_1, \qquad X_i = \max\big\{(1-\theta)X_{i-1},\, \theta Z_i\big\}, \quad i \geq 2.$$

This construction yields a strictly stationary process with unit Fréchet marginal distribution, corresponding to an extreme value index $\xi = 1$. Moreover, for threshold sequences of the form $u_n = nx$, with $0 < x < \infty$, the distribution of the sample maximum

$$M_n = \max(X_1, \ldots, X_n)$$

satisfies

$$P\{M_n \leq u_n\} \longrightarrow \exp\left(-\frac{\theta}{x}\right), \qquad n \to \infty.$$

The EI associated with this process is equal to θ, [2].

Max-Autoregressive Process II (ARMAX II) Let $\{Z_i\}_{i\geq 1}$ be independent random variables with unit Fréchet distribution and let X_0 be a random variable with cumulative distribution function

$$H_0(x) = \exp\big(-x^{-1}(\beta^{-1}-1)\big).$$

For a parameter $\beta \in (0,1)$, define the stochastic process

$$X_i = \beta \max\{X_{i-1}, Z_i\}, \qquad i \geq 1.$$

The EI of this process is given by $\theta = 1-\beta$, [1].

The EI is estimated using the algorithm described in Sect. 2.3, based on samples of size $n = 500,\ 1000,\ 2000$ generated from both ARMAX I and ARMAX II models. The simulation study considers several values of EI together with different choices of the tuning parameter δ. For each model, Tables 1 and 2 report the corresponding estimates obtained using the two estimators, $\widehat{\Theta}^{\mathrm{UC}}$ and $\widehat{\Theta}^{\mathrm{FS}}$, together with their bias-reduced versions for some values of δ. For each configuration, the value of k associated with the longest stability region is also presented for every value of δ considered, along with the corresponding estimate of θ. As an initial assessment of the results, the combinations of δ, k, and the resulting estimates that are closest to the true value of θ are highlighted in **bold**.

These preliminary simulation results suggest that, for smaller sample sizes estimates closer to the true parameter value tend to be obtained for relatively small values of δ. As the sample size increases, however, the values of δ that yield estimates nearest to the true parameter value also tend to increase.

Table 1. EI estimates selecting the parameter δ and the threshold level k, using algorithm described in Sect. 2.3 for ARMAX I process

	$n = 500$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	207	0.1683	340	0.1878	248	0.2691	323	0.3353	161	0.5864	237	0.8244
δ	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	412	0.1722	292	0.0774	**425**	**0.4915**	385	0.4779	**237**	**0.7942**	265	0.6053
0.20	433	0.2106	403	0.2022	293	0.4339	342	0.3923	426	0.7037	**432**	**0.7165**
0.25	**489**	**0.2445**	304	0.1781	290	0.4868	**472**	**0.4857**	458	0.7971	347	0.5937
0.30	390	0.2032	453	0.2943	375	0.5705	336	0.5764	429	0.8097	317	0.5025
0.35	422	0.2297	386	0.2866	298	0.6030	317	0.5949	422	0.8938	343	0.9617
0.40	369	0.2042	470	0.3132	378	0.3893	409	0.0614	487	0.9767	287	0.8856
0.45	314	0.2131	461	0.3122	453	0.4053	349	0.1067	269	0.8259	266	0.5729
0.50	399	0.2894	**434**	**0.2281**	353	0.3238	285	0.1001	474	0.9326	406	1.0000
	$n = 1000$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	825	0.0847	411	0.2709	338	0.3864	567	0.4193	354	0.5127	404	0.7889
δ	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	**754**	**0.2289**	592	0.1984	651	0.5934	335	0.6225	849	0.7904	400	0.9193
0.20	983	0.2760	928	0.3362	526	0.5973	**869**	**0.5060**	**871**	**0.7635**	653	0.8026
0.25	870	0.2775	**640**	**0.3264**	**895**	**0.5081**	972	0.5797	987	0.7321	850	0.9083
0.30	924	0.3092	871	0.4687	757	0.4613	742	0.5979	897	0.7312	658	0.8749
0.35	870	0.3135	980	0.4137	507	0.4795	530	0.5221	923	0.7098	671	0.9558
0.40	656	0.3048	913	0.3979	958	0.6133	934	0.6118	779	0.6833	589	1.0000
0.45	649	0.3244	647	0.4343	910	0.5839	637	0.4900	750	0.7298	**440**	**0.7940**
0.50	732	0.1891	576	0.5207	752	0.6146	693	0.6128	660	0.7189	788	1.0000
	$n = 2000$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	640	0.2197	714	0.2607	491	0.4024	1389	0.3905	265	0.6955	745	0.7637
δ	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	**1972**	**0.2727**	1838	0.2739	1985	0.5331	1758	0.5998	1939	0.8024	1328	0.7138
0.20	1523	0.2756	**1437**	**0.2634**	1988	0.4994	1724	0.5658	1843	0.7841	**1695**	**0.7636**
0.25	1738	0.3180	690	0.2154	1718	0.5061	1654	0.4139	1986	0.8013	1870	0.8905
0.30	1034	0.2665	1845	0.4282	**1324**	**0.5005**	1455	0.4216	1845	0.7982	1140	0.7766
0.35	1705	0.3239	1795	0.4331	1528	0.5214	1815	0.5910	1827	0.7951	942	0.7919
0.40	1638	0.3347	1388	0.4882	1633	0.5321	**1651**	**0.5004**	**1589**	**0.7765**	1462	1.0000
0.45	1840	0.3262	1780	0.4253	1880	0.5994	1762	0.6204	1607	0.8407	1396	0.9129
0.50	1886	0.3209	1839	0.3809	1689	0.5955	1287	0.4938	1357	0.7844	1234	0.8835

Table 2. EI estimates selecting the parameter δ and the threshold level k, using algorithm described in Sect. 2.3 for ARMAX II process

	$n = 500$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	333	0.1287	448	0.1224	153	0.4026	348	0.4169	253	0.3858	248	0.5799
δ	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	**463**	**0.2372**	474	0.2889	277	0.3323	342	0.3985	345	0.8769	465	0.9407
0.20	451	0.2734	378	0.3484	241	0.3068	405	0.5560	461	0.7810	224	0.9358
0.25	474	0.2808	**487**	**0.2263**	**410**	**0.4667**	**446**	**0.5205**	369	0.7922	**385**	**0.7412**
0.30	488	0.3074	487	0.1822	418	0.5973	316	0.5445	274	0.8195	388	1.0000
0.35	410	0.2654	298	0.3184	371	0.5745	303	0.5953	**329**	**0.7255**	170	0.8275
0.40	343	0.2833	464	0.1082	369	0.6429	340	0.6253	306	0.7120	334	0.8150
0.45	270	0.2869	273	0.0780	282	0.5792	462	0.7093	413	0.6139	393	0.7866
0.50	454	0.3156	282	0.0344	347	0.7789	364	0.4394	468	0.6857	359	0.6233
	$n = 1000$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	462	0.1901	653	0.1696	389	0.3589	406	0.4889	647	0.2978	459	0.6431
δ_0	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	**652**	**0.2654**	732	0.4049	**941**	**0.5029**	**602**	**0.5054**	553	0.7123	627	0.6782
0.20	983	0.2735	893	0.3974	705	0.5154	453	0.4158	**848**	**0.7702**	**605**	**0.7220**
0.25	834	0.2668	890	0.3881	985	0.5290	918	0.5993	934	0.7784	864	0.6409
0.30	915	0.2815	858	0.3879	970	0.5079	893	0.7049	965	0.7931	708	0.5778
0.35	986	0.3212	721	0.3789	818	0.5162	850	0.7012	892	0.8285	786	0.7090
0.40	953	0.3301	873	0.3894	915	0.5673	557	0.6783	786	0.8023	590	0.6863
0.45	459	0.1989	741	0.4207	969	0.5903	879	0.7272	730	0.7841	611	0.7846
0.50	841	0.3291	**816**	**0.3167**	948	0.5848	689	0.4659	922	0.8224	545	0.6008
	$n = 2000$											
	$\theta = 0.25$				$\theta = 0.50$				$\theta = 0.75$			
	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$	$\hat{k}$	$\widehat{\theta}^{UC}$	$\hat{k}$	$\widehat{\theta}^{FS}$
	924	0.1676	1319	0.1579	556	0.3896	1141	0.4291	521	0.5594	799	0.6726
δ	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS.}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$	$\hat{k}$	$\widehat{\theta}^{GJ.UC}$	$\hat{k}$	$\widehat{\theta}^{GJ.FS}$
0.15	**1985**	**0.2382**	1365	0.2981	1746	0.4695	1450	0.3937	1232	0.7985	785	0.9122
0.20	1853	0.2257	**1743**	**0.2064**	1988	0.5229	913	0.3188	1688	0.7031	1448	0.8022
0.25	1335	0.2025	1862	0.2033	1886	0.5233	718	0.2974	1840	0.6780	1286	0.7281
0.30	1585	0.2295	1735	0.1948	**1655**	**0.5116**	1692	0.5986	1802	0.7015	**1431**	**0.7349**
0.35	1393	0.2295	1744	0.1791	1329	0.5126	1733	0.5847	1111	0.5899	1584	0.7700
0.40	1766	0.2686	1563	0.1664	1191	0.5238	1389	0.6278	**1770**	**0.7874**	1170	0.7278
0.45	1060	0.2239	1794	0.1722	1694	0.5270	**1701**	**0.5061**	1525	0.7951	1339	0.8071
0.50	1798	0.3222	1453	0.1299	1499	0.4933	1454	0.4775	1372	0.7903	918	0.8348

4 An Application to Real Data

In this section we apply our procedure to the series of daily maximum temperatures (in degrees Celsius) recorded in Uccle, Belgium, during the month of July over the period 1901–1999. The dataset contains $n = 3051$ observations and is available in the `exdex` package in `R` under the name `Uccle`. Estimates are computed over a wide range of threshold values k, allowing the identification of stable regions in the estimation of the EI. The tunparameter δ is selected through the proposed stability-based algorithm, yielding the value $\delta = 0.2$. Conditional on this choice, the algorithm identifies the optimal threshold $k = 2973$ for $\widehat{\Theta}^{GJ_UC}$ estimator, corresponding to a pronounced stability region and producing the estimate $\widehat{\theta}^{GJ_UC} = 0.4804$. Similarly for the $\widehat{\Theta}^{GJ_FS}$ estimator, the optimal threshold is found to be $k = 2915$, leading to the estimate $\widehat{\theta}^{GJ_FS} = 0.4782$. Figure 1 illustrates the procedure considered here and also displays the corresponding extremal index estimates.

This dataset was previously investigated by [2] in the context of extremal dependence analysis using parametric approaches. The estimated values of the extremal index were found to lie approximately between 0.49 and 0.56. Later [5] considered this dataset to illustrate the comparison of some estimation methods. The values obtained with our procedure agree with the values presented in those papers.

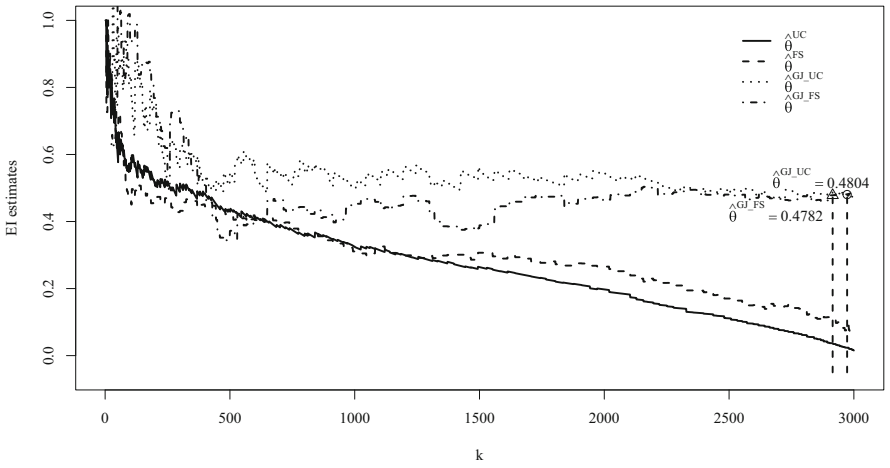

Fig. 1. Adaptive choice of the level k with the optimal value of δ ($\delta = 0.2$) for estimating θ for the daily maximum temperatures (in degrees Celsius) in Uccle, during the month of July, over the period 1901–1999.

5 Concluding Remarks

This work aims to improve the practical estimation of the parameter θ, with particular emphasis on bias reduction and on the data-driven selection of tuning parameters that strongly influence finite-sample performance.

In this context, we investigated the use of the Generalized Jackknife applied to two EI estimators. A stability-based procedure for the automatic calibration of both the tuning parameter δ and the threshold level k used in the estimation of the EI is proposed. The main idea is to identify regions where the estimator exhibits stable behaviour, providing a principled criterion for selecting the relevant tuning parameters.

A preliminary version of this procedure was presented in [21], where only samples of size $n = 1000$ were considered for the UC estimator.

In the present work, the performance of the proposed approach was further examined through simulation experiments under scenarios with temporal dependence. The results suggest that the stability-based calibration provides a systematic and effective strategy for determining the tuning parameters required by the estimator. Further simulation studies involving a wider range of sample sizes, dependence structures, and parameter configurations are currently under development, with the aim of providing a more comprehensive assessment of the robustness and applicability of the proposed methodology. An application to a real dataset already used in previous studies illustrates the good performance of our methodology.

Acknowledgments. The authors thank the referees for their constructive comments and valuable suggestions, which led to substantial improvements to this work.

This work is funded by national funds through the FCT – Fundação para a Ciência e a Tecnologia, I.P., under the scope of the projects UID/00297/2025 (https://doi.org/10.54499/UID/00297/2025) and UID/PRR/00297/2025 (https://doi.org/10.54499/UID/PRR/00297/2025) (Center for Mathematics and Applications – NOVA Math).

References

1. Alpuim, M.T.: An extremal Markovian sequence. J. Appl. Probab. **26**(2), 219–232 (1989). https://doi.org/10.2307/3214030
2. Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J.L.: Statistics of Extremes: Theory and Applications. Wiley (2006). https://doi.org/10.1002/SERIES1345
3. Embrechts, P., Klüppelberg, C., Mikosch, T.: Modelling Extremal Events. AM, vol. 33. Springer, Heidelberg (1997). https://doi.org/10.1007/978-3-642-33483-2
4. Faranda, D., Messori, G., Yiou, P., Thao, S., Pons, F., Dubrulle, B.: Dynamical footprints of hurricanes in the tropical dynamics. Chaos Interdisc. J. Nonl. Sci. **33**(1) (2023). https://doi.org/10.1063/5.0093732
5. Ferreira, M.: Heuristic tools for the estimation of the extremal index: a comparison of methods. REVSTAT-Stat. J. **16**(1), 115–136 (2018)
6. Ferro, C.A.T., Segers, J.: Inference for clusters of extreme values. J. R. Stat. Soc. Ser. B Stat Methodol. **65**(2), 545–556 (2003)

7. Freitas, A.C.M., Freitas, J.M., Soares, J.V.: Rare events for product fractal sets. J. Phys. A: Math. Theor. **54**(34), 345202 (2021). https://doi.org/10.1088/1751-8121/ac16c6
8. Gnedenko, B.: Sur la distribution limite du terme maximum d'une serie aleatoire. Ann. Math. **3**(44), 423–453 (1943). https://doi.org/10.2307/1968974
9. Gomes, M.I.: Statistical inference in an extremal Markovian model. In: Compstat: Proceedings in Computational Statistics, 9th Symposium held at Dubrovnik, Yugoslavia, pp. 257–262. Springer (1990). https://doi.org/10.1007/978-3-642-50096-1-39
10. Gomes, M.I.: In: I Congresso Ibero-Americano de Esdadistica e Investigacion Operativa, pp. 209–220 (1992). Modelos extremais em esquemas de dependência
11. Gomes, M.I.: On the estimation of parameters of rare events in environmental time series. Stat. Environ. **2**, 225–241 (1993)
12. Gomes, M.I., Hall, A., Miranda, M.C.: Subsampling techniques and the jackknife methodology in the estimation of the extremal index. Comput. Stat. Data Anal. **52**(4), 2022–2041 (2008). https://doi.org/10.1016/j.csda.2007.06.023
13. Gomes, M.I., Henriques-Rodrigues, L., Fraga Alves, M.I., Manjunath, B.G.: Adaptive port–MVRB estimation: an empirical comparison of two heuristic algorithms. J. Stat. Comput. Simul. **83**(6), 1129–1144 (2013). https://doi.org/10.1080/00949655.2011.652113
14. Gray, H.L., Schucany, W.R.: The generalized jackknife statistic. M. Dekker, New York (1972). https://doi.org/10.1080/00224065.1973.11980580
15. de Haan, L.: On Regular Variation and Its Application to the Weak Convergence of Sample Extremes, vol. 32. Centrum Voor Wiskunde en Informatica (1970)
16. Hsing, T.: Extremal index estimation for a weakly dependent stationary sequence. Ann. Stat., 2043–2071 (1993). https://doi.org/10.1214/aos/1176349409
17. Leadbetter, M.R., Lindgren, G., Rootzen, H.: Extremes and Related Properties of Random Sequences and Processes. Springer Series in Statistics. Springer (1983). https://doi.org/10.1007/978-1-4612-5449-2
18. Moloney, N.R., Faranda, D., Sato, Y.: An overview of the extremal index. Chaos Interdisc. J. Nonl. Sci. **29**(2) (2019). https://doi.org/10.1063/1.5079656
19. Nandagopalan, S.: Multivariate Extremes and Estimation of the Extremal Index. The University of North Carolina at Chapel Hill (1990)
20. Neves, M.M., Gomes, M.I., Figueiredo, F., Gomes, D.P.: Modeling extreme events: sample fraction adaptive choice in parameter estimation. J. Stat. Theory Pract. **9**(1), 184–199 (2014). https://doi.org/10.1080/15598608.2014.890984
21. Prata Gomes, D., Neves, M.M.: Computational Methods for Extremal Index Estimation. Accepted
22. Prata Gomes, D., Neves, M.M.: Resampling procedures for a more reliable extremal index estimation. In: Data Analysis and Applications 4, chap. 6, pp. 89–101. Wiley (2020). https://doi.org/10.1002/9781119721611.ch6
23. Robert, C.Y., Segers, J., Ferro, C.A.: A sliding blocks estimator for the extremal index. Electron. J. Stat. **3**, 993–1020 (2009). https://doi.org/10.1214/08-EJS345
24. Smith, R.L., Weissman, I.: Estimating the extremal index. J. Royal Stat. Soc. Ser. B (Methodol. **56**(3), 515–528 (1994). https://doi.org/10.1111/j.2517-6161.1994.tb01997.x
25. Weissman, I., Novak, S.Y.: On blocks and runs estimators of the extremal index. J. Stat. Plan. Inference **66**(2), 281–288 (1998). https://doi.org/10.1016/S0378-3758(97)00095-5

Circular Economy and Territorial Competitiveness (CETC 2026)

Measuring Urban Competitiveness: An Application at the Urban Districts Level in the Galicia Cities of Vigo and Coruna (Spain)

Gerardo Carpentieri[1], Carmen Guida[1(✉)], Manuel Fernandez-Grela[2], Yolanda Pena Boquete[2], and Jackiline Salazar-Torres[2]

[1] University of Naples Federico II, Naples, Italy
carmen.guida@unina.it
[2] Universidade de Santiago de Compostela, Santiago, Spain

Abstract. Urban competitiveness, the study of which has broadened significantly in recent decades, is the ability of a specific urban territory to attract investment, people, and new activities. As highlighted from the previous studies, it depends on a multitude of closely interrelated factors that characterise urban areas. This work aims to develop a multidimensional methodology to quantitatively evaluate urban competitiveness at the district level and to support decision-making. The paper proposes a composite index structured into several dimensions that are useful for comparing the competitive performance of cities' districts relative to the city's social, economic, and physical characteristics, thereby providing a multidimensional ranking. The application phase focuses on the comparative analysis of the districts of the two main Galician cities, Vigo and Coruna, in Spain. The overall competitive performance of the city's districts is highlighted, and their strengths and shortcomings across different dimensions are considered. The results aim to emphasise the main components of the districts' competitiveness in two cities and to support decision-makers in improving competitiveness.

Keywords: urban competitiveness · Galician cities · multidimensional methodology

1 Introduction

Over the past 30 years, the concept of urban competitiveness has quietly changed. Initially limited to a macroeconomic perspective, it has gradually evolved into a multi-layered, systematic framework encompassing local capital, social dynamics, environmental sustainability, and the experience economy [1, 2]. Although the concept was originally based on national models such as the Global Competitiveness Index, it gradually changed as scientists and policymakers realised that competitiveness dynamics primarily operate at the urban and regional level and therefore localised analysis methods are required [3, 4]. Early definitions, such as Begg's pioneering research, focused on a city's ability to generate income and employment while competing with the outside world. Competitiveness itself was not regarded as an end, but rather as a means of increasing the well-being of their inhabitants [5].

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 387–401, 2027.
https://doi.org/10.1007/978-3-032-30518-3_25

Later, insights from global urban theory and network analysis illuminated the strategic role of cities. It became clear that the positioning of cities within the cross-border flows of finance, services, and information shapes their competitiveness [6]. In addition, the rise of entrepreneurial urbanism has increased interest in growth strategies, public-private partnerships and local Marketing. At the same time, however, this process has raised concerns about social division, gentrification and the prioritization of competitiveness over equity and environmental goals. In response to these developments, contemporary urban planning, especially the concept of "local capital" developed by Camani and Capello, reassessed the source of competitiveness: it lies in the integrated use of tangible and intangible local resources such as infrastructure, human capital, relational capital, environmental and cultural assets, and institutional quality-all these factors together contribute to the true strength of a city [7].

A significant paradigm shift in the understanding of urban competitiveness was the World Economic Forum's Definition as "the collective force of policies, institutions, strategies, and processes that shape a city 's sustainable productivity" [1]. This formally incorporates economic, social, and environmental aspects into city valuation, consistent with the concept of the "Triple Bottom Line", and reinforces the recognition that competitiveness is inherently multifaceted. Modern research highlights several core areas, including economic structure, innovation capacity, human capital, accessibility, environmental sustainability, quality of life, cultural diversity, Governance, and international cooperation [2]. It is now argued that long-term competitiveness is equally closely linked to social cohesion, inclusion, environmental responsibility and resilience are linked to economic growth [8]. In recent years, social and environmental Competitiveness has become increasingly important. This shows how high levels of social capital, participatory governance, and demonstrable success in environmental protection increase urban Competitiveness. Sustainable tourism has also established itself as a new field that helps link cultural resources, environmental quality and service provision with competitive advantages [9, 10].

Large-scale comparative studies, like the Global Cities Competitiveness Report, have demonstrated a clear connection between urban performance and progress in reaching the Sustainable Development Goals (SDGs) [11]. However, critical research has also highlighted the contradictions, which are inherent in efforts to increase competitiveness. In particular, empirical links between competitive success and increasing inequality and urban polarization have been highlighted. In light of these criticisms, the concept of competitiveness is redefined. It is argued that inclusiveness, environmental health, resilience and Innovation should not be viewed as constraints on economic outcomes, but as frameworks that simultaneously promote them [12].

In this perspective, the transition to a circular economy can be framed as an important driver of sustainable urban competitiveness, as it aims to decouple economic performance from finite resource consumption through regenerative models. The circular economy is commonly articulated around three design-driven principles: eliminating waste and pollution, circulating products and materials at their highest value, and regenerating natural systems.

At the EU level, circular economy policy explicitly links circularity to a "cleaner and more competitive" economic model, reinforcing its relevance for competitiveness-oriented urban strategies.

To align the competitiveness perspective with circular economy (CE) and urban resilience agendas, we interpret district competitiveness as the capacity to generate well-being while reducing resource throughput and environmental externalities. In this sense, CE-related urban assets, such as proximity to enabling reuse/repair services, access to shared mobility, and the presence of green/blue infrastructure, act as drivers of competitiveness by strengthening local value retention, reducing exposure to resource-price shocks, and improving liveability. Therefore, although our baseline index is built on social, economic, and urban-structure dimensions, the empirical framework is designed to be extensible toward CE-specific indicators when district-scale data become available.

This paper contributes by: (1) operationalising competitiveness at district scale using an integrated socio-economic and urban-structure indicator system; (2) enabling cross-city comparability through joint standardisation of all districts; (3) providing a reproducible GIS-based workflow for mapping intra-urban disparities and supporting policy design.

2 Competitiveness and Urban Development

The growing interest in integrating the concepts of competitiveness and urban development arises from the increasingly central role that cities play in contemporary societies [13]. Urban areas have become key reference points for global economic and social dynamics, and in recent years, they have entered into direct competition to offer favourable conditions for attracting investments, residents, and skilled human capital. Within this globalised context, the traditional approach, focused solely on the relationships between a city and its immediate surroundings, has become insufficient [14]. Advances in communication and transport technologies now enable cities to interact with distant territories, thereby allowing geographically separated urban systems to compete across diverse domains, including industrial production, specialised services, and environmental management [15].

Major metropolitan areas such as New York, London, and Tokyo now orient their development strategies towards competing with similarly sized international cities, often giving secondary importance to internal national competition. In the European context, the establishment of the single market and the free movement of goods and people have heightened cities' awareness of opportunities and risks associated with European integration [16]. Similarly, in Asia, cities such as Hong Kong compete with rapidly expanding urban centres, including Guangzhou and Shanghai, to strengthen their roles as continental economic and financial hubs [17].

A widely used definition describes urban competitiveness as a city's ability to maintain or enhance its position within a regional, national, or international context. Competition thus occurs between comparable urban systems pursuing similar objectives related to resource preservation and the improvement of residents' well-being, through the effective management of external and internal drivers of urban development [18].

Before becoming fully established in the urban planning literature, the concept of competitiveness was supported by early studies that focused primarily on economic

factors contributing to territorial attractiveness. Kresl & Singh (1999) proposed a two-component framework distinguishing between economic aspects, such as production, infrastructure, location, economic structure, and urban services, and strategic aspects, including government efficiency, development strategies, public–private cooperation, and institutional flexibility [19]. As a result, dimensions not directly linked to business performance were typically excluded.

The subsequent scientific debate expanded the scope of analysis. Martin and Simmie (2008) defined urban competitiveness as the capacity of cities to continually improve their business environment, human capital, and physical, social, and cultural infrastructures [20]. According to their view, a competitive city is capable of attracting and retaining innovative, high-growth firms and a skilled, entrepreneurial workforce, ultimately achieving high productivity, employment, and wages while reducing socio-economic disparities [21, 22].

Further contributions incorporated both tangible and intangible factors, such as environmental quality, cultural and technological resources, human capital, and aesthetic attributes, that directly or indirectly influence the locational choices of firms and individuals seeking to meet economic and personal needs [23, 24].

In urban contexts, circular economy principles translate into "value-retention" loops that reduce material throughput and waste while potentially supporting local economic dynamism and resilience.

Across scientific disciplines, increasing attention has been devoted to developing quantitative and qualitative methods for measuring urban competitiveness. This trend is driven by the need of public and private decision-makers for reliable information systems that support development strategies [25]. Nevertheless, assessing the competitiveness of territorial systems, whether regions, metropolitan areas, or cities, remains challenging due to the multidimensional nature of the phenomena involved, including well-being, quality of life, infrastructure, and service provision [26]. Consequently, analyses relying on single indicators are insufficient, underscoring the need for multidimensional indicator sets that capture a broader range of structural attributes [27].

Today, there is broad agreement that meeting the needs of individuals and groups for an improved quality of life is central to the competitiveness of cities. These new needs are leading to a shift in urban priorities: from production-oriented approaches such as infrastructure development and investment promotion to strategies that promote sustainable, people-centered development. Porter described this transition in four phases: production, investment, Innovation and ultimately quality of life.

3 Methodology

A concise and robust methodological framework is essential to comprehensively understand the competitiveness of urban districts. This study uses a composite indicator approach that integrates socio-economic and spatial variables to create a consistent evaluation model [9]. By applying clear procedures such as indicator selection, data normalization, category aggregation, and the creation of a comprehensive competitiveness index, it ensures comparability between different intervals and establishes an analysis method that is transferable to other urban environments [28, 29]. This well-structured approach

serves as a transparent analysis tool that supports evidence-based planning and helps highlight regional strengths and challenges.

While the composite-index approach is established in the literature, this study contributes by operationalising a district-level competitiveness assessment in a way that is (i) transferable across cities, (ii) comparable through joint standardisation of all districts from both case studies, and (iii) statistically streamlined by screening multicollinearity to limit double counting before aggregation. In addition, the workflow is designed to be reproducible and GIS-ready, enabling planners to visualise intra-urban disparities and relate them to policy-relevant dimensions (Human & Society, Economy, Urban Structure).

The set of 21 indicators used in this study derives from a broader initial pool of 35 indicators identified through an extensive review of the scientific literature on urban competitiveness and multidimensional territorial assessment. The original set was designed to capture the full spectrum of social, economic, and spatial dimensions typically considered in composite index methodologies. However, during the data collection phase, several indicators had to be excluded due to the lack of reliable, systematically available data at the suburban district scale used for the analysis. For some variables, data were either not collected by official statistical agencies at this level of detail, presented inconsistencies across the two case-study cities, or were available only at coarser administrative scales, preventing meaningful intra-urban comparison.

The normalisation of indicators was necessary to make characteristics comparable and amenable to aggregation. The methodology used the min-max method (1) because it is applicable to indicators with positive, negative or zero values and because it allows one to widen the variability of indicators lying within a small interval.

It proposes using min–max normalisation to place indicators on a common scale while preserving relative differences. However, min–max can be sensitive to extreme values. This study complements the baseline results with a robustness check using an alternative normalisation (e.g., z-scores or trimmed min–max) to verify that the main spatial patterns are not driven by outliers.

$$y_{SCi} = \frac{x_{SCi} - \min(x_{SCi})}{\max(x_{SCi}) - \min(x_{SCi})} \tag{1}$$

S indicates the statistical unit, C the category and i the indicator.

The distances obtained through normalisation represent the absolute measures of the gap between each statistical unit and the "ideal" one. The methodology considered indicators that negatively affect competitiveness.

For the finalisation of the dataset, a correlation analysis is also necessary to evaluate relationships among indicators and assess their impact on the overall structure of the indicator set (Table 1).

We proceeded with aggregating normalised indicators across the different categories. In the literature, there are numerous criteria for weighting and aggregating variables, ranging from ex ante systems of weights to criteria that infer the meaningfulness of indicators from data analysis (e.g., using multivariate statistical analysis). However, these aggregation methods sometimes implicitly involve a degree of subjectivity. We did not develop a system of weights, as this paper is a first approach to the research. Furthermore, we wanted to develop an innovative application at the district scale. Therefore,

Table 1. The set of indicators included in the methodology.

Category	ID	Indicator	Source
Human & Society	HS1	Population density	INE, 2023
	HS2	Population with Education superior/Total population	INE, 2023
	HS4	Average age of the population	IGE, 2023
	HS5	Number of pensioners/Total population	IGE, 2024
	HS7	Population under 15 years old/Total population	IGE, 2024
	HS8	Voters	IGE, 2023
	HS9	Asociationes/total population	IGE, 2023
Economy	E2	Persons affiliated to the Social Security/Active population	IGE, 2023
	E5	Business density	IGE, 2023
	E6	Beds per tourist dwelling	INE, 2025
	E8	People affiliated with Social Security/Active population	IGE, 2023
Urban structure	U1	Percentage of Buildings surface	IGE, 2023
	U4	Percentage of Green areas surface for each district	Uban Atalas, 2018
	U5	Percentage of Industrial buildings	IGE, 2023
	U6	Percentage of Commercial buildings	IGE, 2023
	U7	Number of Schools/Population 5–19	IGE, 2025
	U8	Health services density	Sergas, 2025
	U10	Land use mix	Uban Atalas, 2018
	U11	Roads density	OSM, 2025
	U12	Bus/train stops density	OSM, 2025
	U13	Cycling and footway density	OSM, 2025

we put forward the hypothesis of interchangeability between dimensions, giving equal importance to different indicators.

Hence, we used simple averages to calculate a partial indicator of competitiveness for single categories (SC) (2). This operation is conceptually equivalent to putting all indicators on an equal footing.

$$M_{SCj} = \frac{x_{SC1} + x_{SC2} + \cdots + x_{SCn}}{n} = \frac{\sum_{I=1}^{n} x_{SCI}}{n} \tag{2}$$

Subsequently, in order to obtain a general measurement of competitiveness, we proceeded with the aggregation of the M_{scj} indexes, which represent the competitiveness of statistical units within single categories. We obtained a composite index of district

competitiveness based on their performance in sustainable development.

$$I_{UCSDG} = \frac{M_{S1} + M_{S1} + \cdots + M_{SM}}{m} = \frac{\sum_{C=1}^{m} M_{SC}}{m} \quad (3)$$

While Mscj indicates the level of competitiveness in the C category, IUCSDG shows the level of competitiveness of the S district, taking into account all the categories of the model. The results can be represented in GIS and in radar graphs.

Because min–max normalisation may amplify outliers, we implement a robust variant using a trimmed min-max approach (e.g., the 5th–95th percentiles) and verify whether district rankings and spatial clusters remain consistent. This check ensures that the observed centre–periphery gradients are not artefacts of extreme values.

4 The Case Studies

The cities of Vigo and A Coruña play central and complementary roles in Galicia, serving as key administrative, cultural, economic, and logistical hubs. Their differing spatial configurations are also reflected in their demographic and territorial profiles. Vigo, which covers 109.1 km^2 and has a population of over 293,000 inhabitants, is characterised by a diversified urban structure that combines a dense central area with extensive suburban districts. This structure forms the backbone of a vibrant city, supported in particular by a strong economy based on industry, logistics and shipping. Although the surrounding areas are more extensive, they are functionally connected to the city centre by a well-developed transport network, thus preserving the unity of the urban system.

A Coruña, on the other hand, has a remarkably high population density, despite its small area of 38.64 km^2, with over 249,000 inhabitants, making it the second largest city in Galicia after Vigo. Its urban structure is more extensive and multicentric, with a densely populated, economically vibrant centre bordered by peripheral areas characterised by port and industrial zones. This pronounced differentiation between centre and periphery adds complexity to the city while strengthening its role as the hub of the most significant metropolitan area in northern Galicia.

Due to their contrasting characteristics in terms of area and population structure, Vigo and La Coruña provide an excellent test field for the proposed multidimensional analysis method. By comparing the differences in size, population distribution and internal spatial structure of the two cities, competitiveness in different urban environments can be comprehensively evaluated and the reliability of the overall index developed in this study can be increased (Fig. 1).

When applying this method to the two target cities (9 constituencies in Vigo and 11 constituencies in La Coruña), the constituencies were chosen as the spatial analysis unit. This choice was made because socio-economic and spatial data are consistently available at the constituency level, these units have suburban characteristics and have an appropriate size to capture differences in internal competitiveness while ensuring the comparability of the analysis.

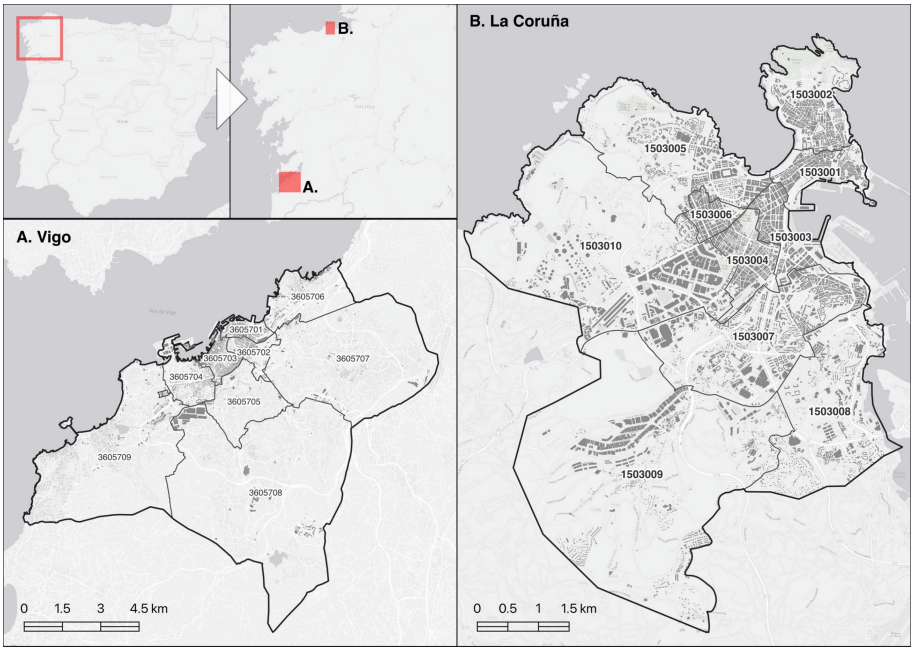

Fig. 1. The municipal electoral district subdivisions of Vigo and La Coruña in the Galicia region, Spain.

5 Measuring Urban Competitiveness

In this analysis, the Index was standardized using all constituency data from both cities to bring all values to a common scale. This method allows direct comparison and clearly highlights the relative differences in competitiveness between the two urban environments. Figure 2 shows the spatial distribution of the standardized competitiveness Index for each constituency in Vigo (a) and A Coruña (B). Based on the correlation analysis, variables with correlation coefficients above 0.7 were excluded. Specifically, the indices HS4, U6, U7, U11, and U13 exhibited strong multicollinearity, which could affect the significance of the overall Index. Therefore, they were removed from the final dataset.

In Vigo, competitiveness values show a clear centre-periphery gradient: the central districts display higher scores, reflecting stronger economic activity, better service provision, and a more consolidated urban fabric. Moving towards the outer suburban areas, the index decreases progressively, indicating more dispersed settlement patterns and comparatively lower concentrations of socioeconomic opportunities. This pattern highlights the role of the compact urban core as the main driver of the city's competitive performance.

In A Coruña, the spatial configuration is more heterogeneous. The central districts also exhibit higher levels of competitiveness, but the variation across peripheral areas is less uniform than in Vigo. Some outer districts maintain relatively positive scores, likely due to the presence of mixed-use environments, coastal accessibility, or consolidated neighbourhood services, while others show lower values associated with more

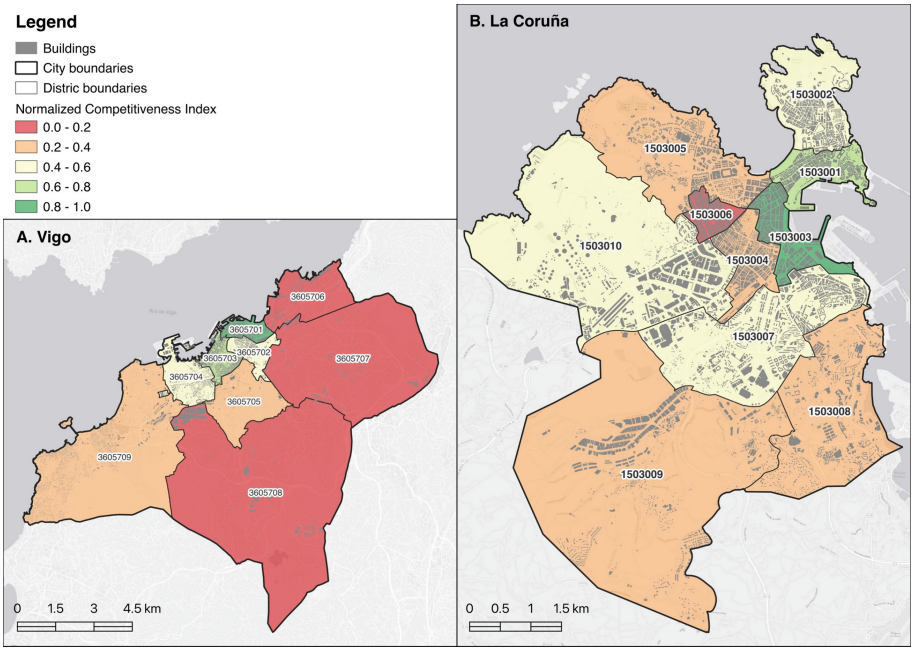

Fig. 2. The value of the synthetic competitiveness index for the municipal electoral districts of Vigo and La Coruña in the Galicia region, Spain.

fragmented urban morphologies. This internal diversity suggests a more polycentric structure, where multiple nodes contribute to the city's overall competitiveness.

Overall, the maps effectively show how the composite index captures intra-urban differences in competitiveness, emphasising how spatial structure, service distribution, and socio-economic conditions interact differently in the two cities.

The bar charts (Fig. 3) for Vigo and A Coruña highlight distinct patterns in the distribution of competitiveness indicators across the three analytical dimensions (Human & Society, Economy, and Urban Structure). Overall, the two cities display heterogeneous profiles, reflecting their different demographic structures, economic dynamics, and spatial configurations.

In Vigo, the standardized values for Human and social indicators (especially HS2 and other indicators on population structure and social participation) are relatively high, indicating excellent results in terms of education level and population distribution. Social indicators such as HS5 and HS9, on the other hand, remain at a medium level, indicating internal disparities between neighborhoods. Economic indicators show a complex picture: employment-related variables (such as E2) are relatively good, while indicators on industrial density and tourism potential are lower, indicating differences in economic power between neighborhoods. Significant internal differences can also be observed in the urban structure. Indicators such as the building area ratio (U1), green spaces (U4), industrial buildings (U5), use diversity (U10) and the number of public transport stops

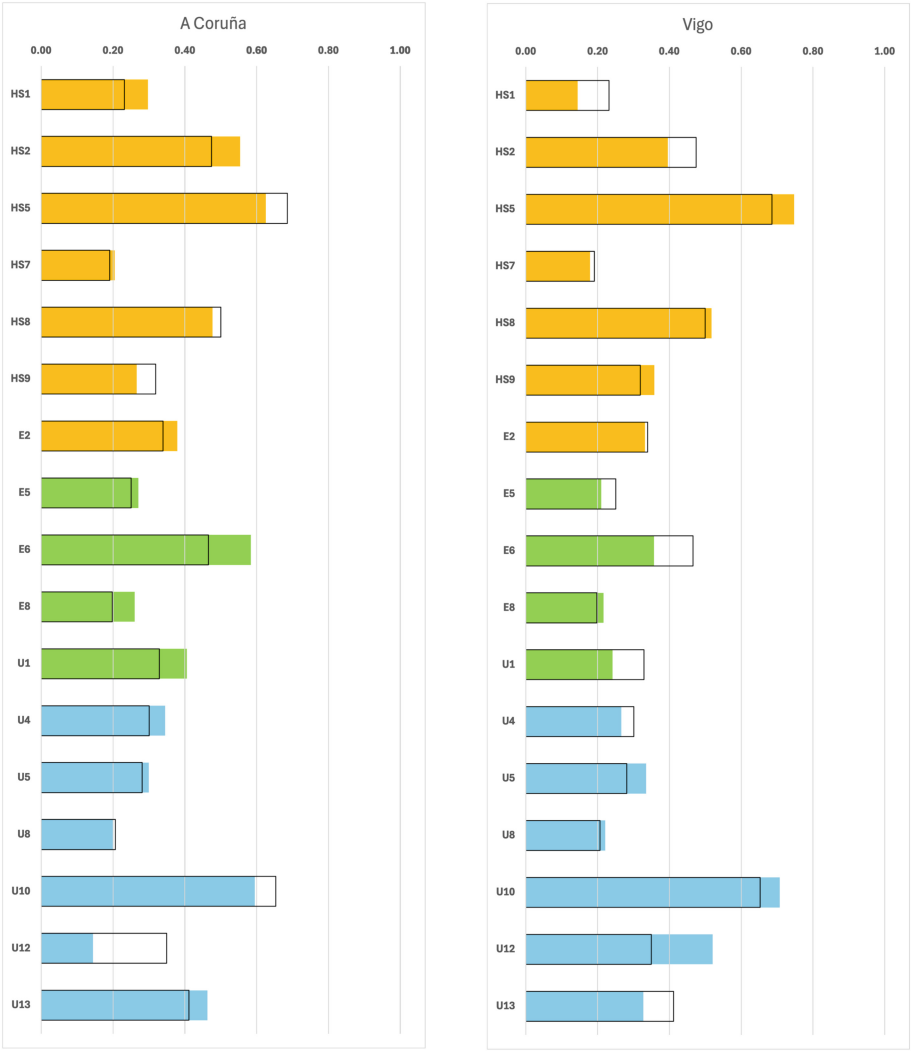

Fig. 3. The values of the competitiveness indicators and their comparison with the weight values for the municipal electoral districts of Vigo and La Coruña in the Galicia region, Spain.

(U12) show, that the centre is more integrated, while the urban structure is more fragmented in the periphery By excluding indicators on road network density and cycling infrastructure Correlation analysis has enabled a more concise but informative structural profile to be drawn up In A Coruña, the balance between human and social aspects is better, but the overall values are slightly lower than in Vigo, especially in terms of education level (HS2). Indicators such as the proportion of the population under 15 (HS7) and the number of eligible voters (HS8), on the other hand, show good values in several districts, indicating active civic participation and a healthy age structure of the urban population.

Economic indicators are more diverse than in Vigo, with a relatively high industrial density and tourism indicators in some districts, which reflects a compact urban structure and the concentration of economic activities Regarding the urban structure, A Coruña has a higher internal consistency than Vigo. Indicators such as U1, U4, U5 and U10 are relatively stable across the city districts, indicating a multi-core and evenly developed cityscape. The smaller difference between the center and the peripheral area seen in A Coruña 's indices could be partly due to the spatial arrangement in connection with the port. in addition to the historic central port, a second large port forms just outside the city limits (created to address the access problems caused by an oil refinery). A functionally dual center structure, which influences urban development. This dynamic must be seen in the broader institutional context of the Spanish metropolitan region. In the absence of a formal regulatory or administrative framework, each municipality acts independently despite its function as an integrated metropolitan system. As the examples of Vigo and A Coruña show, this administrative division makes it difficult to coordinate services and infrastructure between municipalities.

The higher demographic dynamism observed in the city of A Coruña is likely associated with recent immigration flows. Given the city's ageing native population and persistently low fertility rates, population renewal increasingly depends on international migrants. In northern Spain, where immigration tends to correlate with local economic vitality, this dynamic is captured indirectly by indicators of the youth population included in the Human & Society dimension.

When comparing the two cities, the charts clearly illustrate that Vigo shows more pronounced contrasts between central and peripheral districts across several indicators, while A Coruña exhibits a more uniform but less polarised structure. These differences align with the broader spatial organisation of the two cities and confirm the value of the composite index in capturing intra-urban disparities arising from demographic, economic, and morphological characteristics (Fig. 4).

An analysis of the top-performing constituencies in Vigo and A Coruña shows different but comparable patterns in the urban competitiveness of both cities. In both cities, the top-performing constituencies consistently show high values on key indicators such as economic strength, accessibility, and quality of urban services, indicating a balanced and robust regional structure. The top performing constituency in A Coruña is characterized in particular by socio-economic indicators (HS series) and urban services (U series), indicating a mature urban structure with efficient infrastructure and favorable demographic development. The top performing constituency in Vigo, on the other hand, shows a stronger concentration of high values in indicators of mobility, economic activity and urban environment, which reflects the dynamic economic structure of the city and the strategic importance of its city center. These constituencies largely meet the criteria for evaluating urban competitiveness, demonstrating that their strengths are closely linked to the dimensions of the composite indicators. Overall, this analysis confirms the effectiveness of the indicators in capturing differences within cities and shows that certain population groups in both cities emerge as clear centers of competitiveness within the urban system.

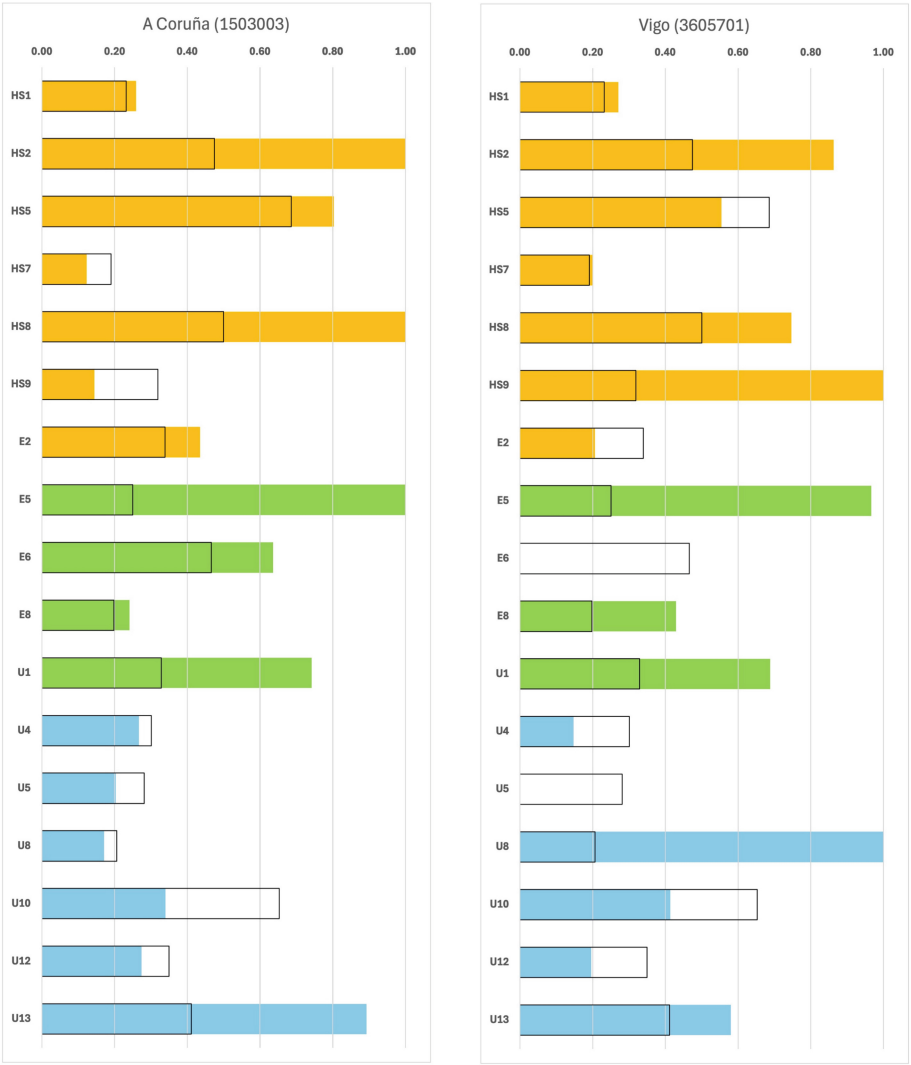

Fig. 4. The values of the competitiveness indicators and their comparison with the corresponding weights for the two best municipal electoral districts in Vigo and La Coruña in the Galicia region, Spain.

6 Conclusions

This study develops and applies a multi-layered conceptual framework to measure urban competitiveness at the neighbourhood level in the Spanish cities of Vigo and A Coruña. It shows how a composite index approach can effectively capture the structural, socio-economic and spatial diversity that characterizes the competitiveness of modern urban systems. The methodology, which integrates 21 indicators in three dimensions-human capital/social, economic and urban structure-provides a transparent and reproducible

analytical tool that highlights both the diversity within each city and the differences due to their respective location conditions. The results show clear Trends in both cities. their more established urban structure and stronger socio-economic base However, while Vigo exhibited a significant discrepancy between the centre and the periphery, A Coruña exhibited a multicore and internally more balanced structure in which several surrounding neighbourhoods achieved comparable competitiveness scores as the central neighbourhoods. These results demonstrate the power of neighbourhood-level analyses to uncover urban asymmetries that were not apparent in more comprehensive regional analyses. in addition to its empirical contributions, this study provides methodological added value by showing how composite indicators at the district level can support the development of more differentiated and evidence-based urban development policies by identifying local strengths and weaknesses. By identifying local strengths and weaknesses, planners and decision-makers can focus their efforts on areas where demographic weaknesses, Supply bottlenecks or low economic activity affect competitiveness. At the same time, the results highlight the importance of incorporating spatial structure into competitiveness analysis, as the diversity of building forms, use mix, accessibility and distribution of public services significantly affect local performance. In this sense, these composite indicators support a more comprehensive and sustainability-oriented Interpretation of competitiveness and are in line with the evolving multidimensional perspective in contemporary urban research.

Although this study makes an important contribution, it also has some limitations. First, the analytical framework is based on the availability of county-level data, which necessitated the exclusion of some of the original indicator sets. This may have limited the thematic coverage of the analysis, especially in areas where there were no consistent county-level data sets, such as indicators on Innovation, environmental performance and digital infrastructure. Second, the decision to assign equal weights to all indicators, while understandable given the first exploratory application, leads to a simplification that may not fully reflect the relative importance of each competitive component. Third, while Minimum-Maximum normalization is a common method, susceptible to extremes and can amplify differences caused by outliers. These limitations open up different approaches for future research. Firstly, expanding the indicator set (especially by adding variables from the environment, digitalization and Innovation) could increase the complexity of the Index. Secondly, future applications could explore alternative weighting methods, such as expert weighting, data-driven approaches such as principal component analysis and factor analysis, and hybrid models that reconcile statistical robustness with policy relevance. Thirdly, longitudinal analyses should be developed to identify Trends, Fourthly, the extension of this methodology to other urban contexts inside and outside Spain allows for more comprehensive comparative analyses and strengthens the external validity of the Index Finally, the Integration of competitiveness results with spatial analyses in a GIS environment (such as cluster analysis and spatial autocorrelation) provides deeper insights into the regional logic that shapes competitiveness and contributes to more spatially precise planning. Overall, this study lays the foundation for a scalable and adaptable analytical framework and deepens our understanding of competitiveness in cities by highlighting the structural, socio-economic and spatial factors that contribute to differences in regional performance.

Indicator availability at the district level may bias the index toward dimensions that are better measured (e.g., accessibility and built form) while under-representing innovation and environmental performance. Similarly, joint standardisation improves cross-city comparability but may compress within-city variance when one city exhibits more extreme values. Finally, indicator screening reduces multicollinearity but can remove conceptually meaningful variables; therefore, we report excluded indicators and correlations in an appendix and provide sensitivity checks to confirm that substantive interpretations are stable.

This study contributes to both academic discussion and practical policymaking and provides a tool to support more equitable, sustainable and contextual urban development strategies.

References

1. World Economic Forum: The competitiveness of cities. World Economic Forum (2014). https://www.weforum.org/publications/competitiveness-cities/
2. World Bank: Competitive cities — Facilitating economic growth and job creation by city and subnational governments. World Bank Group (2025). https://www.worldbank.org/content/dam/Worldbank/document/Trade/CompSector_CompCities.pd
3. UN-Habitat & Economist Impact: Urban Performance Index – Key Findings Report (2024). https://impact.economist.com/sustainability/urbanperformanceindex/downloads/UN-H_UPI-KeyFindingsReport-2024.pdf
4. Porter, M., Warner, A., Sachs, J.: Global Competitiveness Report 2000. McGraw-Hill (2000)
5. Begg, I.: Cities and competitiveness. Urban Stud. **36**(5–6), 795–809 (1999). https://doi.org/10.1080/0042098993222
6. UN-Habitat: Global Urban Competitiveness Report 2019–2020: The world—300 years of transformation into city. UN-Habitat (2020). https://unhabitat.org/sites/default/files/2020/10/global_urban_competitiveness_report_2019-2020_the_world_300_years_of_transformation_into_city.pdf
7. Camagni, R., Capello, R.: Regional competitiveness and territorial capital: a conceptual approach and empirical evidence from the European Union. Reg. Stud. **47**(9), 1383–1402 (2013). https://doi.org/10.1080/00343404.2012.681640
8. Camagni, R.: Regional competitiveness: towards a concept of territorial capital. In: Modelling Regional Scenarios for the Enlarged Europe, pp. 33–47. Springer (2008). https://doi.org/10.1007/978-3-540-74737-6_3
9. Sgambati, S., Carpentieri, G., Gargiulo, C.: Measuring urban competitiveness through the lens of sustainability: an application at the urban districts level in the city of Naples (Italy). In: International Conference on Computational Science and its Applications, pp. 93–108. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-031-10592-0_9
10. Song, J., Xu, B.: Evaluation model of urban tourism competitiveness in the context of sustainable development. Front. Public Health **12**, 1396134 (2024). https://doi.org/10.3389/fpubh.2024.1396134
11. Ni, P., Kamiya, M., Guo, J., Xu, H.: Progress in the pursuit of global urban sustainable development goals (SDGs): from the perspective of sustainable urban competitiveness. In: Global Urban Value Chain in History of Human Civilization: Global Urban Competitiveness Report (2020–2021), pp. 85–111. Springer Nature Singapore, Singapore (2024). https://doi.org/10.1007/978-981-97-1402-5_3

12. Komasi, H., Zolfani, S.H., Nemati, A.: Evaluation of the social-cultural competitiveness of cities based on sustainable development approach. Decis. Making: Appl. Manag. Eng. **6**(1), 583–602 (2023). https://doi.org/10.31181/dmame06012023k
13. Kresl, P., Ietri, D.: Urban Competitiveness: Theory and Practice. Routledge (2014)
14. Deas, I., Giordano, B.: Conceptualising and measuring urban competitiveness in major English cities: an exploratory approach. Environ. Plan A **33**(8), 1411–1429 (2001). https://doi.org/10.1068/a33142
15. Nasi, G., Choi, H., Cucciniello, M., Christensen, R.K.: A systematic literature review of city competitiveness. Public Manag. Rev. (2022). https://doi.org/10.1080/14719037.2022.2029550
16. Shen, J.: Urban competitiveness and urban governance in the globalizing world. Asian Geogr. **23**(1–2), 19–36 (2004). https://doi.org/10.1080/10225706.2004.9684110
17. Jiang, Y., Shen, J.: Measuring the urban competitiveness of Chinese cities in 2000. Cities **27**(5), 307–314 (2010). https://doi.org/10.1016/j.cities.2010.02.004
18. Ginevičius, R., Žemaitis, E., Šuhajda, K., Martin, S.: Assessing the global competitiveness of European countries. J. Competitiveness **15**(4), 36–52 (2023). https://doi.org/10.7441/joc.2023.04.03
19. Kresl, P.K., Singh, B.: Competitiveness and the urban economy: twenty-four large US metropolitan areas. Urban Stud. **36**(5–6), 1017–1027 (1999). https://doi.org/10.1080/0042098993330
20. Martin, R., Simmie, J.: Path dependence and local innovation systems in city-regions. Innovation **10**(2–3), 183–196 (2008). https://doi.org/10.5172/impp.453.10.2-3.183
21. Tennøy, A., Gundersen, F., Øksenholt, K.V.: Urban structure and sustainable modes' competitiveness in small and medium-sized Norwegian cities. Transp. Res. Part D: Transp. Environ. **105**, 103225 (2022). https://doi.org/10.1016/j.trd.2022.103225
22. Sgambati, S., Carvalho, L.: Benchmarking urban competitiveness for economic recovery: an application to Porto and Lisbon. J. Place Manag. Dev. **17**(1), 90–110 (2024). https://doi.org/10.1108/JPMD-03-2023-0019
23. Sgambati, S., Gargiulo, C.: The evolution of urban competitiveness studies over the past 30 years. A Bibliometric analysis. Cities **128**, 103811 (2022). https://doi.org/10.1016/j.cities.2022.103811
24. Nasi, G., Choi, H., Cucciniello, M., Christensen, R.K.: A systematic literature review of city competitiveness: a 30-year assessment and future agenda for public administration. Public Manag. Rev. **25**(8), 1562–1586 (2023). https://doi.org/10.1080/14719037.2022.2029550
25. Gargiulo, C., Sgambati, S.: The investments of NextGenerationEU vs urban competitiveness of Italian metropolitan areas. TeMA – J. Land Use, Mobility Environ. **16**(3), 545–564 (2023). https://doi.org/10.6093/1970-9870/10271
26. Charnock, G., Purcell, T.F., Ribera-Fumaz, R.: City of rents: the limits to the B arcelona model of urban competitiveness. Int. J. Urban Reg. Res. **38**(1), 198–217 (2014). https://doi.org/10.1111/1468-2427.12103
27. Duka, J.A.O., Guevarra, F.M.C., Balisacan, M.T., Encarnacion, R.H.: Spatial analysis of urban competitiveness in metro Manila and its periphery: identifying leading and lagging clusters through CMCI indices (2015-2023). In: ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 10, pp. 221–228 (2026). https://doi.org/10.5194/isprs-annals-X-5-W4-2025-221-2026
28. OECD: Handbook on constructing composite indicators: Methodology and user guide. OECD Publishing (2018). https://www.oecd.org/en/publications/handbook-on-constructing-composite-indicators-methodology-and-user-guide_9789264043466-en.html
29. Bruneckienė, J., Guzavicius, A., Cincikaite, R.: Measurement of urban competitiveness in Lithuania. Inzinerine Ekonomika–Eng. Econ. **21**(5), 493–508 (2010) https://epubl.ktu.edu/object/elaba:3938956/3938956.pdf

Mapping Territorial Competitiveness and Circular Economy Through Inclusive Management: A Bibliometric Analysis and Case Study of the CirculEC Project

Madina Yussubaliyeva(✉), Zhadyra Konurbayeva, Madina Rakhimberdinova, Oxana Denissova, Allfiya Zakimova, and Darya Surova

D. Serikbayev East Kazakhstan Technical University, Ust-Kamenogorsk, Kazakhstan
{myussubaliyeva,zhkonurbaeva,mrakhimberdinova,odenisova, azakimova,dsurova}@edu.ektu.kz

Abstract. This article explores how Central Asian states are shifting from resource-based economies to knowledge-based economies. As a crucial prerequisite for regional resilience and sustainable growth, the study investigates territorial competitiveness alongside the principles of circular economy. By incorporating academia into governance structures, the Triple Helix model reshapes the functions of universities, industry, and government as mutually linked actors that drive innovation. To foster regional innovation clusters-particularly in mining and agriculture-universities need to evolve into entrepreneurial ecosystems and strategic intermediaries. To identify thematic developments and clusters for 20202025, we conducted a Scopus-based bibliometric analysis using VOSviewer and Biblioshiny. The real-world application of these concepts is demonstrated through the Erasmus+ CirculEC project (20232026), which seeks to update curricula and strengthen competencies in circular economy in universities in Kazakhstan, Uzbekistan, Tajikistan and Turkmenistan, aligning education with regional labour market needs. The study introduces a new paradigm-"circular territorial competitiveness"-in which sustainable development is embedded in economic policy. Long-term regional growth is closely related to the success of digital and green transformations, which activate scientific capacity and enhance human capital.

Keywords: Territorial competitiveness · Circular economy · Inclusive management · Triple helix model · Bibliometric analysis · Central Asia

1 Introduction

In the new global economy, territorial competitiveness has become a central issue for emerging economies in the Central Asian region, as they navigate the systemic structural transformation necessitated by the inevitable conclusion of the

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 402–418, 2027.
https://doi.org/10.1007/978-3-032-30518-3_26

"oil era". In this context, the competitiveness of a region is no longer defined solely by its economic output, but emerges from the dynamic interaction between spatial development, innovation systems, and circular resource use. For these resource-dependent nations, the adoption of the principles of a circular economy (CE) has emerged as a key strategy to improve economic resilience, reduce dependence on imports of raw materials, and foster regional innovation. Recent bibliographic trends indicate that research on territorial competitiveness and the circular economy in Central Asia has practically merged into a single trajectory focused on sustainable development, where innovation acts as the unifying foundation. This convergence establishes the basis for a new integrated paradigm known as "circular territorial competitiveness," where economic performance and sustainability are co-developed within a unified regional system. However, achieving such a systemic transformation requires a radical departure from conventional business and educational practices to address the persistent "skills mismatch" between academic outputs and the needs of a transitioning "green" labour market. Higher education institutions (HEIs) have consequently assumed expanded roles as central nodes in regional innovation ecosystems, acting as "strategic generators of knowledge" that bridge the gap between academic research, industrial application, and national sustainability goals. Through the implementation of international capacity-building projects like **Erasmus+ CirculEC** (Project No. 101129169), universities are transforming into catalysts for reform, developing the interdisciplinary human capital and "heterogeneous engineers" necessary to drive the regional transition toward a climate-neutral and resource-efficient economy. The project relies on the Triple Helix model of university-industry-government interaction, redefining these actors as collaborative engines of innovation capable of breaking territories free from "low-growth traps".

A critical success factor in this transformation is **inclusive academic management** - a governance approach grounded in distributed leadership, participatory decision-making, and universal design principles. Inclusive governance serves as the essential "bridge" that converts Triple Helix trilateral interactions into sustainable regional outcomes, ensuring that academia remains agile and responsive to labour market signals. This study argues that inclusive management is not peripheral to territorial competitiveness but constitutes its institutional foundation: without the capacity of HEIs to govern inclusively and engage diverse stakeholders, the transition to a circular, knowledge-based regional economy cannot be sustained. The CirculEC project is examined as empirical evidence of this proposition, illustrating how inclusive academic governance translates bibliometrically-identified research priorities into measurable capacity-building outcomes across Central Asia.

2 Theoretical Background

Innovation as a joint product of innovation between universities, industry, and government is the basis for the Triple Helix model and is the primary framework for this study. In this model, they co-create new knowledge, technologies

and regional strategies [13,26,35]. Entrepreneurial universities, as described by Etzkowitz, go beyond traditional education by creating spinoffs and licencing patents. Gleichzeitig notes that governments act as "public entrepreneurs" and venture capitalists, synchronising national strategies with the needs of a knowledge-based economy. Modern scientific discussions significantly expand the classical triple helix system model, interpreting it as a complex innovation system [35] that integrates ecological and social dimensions through the quadruple/quadruple helix concepts [7,48]. This theoretical reasoning aligns with the central emphasis on the intersection between territorial competitiveness and the circular economy in Central Asia, where sustainable development is no longer viewed as an external limitation but is integrated as a foundational component of economic policy. Within the broader shift of the region from a resource-based to a knowledge-driven economy, the Triple Helix framework supports a reimagining of universities not as standalone entities, but as strategic actors that can help territories break free from "low-growth traps" [31]. Nurmukhanova et al. argue that strategic governance of higher education institutions is crucial for territorial competitiveness because these institutions create and sustain entrepreneurial ecosystems within regional business support networks [32]. In Central Asia, this appears in the creation of industrial clusters around regional universities, especially in mining and agriculture. According to Taikulakova et al., these clusters speed up the application of scientific developments and improve resource efficiency in the real economy [42]. Akhmedov et al. emphasise that as resource revenues decrease, it becomes crucial to actively improve competitiveness indicators [1]. The Triple Helix model of universityindustrygovernment interaction facilitates the transition to a knowledge-based economy, in which universities function as the main producers and disseminators of knowledge that underpin regional green and digital transformations. Yessenbek et al. argue that effective mobilisation and utilisation of this scientific potential constitute a critical determinant of national resilience and long-term viability in emerging economies [47]. Inclusive academic management is the key "bridge" to turn these trilateral interactions into sustainable regional outcomes. Reorganises university leadership, policies, and daily practices so that all students and staff can participate and succeed, regardless of disability, age, background, or status. For academic reforms to drive territorial development, the gap between declarations and implementation must be closed. Therefore, researchers propose frameworks at national, regional and institutional levels to align inclusive strategies with curricula [14,33]. Specific strategic roadmaps are proposed to harmonise management practices, universal design principles, and service provision, a process that is critically important for the effective integration of interdisciplinary programmes within the framework of the circular economy [37]. Effective governance that reconciles environmental and economic objectives requires a specific constellation of leadership competencies. High-quality inclusive management requires leaders to integrate pedagogical, psychological, and innovation-related competencies and to employ a democratic leadership style [30]. Management grounded in innovative technologies facilitates participatory governance models, actively engaging faculty members

and subject-matter experts in decision-making processes to ensure timely and effective responses to external challenges [44]. Inclusive management reimagines the university as a "growth pole" able to compete effectively in the global academic and scientific arena. In this sense, inclusive education serves as a strategic catalyst for international cooperation, strengthening and diversifying the external ties of the region and its integration into transnational networks [22]. The systematic application of Universal Design for Learning (UDL) principles in the redesign of academic procedures and policies mitigates alienation experiences and improves the probability of successful outcomes for all stakeholders involved in these processes [9]. An inclusive organisational culture is grounded in explicitly articulated values and a pronounced sense of social responsibility, thus establishing an ethical framework that underpins sustainable regional development. [17,41].

3 Methodology

The methodology of this study incorporates bibliometric analysis and science mapping, specifically employing VOSviewer [10] for keyword co-occurrence analysis and Biblioshiny [4] for thematic mapping and trend identification. The Scopus database by Elsevier was used to collect the research literature on territorial competitiveness and the circular economy in the context of Central Asia, as it is a widely used source for bibliometric analyses in management education and the social sciences [15]. To focus on research related to territorial competitiveness and circular economy in Central Asia, we conducted two searches, concentrating on subject areas within Social Sciences, Business, Management, and Accounting, Economics, Econometrics and Finance and Engineering. The scope covered research in the last five years. Our literature search strategy was implemented in two stages. The first search aimed to find publications that contain the terms "territorial competitiveness," "regional competitiveness" or "economic competitiveness." The second search combined "Circular Economy" AND "Central Asia" AND ("territorial competitiveness" OR "regional competitiveness" OR "economic competitiveness"). We used the PRISMA 2020 framework [34], thereby ensuring that the research design was systematic, transparent, and reproducible. Bibliometric techniques were applied to generate quantitative evidence on the degree of convergence between the investigated research domains. Overall, we reviewed 54 academic works on territorial competitiveness and 69 on the circular economy in Central Asia. This combined methodological approach makes it possible to detect intersecting thematic clusters and facilitates the framing of circular territorial competitiveness as a developing interdisciplinary field of inquiry. The study goes beyond a bibliometric assessment by also using an exploratory single case study design to obtain an in-depth view of how theoretical concepts are implemented in practice. The case examined concerns the implementation of the project "Development of innovative curricula and modules in the field of circular economy and sustainable development," a three?year initiative (20232026) designed to strengthen human capital in universities in Kazakhstan, Tajikistan,

Uzbekistan, and Turkmenistan by embedding circular economy (CE) principles into academic programmes. The project is directed towards updating curricula and fostering interdisciplinary skills that align with global sustainable development priorities as well as the demands of regional labour markets. The CirculEC project is selected as an illustrative case because it reflects the integration of circular economy principles, innovation systems, and regional development strategies in Central Asia. This approach enables the analysis of complex interactions between institutional stakeholders and supports the validation of the proposed concept of circular territorial competitiveness.

4 Results and Discussion

4.1 Trend Topics

Territorial competitiveness is a decisive factor for the survival of emerging economies in the context of global uncertainty and the structural transformation of world markets [18,47]. Territorial competitiveness is a key tool for Central Asian countries in overcoming resource dependence and diversifying their economic structure [1,18].

A comparative bibliometric analysis was conducted using Biblioshiny, and Figs. 1 and 2 present the main trends (20202025) identified regarding Territorial Competitiveness and the Circular Economy in the context of Central Asia.

The topics of trend in territorial competitiveness evolved from the concept of innovation moving to research on competitiveness indices to tools for regional planning and sustainable development. Although Krammer and Jimenez noted that the institutional environment defines the type of innovation in the region [16,24] emphasise a stable influx of investment to small and medium enterprises. Androniceanu et al. used correlation analysis to demonstrate a strong and direct link between innovation, product competitiveness, and the well-being of the population [3]. Their research indicates that countries investing in research and development and human capital tend to have higher GDP levels and a better quality of life for their citizens. Mukhametzhan et al. created an econometric model to assess urban development disparities in Kazakhstan, revealing that spatial inequality and wealth concentration in major cities such as Almaty and Astana significantly affect regional economic growth [28]. Akhmedov et al. focused on examining the key parameters of the Global Competitiveness Index (GCI) for oil-producing nations such as Kazakhstan. They concluded that, with the end of the "oil era,"institutions, higher education, market efficiency, and business sophistication are essential for competitiveness [1]. Nurmukhanova et al. emphasise the importance of strategic management in universities as a crucial element in territorial development, advocating for universities to establish entrepreneurial ecosystems integrated into the regional business support system [32].

The concept of green competitiveness connects a company's performance with stakeholder engagement and the incorporation of environmental factors into development policy. Meanwhile, Chechelashvili et al. underscore the need

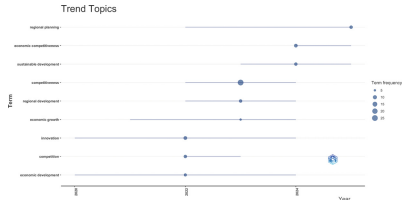

Fig. 1. Trend topics on Territorial Competitiveness.

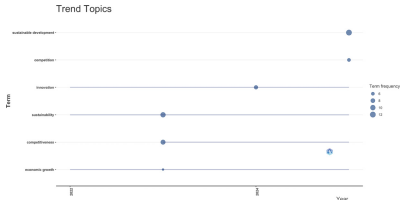

Fig. 2. Trend topics on Circular Economy.

to transition from tactical to strategic management of social responsibility, incorporating ESG principles into corporate strategies to ensure sustainable growth [8]. Muratbekova conducted a comprehensive analysis of digital transformation in Central Asian countries, highlighting that regional governments need to increase infrastructure investment and eliminate connectivity barriers to accelerate digital advancement [29]. The Enterprise Digital Competitiveness Index (EDCI) was developed, showing that nations with the highest levels of business digitalisation rank among the top five in GDP per capita [38]. In their 2025 study, Taikulakova et al. concentrate on strategic planning through the creation of industrial clusters around universities. This partnership accelerates the commercialisation of scientific innovations and significantly improves the efficiency of production resources [42]. Recent research by Yessenbek et al. considers the mobilisation of scientific potential as a crucial factor for economic survival amid global changes [47]. The authors confirm a positive relationship between the regional RandD and national competitiveness indicators. The traditional analysis of territorial competitiveness in Central Asian countries is shifting from monitoring macroeconomic indicators to the search for new models of sustainable growth [1,36]. This transition being carried out through the integration of the principles of circular economy as a key driver of "green"competitiveness, where production efficiency and environmental responsibility in business become mandatory elements of the regional strategy [27]. Green taxation is examined as a key instrument for steering EGG management in environmentally orientated companies ("green") in the region, taking into account the need to revise management strategies in response to the imperatives of decarbonisation and the promotion of a circular economy within the EAEU [5]. Ergasheva recognises the use of digital technologies to automate "green"accounting [11]. Finally, integrating regional higher education systems to train specialists in the fields of digital and circular economy is widely discussed [45].

4.2 Thematic Map

To identify the motor, basic, niche and emerging/declining themes of the research from 2020–2025, we performed an analysis and created a thematic map using Biblioshiny, as shown in Fig. 3 and Fig. 4. The analysis of the thematic map revealed the primary themes: Economic Competitiveness and Innovation, as well

as International Trade and Regionalisation. Research proves the direct connection to innovation, Research confirms a direct link between innovation, product competitiveness, and the well-being of the population [3]. The mobilisation of scientific potential is considered a decisive factor for the survival of economies during the transition to a knowledge-based economy [47]. Stratan et al. emphasise the relationship between the innovation environment, productivity, and the development of infrastructure [40]. The study examines trade patterns in specific sectors, such as automobile manufacturing [39], as well as the impact of geopolitics on major transit corridors such as the Belt and Road Initiative [46]. Territorial competitiveness is considered a factor in improving production culture in integrated sectors of the economy [43] and a tool in the regional tourism industry [6].

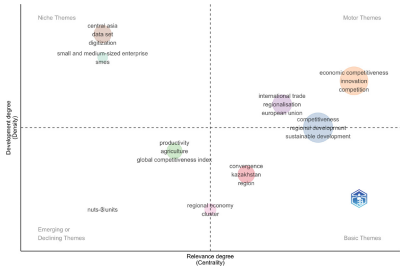

Fig. 3. Thematic map of the research literature on Territorial Competitiveness.

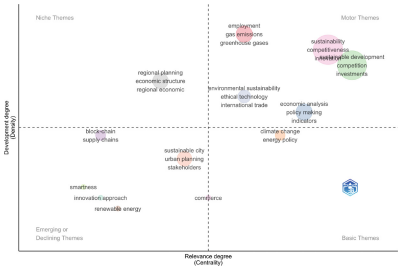

Fig. 4. Thematic map of the research literature on Circular Economy in Central Asia.

The second set of themes covers regional development and inequality, sustainable development, and clusters and convergence. Territorial competitiveness is examined through spatial inequality and wealth concentration in cities [36]. Competitiveness planning now routinely integrates economic goals with social responsibility and environmental safety [23]. Studies address the asymmetry of urban development in Kazakhstan [28], the role of entrepreneurial universities [32], and the use of deep learning technologies [21].

The research of Andekina et al. occupies a central place in circular economy research in Central Asia focusing on human capital analysis, substantiating the transformation of Kazakhstani universities into entrepreneurial universities, which serve as strategic generators of knowledge for the transition to new economic models [2]. The author views universities as catalysts for reform, which, through the implementation of international projects (such as Erasmus+ capacity bulding in higher education), introduce interdisciplinary programmes on the circular economy and cultivate an adaptive workforce. According to Lambekova and Andekina, creating such an innovative environment within higher education institutions is a necessary condition to ensure the territorial competitiveness of Central Asian countries [25]. Nosirov links macroeconomic instruments (taxes)

with micro-level innovations (ESG), which serves as the foundation for enhancing territorial competitiveness through the lens of the circular economy [5]. The financial stability of companies [12] and the stability of public finances [43] create the necessary conditions for the implementation of territorial competitiveness strategies and the transition to a circular economy in the region. In Central Asia, the emphasis is on the creation of cluster structures in the food and agriculture sectors to implement circular technologies [19,20]. By 2024–2025, research on territorial competitiveness and the circular economy in Central Asia has converged into a unified sustainable development trajectory, with innovation as the common foundation. Both fields have shifted toward applied regional planning: territorial competitiveness as a tool to overcome resource dependence, and the circular economy as the primary diversification mechanism. Key instruments include the GCI for benchmarking institutional constraints [1], digital competitiveness indices [38], green taxation within the EAEU [5], and university-led industrial clusters in agriculture and mining [42].

Table 1 presents a synthesis of Territorial Competitiveness and Circular Economy Development in Central Asia.

Table 1. Synthesis of Territorial Competitiveness and Circular Economy in Central Asian Development

Characteristic	Territorial competitiveness	Circular economy
Driver	Innovation and SMEs	Innovation and SMEs
Current focus	Sustainable Development (ESG)	Sustainable Development (ESG)
Aim in Central Asia	Diversification of the economy	Overcoming dependence on natural resources
Implementation tool	Implementation tool	System Resource Management

The comparison demonstrates that territorial competitiveness and the circular economy are structurally interconnected and mutually strengthening. Their convergence creates the foundation for an integrated development model, where economic diversification, resource efficiency, and innovation jointly contribute to sustainable regional transformation. This alignment further supports the emergence of the concept of circular territorial competitiveness. The intersection of territorial competitiveness and circular economy forms the basis of a new integrated paradigm, circular territorial competitiveness, where economic performance, sustainability, and innovation are co-developed within a unified regional system.

In addition to examining which topics are most relevant to regional competitiveness and circular economies, as well as how interest in these topics has changed over time (trend analysis), it is crucial to comprehend how these topics interact. By using VOSViewer to analyse co-occurrence, we were able to group them into stable thematic blocks, clusters. In contrast to trend analysis, which shows “when” and “what” became popular, Fig. 5 and Fig. 6 reveal the internal structure of research: which ideas always go together and form the basis for regional development today.

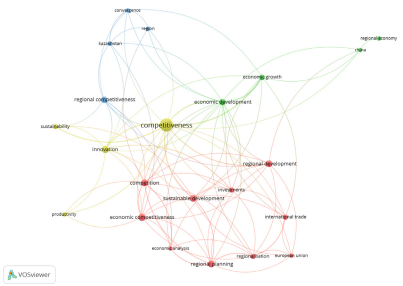

Fig. 5. Map of co-occurrence of author keywords on research territorial competitiveness in Central Asia.

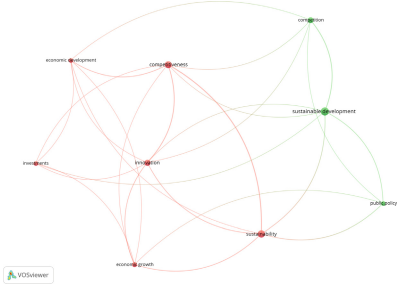

Fig. 6. Map of co-occurrence of author keywords on Circular Economy in Central Asia.

To cluster current research on territorial competitiveness, we made analysis of author keywords co-occurrence using VoSviewer and created map. With minimum occurrence of author keywords kept at 6. Table 2 Summary of identified clusters in the field of territorial competitiveness (Fig. 7 and Table 3).

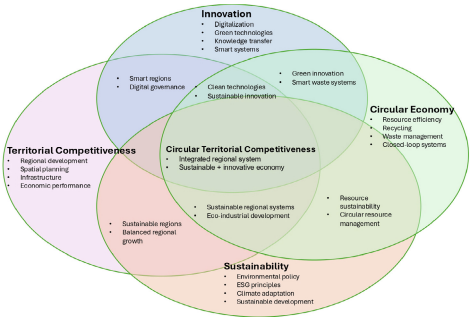

Fig. 7. Cluster Overlap and Knowledge Convergence. Source: Own elaboration.

The overlap of research clusters identified through bibliometric analysis provides strong evidence of an ongoing transformation in the knowledge structure of the field. The blurring of the boundaries between thematic areas signals the emergence of a new interdisciplinary domain, where territorial competitiveness is increasingly shaped by the integration of circular economy principles, innovation systems, and sustainability frameworks. This convergence reflects a transition from fragmented research perspectives to a holistic and system-orientated understanding of regional development.

Table 2. Summary of identified clusters in the field of territorial competitiveness

Cluster	Analysis/Remarks
Regional planning and sustainable development	Researchers emphasize that environmental responsibility and the implementation of ESG principles are becoming mandatory elements of regional competitive strategy. Within this area, the processes of regionalization and the impact of international trade on the competitiveness of territories are also studied, including the experience of the European Union
Competitiveness and innovation	Research confirms that a region's capacity for innovation directly determines the population's welfare and overall productivity. Special importance in this cluster is given to mobilizing scientific potential and to the role of "entrepreneurial universities" as the main generators of knowledge for business
Macroeconomic dynamics and the influence of external markets	In this field, the factors ensuring stable economic growth and development are studied. An important aspect is the role of China and its influence on the regional economy through initiatives such as the "One Belt, One Road". Scholars emphasize the need for balanced development of infrastructure and productivity to maintain positive momentum in national economies amid the global transformation of markets.

4.3 Case Discussion: Practical Implementation of the "Triple Helix" Model and Inclusive Academic Governance: The CirculEC Project Case for Enhancing Territorial Competitiveness. The initiatives of the European Union to modernise higher education systems in third countries are associated with the European Green Deal and the United Nations Sustainable Development Goals (SDGs) and emphasise projects advancing transition to circular economy. Capacity building in higher education action aims to transfer knowledge to build a "green" workforce and to bridge academia and the labour markets to market skills. The "Development of Innovative Curricula and Modules on Circular Economy and Sustainable Development" project (CirculEC) (2023–2026) is being implemented by the Consortium of 10 leading universities in Kazakhstan, Uzbekistan, Tajikistan, and Turkmenistan and 3 European universities from Germany, Italy, and Spain since 2023. The project aims to build the capacity for human capital in HEIs of Central Asian countries in a circular economy to provide staff, students, and the general public with skills and competencies aligned with the needs of local labour markets and environmental challenges faced by CA countries. The principles of circular economy are integrated into the undergraduate, graduate degree programmes and life-long learning programmes and are adapted to the needs of Kazakhstan, Uzbekistan, Tajikistan and Turkmemistan. The project relies on collaboration of academia, industry and government entities, which provides approach of Triple Helix model. The universities in Central Asia aimed to become regional innovation ecosystems driving

Table 3. Summary of identified clusters in the field of Circular Economy in Central Asia

Cluster	Remarks
Innovation cluster	The most closely connected keywords deal with the innovative-economic cluster of territorial development. The research in this cluster shows the direct relationship between technological progress, productivity, and the population's standard of living. They analyze business sustainability as an internal driver of competitiveness, where investments and the adoption of ESG principles become the foundation for long-term growth.
Sustainable development cluster	Research in this cluster deals with institutional frameworks and strategic planning. The research literature explores mechanisms for implementing sustainable development goals through tools of government regulation and regional planning (Kapesa, 2025; Muratbekova, 2024). Authors evaluate how effectively the Triple Helix model interacts and the universities as strategic drivers capable of leading regions out of "low-growth traps" (Andekina et al., 2024; Nurmukhanova et al., 2021; Taikulakova et al., 2024).

regional competitiveness through implementing new courses adopting the principles of circulation economy, building the capacity of academic staff, preparing innovative decisions for the demands of the real private sector during regional hackathons, disseminating knowledge and sustainable development. CirculEC provides a comprehensive market analysis that was incorporated into 27 newly developed courses comprising modules on Environmental Management, Sustainable Competetiveness, and Environmental Entrepreneurship. The modules cover resource efficiency and environmental risk management, market value, transformative management, and sustainable business models. We achieved a wide range of inclusion by updating undergraduate, graduate, and lifelong learning programmes where the 27 courses were interrelated. The needs analysis involved 130 stakeholders from industry, government entities, and academia to meet the demands of the labour market in mining, agriculture, and construction in the domain of regional economies. One of the examples of an engaging private sector company is Izoplus Central Asia LLP that actively participated in workshop validation, content design, and participating in European trainings on their own expense. The core initiative of the project is to build capacity in higher education in Central Asia by transfer of knowledge from European partner universities. The capacity building initiative included training of 78 academic staff from 10 Central Asian universities. The format of "Train the Trainer" workshops ensures that trained trainers act as local disseminators of knowledge and uphold content quality standards. Post-training evaluations confirmed high perceived relevance, with participants reporting increased confidence in delivering circular economy topics and readiness to contextualise materials for their institutional settings.

Table 4 maps the bibliometric clusters identified through VOSviewer analysis (Figs. 5–6) to the three core module groups of the CirculEC project, illustrating how curriculum design is grounded in the state-of-the-art literature.

Table 4. Mapping of Bibliometric Clusters (Figs. 5–6) to CirculEC Curriculum Modules

Bibliometric Cluster	CirculEC Module Group	Rationale for Alignment
Innovation cluster (TC): technological progress, productivity, ESG investment	Sustainable Competitiveness	Cluster themes on innovation-driven competitiveness and ESG principles directly informed modules on sustainable business models and market value creation.
Sustainable development cluster (TC): Triple Helix, regional planning, institutional frameworks	Environmental Management	The focus on university–industry–government interaction and strategic regional planning shaped modules on environmental risk management and resource efficiency governance.
Innovation-economic cluster (CE): circular business models, green entrepreneurship	Environmental Entrepreneurship	Cluster keywords on circular business design, green startups, and stakeholder engagement are reflected in eco-auditing modules and regional hackathon activities.

The final phase of the project was to test the developed courses within the undergraduate, graduate and LLL programmes. The original target of 1,100 students was exceeded and 3882 students were reached, namely from Kazakhstan 2008, Uzbekistan 880, Turkmenistan 781 and Tajikistan 213. The implementation of pilot phase followed by the developed guidelines establishing the assessment of interdisciplinary modules, continuous monitoring, making adjustments to align with regional contexts. The visibility of the project is ensured by digital presence on the Internet platform and social media with 1800 users, which is a valuable contribution to interaction and dissemination between industries, government and academia. The project maintains its digital ecosystem to engage the stakeholders when academia, industry, and government are able to continuously interact. The innovative projects developed by graduates of the programmes were presented during regional Hackathons and focused on regional demands such as agro-processing, sustainable construction, circular business models, energy and waste management. This initiative improves the employability of graduates and the rise of environmentally conscious startups, increasing the competitiveness of universities in Central Asia. One of the key elements of the project is the sustainability plan which aims to enhance collaboration internationally, monetise academic expertise to move to self-sustaining development. For long term

sustainability and alignment with regional development goals, the sustainability plan is structured around financial, institutional, and operational dimensions. The project supports the off-budget revenues through the commercialisation of LLL programmes and consulting services in eco-auditing and circular economy to businesses. In the context of Triple Helix model universities act as enterprenuers. As a result, these mechanisms collectively contribute to territorial competitiveness by enhancing innovation capacity, knowledge transfer, and circular solutions that are embedded locally. The CirculEC project illustrates how the Triple Helix model is used to promote sustainable development driven by innovation. The project becomes a scalable model for converting ecological values into economic assets and transforms higher education that shapes a sustainable future.

4.3 Conclusion

The core finding is a thematic convergence: sustainability is no longer an external constraint but an integral part of competitiveness, and inclusive academic management is its institutional driver. Universities in Central Asia are transitioning from isolated education centres to active nodes in regional innovation networks - a transformation made possible by governance structures that are participatory, responsive, and grounded in Triple Helix collaboration. The bibliometric analysis confirms that research on territorial competitiveness and circular economy in Central Asia has merged into a unified sustainable development trajectory, with inclusive HEI governance emerging as the critical enabling condition. The CirculEC project demonstrates that when universities govern inclusively - engaging 130 stakeholders, training 78 staff, and reaching 3,882 students across four countries - the result is not only curriculum reform but a measurable strengthening of regional circular territorial competitiveness. Future research should examine how digital accessibility and inclusive governance frameworks within HEIs can be systematically scaled to further accelerate this regional transition.

Acknowledgments. The CirculEC project is funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

References

1. Akhmedov, E., Mukashev, Y., Akhmedov, A.: Competitiveness of oil-exporting developing and emerging countries and parameters critical for increasing it. Int. J. Bus. Global. **31**(2), 231–249 (2022). https://doi.org/10.1504/ijbg.2022.125978
2. Andekina, R., Rakhmetova, R., Kaliyeva, S., Maxyutova, A.: Entrepreneurial universities of the republic of Kazakhstan: the essence and features of development. Sci. Herald Uzhhorod Univ. Ser. Phys. **55**, 2690–2700 (2024). https://doi.org/10.54919/physics/55.2024.269an0

3. Androniceanu, A., Gherghina, R., Ciobănaşu, M.: The interdependence between fiscal public policies and tax evasion. Administratie si Manag. Publ. **2019**(32), 32–41 (2019). https://doi.org/10.24818/amp/2019.32-03
4. Aria, M., Cuccurullo, C.: bibliometrix: An r-tool for comprehensive science mapping analysis. J. Informet. **11**(4), 959–975 (2017). https://doi.org/10.1016/j.joi.2017.08.007
5. Bogoviz, A.V., Nosirov, J.T., Litvinova, T.N., Ostrovskaya, V.N.: Environmental taxation as an incentive to ESG management of green business in Russia and central Asia. In: Popkova, E.G., Sergi, B.S. (eds.) ESG Management of the Development of the Green Economy in Central Asia, pp. 67–74. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-46525-3_8
6. But, T.: An analysis of the tourism industry competitiveness of Czech regions. J. Competit. **16**(3), 1–25 (2024). https://doi.org/10.7441/joc.2024.03.01
7. Cai, Y.: Neo-triple helix model of innovation ecosystems: integrating triple, quadruple and quintuple helix models. Triple Helix **9**(1), 76–110 (2022). https://doi.org/10.1163/21971927-bja10029
8. Chechelashvili, M., Chikviladze, N., Iashvili, I., Kajaia, T., Kharkhelauri, K., Malania, E.: Creation of a strategically oriented social responsibility management system in a Georgian organization. Qubahan Acad. J. **4**(3), 694–712 (2024). https://doi.org/10.48161/qaj.v4n3a246
9. Davis, M.: Examining and improving inclusive practice in institutional academic integrity policies, procedures, teaching and support. Int. J. Educ. Integr. **18**(1), 14 (2022). https://doi.org/10.1007/s40979-022-00108-x
10. van Eck, N.J., Waltman, L.: Software survey: VOSviewer, a computer program for bibliometric mapping. Scientometrics **84**(2), 523–538 (2010). https://doi.org/10.1007/s11192-009-0146-3
11. Ergasheva, S.: Green accounting and digitalization of E-business for sustainable development. In: Green Accounting and Digitalization of E-Business for Sustainable Development. IGI Global (2025). https://doi.org/10.4018/979-8-3693-8362-9
12. Eshov, M.P., Abdurakhmanova, G.K., Burkhanov, A.U., Abdusalomova, N.B., Ergasheva, S.T. (eds.): Development of International Entrepreneurship Based on Corporate Accounting and Reporting According to IFRS: Part B. Emerald Publishing Limited (2024). https://doi.org/10.1108/S1877-6361202433B
13. Etzkowitz, H.: Innovation in innovation: the triple helix of university-industry-government relations. Soc. Sci. Inf. **42**(3), 293–337 (2003). https://doi.org/10.1177/05390184030423002
14. Filippou, K., Acquah, E.O., Bengs, A.: Inclusive policies and practices in higher education: a systematic literature review. Rev. Educ. **13**(1), e70034 (2025). https://doi.org/10.1002/rev3.70034
15. Ghani, M., Razak, N., Tahir, P., Alshurideh, M.: Sustaining competitive advantage during COVID-19 pandemic: a study of aerospace manufacturing industry in Malaysia. Stud. Comput. Intell. **1056**, 2287–2305 (2023). https://doi.org/10.1007/978-3-031-12382-5_125
16. Gherghina, S., Botezatu, M., Hosszu, A., Simionescu, L.: Small and medium-sized enterprises (SMEs): the engine of economic growth through investments and innovation. Sustainability (Switzerland) **12**(1) (2020). https://doi.org/10.3390/SU12010347
17. Gil-Álvarez, J.L., Morales-Cruz, M., León-González, J.L.: Educational management in the inclusive university. Transformation of their processes from a person-

alized and participatory education. Revista Metropolitana de Ciencias Aplicadas **6**(1), 47–54 (2023). https://doi.org/10.62452/96dn3m03
18. Hasanova, J., Najafova, K.: Research and analysis of opportunities for regional economic integration among the countries of the organization of Turkic states. Econ. Dev. **24**(2), 54–67 (2025). https://doi.org/10.63341/econ/2.2025.54
19. Hernandar, M., Fontana, A., Wijanto, S.: The role of innovation ecosystem in an open innovation model of state-owned enterprises to build collaborative advantage. J. Infrastruct. Policy Dev. **8**(13) (2024). https://doi.org/10.24294/jipd6983
20. Hrosul, V., Kruhlova, O., Mkrtchyan, T., Zubkov, S., Timchenko, O.: Integration of enterprises and formation of cluster structures in the food sector. Agric. Resourc. Econ. Int. Sci. E-Jo. **7**(2), 119–140 (2021). https://doi.org/10.51599/are.2021.07.02.07
21. Issabayeva, S., Katyetova, A.: Deep learning of AI: Kazakhstan's case. In: Arai, K. (ed.) FICC 2025. LNNS, vol. 1285, pp. 586–602. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-84460-7_37
22. Juanzo, C.J.: Inclusive education: the strategic catalyst of internationalization at a Philippine state university. Romblon State Univ. Res. J. **6**(1), 71–76 (2024). https://doi.org/10.58780/rsurj.v6i1.175
23. Kapesa, T.: Collaboration and partnerships: fostering regional cooperation for sustainable development. In: Sustainable Development in Southern Africa: The Critical Role of Public Governance, pp. 239–261 (2025). https://doi.org/10.4324/9781003596011-11
24. Krammer, S., Jiménez, A.: Do political connections matter for firm innovation? Evidence from emerging markets in central Asia and eastern Europe. Technol. Forecast. Soc. Change **151** (2020). https://doi.org/10.1016/j.techfore.2019.05.027
25. Lambekova, A., Shakenova, G., Rakhmetova, R., Maxyutova, A., Andekina, R.: Higher education in the era of the sixth information revolution: digital competencies as a catalyst for economic transformation and business innovation. Procedia Comput. Sci. **272**, 469–474. Elsevier B.V. (2025). https://doi.org/10.1016/j.procs.2025.10.233
26. Leydesdorff, L., Ivanova, I.: 'open innovation' and 'triple helix' models of innovation: can synergy in innovation systems be measured? (2016). https://doi.org/10.2139/ssrn.2791914, https://papers.ssrn.com/abstract=2791914
27. Lyulyov, O., Chygryn, O., Pimonenko, T., Kwilinski, A.: Stakeholders' engagement in the company's management as a driver of green competitiveness within sustainable development. Sustainability **15**(9), 7249 (2023). https://doi.org/10.3390/su15097249
28. Mukhametzhan, S., Junusbekova, G., Daueshov, M.: An econometric model for assessing the asymmetry of urban development as a factor of regional economic growth: the case of Kazakhstan. Ind. Eng. Manag. Syst. **19**(2), 460–475 (2020). https://doi.org/10.7232/iems.2020.19.2.460
29. Muratbekova, A.: Digitalization in central Asia: progress and potential. In: Akıllı, E., Güneş, B., Güner, O. (eds.) Digital Diplomacy in the OSCE Region. Contributions to International Relations, pp. 33–44. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-50966-7_4
30. Nikolaesku, I., Budnyk, O., Bondar, V., Tepla, O., Berezovska, L.: Pedagogical management in inclusive process of the educational institution. Revista Amazonia Investiga **10**(39), 76–85 (2021). https://doi.org/10.34069/AI/2021.39.03.7
31. Nurmukhametov, N.: State and problems of the machine-building complex of the republic of Kazakhstan. Sci. Horiz. **25**(7), 115–122 (2022). https://doi.org/10.48077/SCIHOR.25(7).2022.115-122

32. Nurmukhanova, G., Alibekova, G., Tamenova, S., Niyetalina, G.: Strategic management of universities for regional competitiveness. J. Asian Financ. **8**(1), 551–562 (2021). https://doi.org/10.13106/jafeb.2021.vol8.no1.551. Economics and Business
33. Oswal, N., Al-Kilani, M.H., Faisal, R., Fteiha, M.: A systematic review of inclusive education strategies for students of determination in higher education institutions: current challenges and future directions. Educ. Sci. **15**(5), 518 (2025). https://doi.org/10.3390/educsci15050518
34. Page, M.J., et al.: The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ **372**, n71 (2021). https://doi.org/10.1136/bmj.n71
35. Ranga, M., Etzkowitz, H.: Triple helix systems: an analytical framework for innovation policy and practice in the knowledge society. Ind. High. Educ. **27**(4), 237–262 (2013). https://doi.org/10.5367/ihe.2013.0165
36. Rodríguez-Pose, A., Bartalucci, F., Rau, G.: One Kazakhstan, multiple nations: on a growing regional divide amidst economic dynamism. Asia-Pac. J. Region. Sci. **9**(3), 629–651 (2025). https://doi.org/10.1007/s41685-025-00385-7
37. Sargsyan, L.: The problem of inclusive education in higher schools of Armenia. Educ. 21st Century **6**(2), 60–73 (2024). https://doi.org/10.46991/educ-21st-century.v6.i2.60
38. Skvarciany, V., Jurevičienė, D.: Comprehensive assessment of enterprise digital competitiveness. Prague Econ. Papers **33**(2), 220–243 (2024). https://doi.org/10.18267/j.pep.857
39. Smith, M., Sarabi, Y.: UK trading patterns within and between regions in the automotive sector—a network analysis. World Econ. **44**(2), 510–529 (2021). https://doi.org/10.1111/twec.13006
40. Stratan, A., Duca, D., Duca, D.: Exploring drivers of national competitiveness: Insights from productivity, innovation and infrastructure. In: Xu, J., Binti Ismail, N.A., Dabo-Niang, S., Ali Hassan, M.H., Hajiyev, A. (eds.) ICMSEM 2024. LNDECT, vol. 215, pp. 1344–1355. Springer, Singapore (2024). https://doi.org/10.1007/978-981-97-5098-6_92
41. Supriyanto, S., Wiyono, B.B., Muhimmah, H.A.: Strategies for building an inclusive culture in higher education. Glob. J. Sociol.: Curr. Issues **15**(1), 12–21 (2025). https://doi.org/10.18844/gjs.v15i1.9711
42. Taikulakova, G., Myrzakhmet, M., Myrzakhmet, B., Bazarkhanova, G.: Analysis of the entrepreneurial activity of regional universities in the process of creation industrial clusters. Econ. Strateg. Pract. **19**(2), 100–122 (2024). https://doi.org/10.51176/1997-9967-2024-2-100-122
43. Tashbaev, B., Usmonov, B., Omanov, S.: Analysis of public debt policy and its effectiveness in Uzbekistan. In: Development of International Entrepreneurship Based on Corporate Accounting and Reporting According to IFRS, pp. 191–195. Emerald Publishing Limited (2024). https://doi.org/10.1108/S1877-63612024000033B024
44. Vakaliuk, T., Novitska, I., Verbovsky, I., Rozhnova, T., Kontsedailo, V.: Management of educational institutions with inclusive education based on innovative technologies. In: SHS Web of Conferences, vol. 104, p. 03008 (2021). https://doi.org/10.1051/shsconf/202110403008
45. Vanchukhina, L.I., Leybert, T.B., Ergasheva, S.T., Khalikova, E.A., Khanafieva, I.R.: Integration of the higher education systems of Russia and the Republic of Uzbekistan in training for the digital economy. In: Popkova, E.G., Sergi, B.S. (eds.) Digital Education in Russia and Central Asia, pp. 3–13. Springer, Singapore (2022). https://doi.org/10.1007/978-981-16-9069-3_1

46. Wensing, A., Wensing, E., Virgo, M.: Towards a core curriculum for civic engagement on appropriate technology: characterizing, optimizing and mobilizing youth community service learning. Afr. J. Sci. Technol. Innov. Dev. **10**(7), 867–877 (2018). https://doi.org/10.1080/20421338.2018.1439279
47. Yessenbek, C., Bekbolsynova, A., Numanova, F.: Analysing the efficiency of scientific potential in driving national competitiveness: evidence from emerging economies and the case of Kazakhstan. J. Appl. Econ. Sci. **20**(4), 975–988 (2025). https://doi.org/10.57017/jaes.v20.4(90).19
48. Zhou, W., Li, H.: A study on the multidimensional driving mechanism of cross-regional scientific collaboration network in China. Technol. Anal. Strateg. Manag. **36**(11), 3667–3681 (2024). https://doi.org/10.1080/09537325.2023.2220824

Region-Focused Architectural Design in the Context of the Circular Economy

Oxana Gurskaya(✉), Yelena Feoktistova, Tatiana Inozemtseva, and Madina Yussubaliyeva

D.Serikbaev East Kazakhstan Technical University, Ust-Kamenogorsk, 070004, D. Serikbayev STR, 19 Oskemen, Kazakhstan
{ogurskaya,myussubaliyeva}@edu.ektu.kz, {EFeoktistova, TInozemtseva}@ektu.kz

Abstract. The transition of the construction industry towards the principles of the circular economy is regarded as one of the key areas for ensuring the sustainable development of the urban environment. In industrial regions, the issue of reusing existing industrial buildings is of particular relevance, as these can be viewed as a significant resource for secondary building systems. One promising approach is the concept of 'urban mining', which involves the use of material resources accumulated within the urban environment. However, in most studies, urban mining is viewed primarily as a tool for analyzing material flows, whilst its potential as a strategy for architectural design remains insufficiently explored.

The aim of this study is to develop a conceptual approach to regionally-oriented architectural design based on the integration of circular economy principles and the utilization of the potential of the industrial building stock. The research methodology includes an analysis of the scientific literature, a structural-typological analysis of industrial building systems as well as methodological modelling at the intersection of regional planning and architectural design. The research has identified architectural principles and design tools that enable the integration of circular economy strategies into the transformation processes of industrial territories. A conceptual model is proposed for utilizing the industrial building stock as a spatially organized resource for secondary structural systems, and has been tested using a case study of a facility reorganization. The findings expand architectural approaches to implementing circular economy principles and can be applied in the development of strategies for the regeneration of industrial regions.

Keywords: Circular economy · Urban mining · Regional architectural design

1 Introduction

As urbanization accelerates and natural resources are consumed at a growing rate, the construction sector is becoming a critical component of urban environmental sustainability. It is evident that the construction sector has a high resource intensity - consuming up to 40% of natural resources and energy [6], as well as the lack of a unified conceptual model for implementing the circular economy. This reinforces the need to transition

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 419–436, 2027.
https://doi.org/10.1007/978-3-032-30518-3_27

from linear models to circular economic principles [7]. Among the promising areas is urban mining. This views the existing urban environment as a source of secondary resources accumulated in buildings, infrastructure and engineering systems. According to industrial ecology, circular economies are analogous to natural ecosystems with closed material and energy flows. In a study developed by [8], many key principles of the circular economy - systems thinking, resource efficiency, and circularity - are derived from industrial ecology.

The importance of sustainable and adaptive architectural solutions has been highlighted in research on regional architecture in Kazakhstan. G.S. Abdrasilova and L.M. Aukhadiyeva [9] note that the contemporary development of the country's architecture is closely related to the implementation of sustainable design principles. It includes the adaptive reuse of buildings and the development of a national aesthetic identity. These approaches form the basis for integrating urban mining into the architectural design of industrial regions. Regional studies highlight the importance of taking into account the economic, social, natural, climatic, and historical and cultural conditions of a region's development [10, 11], since these studies constitute the basis of a practice-oriented methodological framework for implementing circular economy principles and urban mining in the architectural design and management of material resources in industrial regions.

The authors believe that the integration of circular economy principles into the architectural design of industrial areas is possible through the development of regionally adapted design strategies, which utilize the existing industrial building stock as an urban mining resource and use design tools for the regeneration of industrial areas as an urban mining resource. This hypothesis defined the objective of the study, which is to develop a methodological strategy for the renovation of industrial facilities that takes regional characteristics into account. According to the stated research objective, the following research questions are formulated:

In what ways can the principles of the circular economy be incorporated into regional architectural design at industrial facilities?

What potential does the existing industrial building stock hold as a source of secondary resources?

What architectural and design tools are available to facilitate the integration of circular economy principles into industrial site regeneration programs?

The answers to these research questions formed the basis for the development of a methodological strategy aimed at improving resource efficiency, reducing environmental impact, and creating a sustainable architectural environment in Kazakhstani industrial regions.

2 Literature Review

According to Guerra et al. [12], a systematic review of over 300 publications, circular construction is still at its infancy and focuses primarily on waste management and material reuse. Finamore and Oltean-Dumbrava [6] highlight the lack of a single definition of the circular economy in relation to construction and stress the need for a holistic approach that encompasses the entire building value chain. There is also a need for a

comprehensive approach that considers all aspects of building life cycle management, including design, operation, renovation, and demolition [13, 14].

The integration of circular economy principles into the construction sector is primarily focused on the management of construction and demolition waste. The bibliographic analysis by Hasibuan G. et al. [15] examines the main research areas related to the use of circular economy principles in the management of construction and demolition waste, as well as sustainable engineering practices in buildings.

A number of academic publications have addressed the methodology of circular architectural and construction design in the past 10–15 years. In the present day, frameworks for strategies integrating circularity principles into a design object's life cycle are being developed most actively, including the formulation of revised requirements for architects' professional competencies [16–18], as well as assessment methodologies and indicators of building circularity [19] and industrial areas [20]. An approach to territorial and regional planning of industrial areas is defined by the circular economy. The method replaces linear location factors with closed cycles that are intended to produce "economic symbiosis." A number of modern industrial areas are already designed to function as ecosystems based on resource distribution and regional specificities [20, 21]. As part of this effort, circular revitalization practices are being implemented in old or degraded industrial areas [22]. Despite an increasing number of publications, the spatial aspects of circular economy implementation at the regional level remain largely unexplored. Several researchers report that studies have been limited in scope, particularly because the spatial organization of industrial territories is fragmented and there are no standardized circular economy indicators at the regional level [10, 11].

The literature on spatial aspects and regional planning of industrial areas in Kazakhstan has not yet addressed these issues. In the research, the macro-region is represented by issues of the circular economy at the sectoral, management, or entrepreneurial levels [23, 24], without addressing territorial and industrial planning at the regional level. The Republic of Kazakhstan has formed a comprehensive database of strategic documents [1–4], focused on the consistent widespread implementation of the principles of sustainable development and green economy. Key government initiatives and regulations, from the long-term strategy "Kazakhstan – 2050" [1] to the Environmental Code of the Republic of Kazakhstan [3] and international climate obligations [4], create institutional conditions for the transition to circular architectural and construction practices. At the same time, the analysis of the current design regulatory documentation of industrial facilities [32–35] shows that the concept of urban mining has not yet been properly reflected and applied in practice.

The authors of this article have attempted to contribute to the growth of research on regional planning of industrial territories in the context of the circular economy by filling a part of the research gap. Ust-Kamenogorsk, a city with a history of metallurgy and environmental problems, has been used as an example. This study aims to develop a circular economy-based planning strategy. With this approach, it is possible to enhance the sustainability of industrial zones, optimize resource use, and revitalize territories while taking into consideration local specificities.

3 Methods and Materials

A hybrid methodology was selected due to the interdisciplinary nature of the research topic, which is situated at the intersection of concepts relating to the circular economy, urban mining, architectural and urban design, and spatial planning. This article utilized a systematic qualitative analysis of scientific literature and regional strategic documents in conjunction with a deductive-inductive approach in identifying the criteria, characteristics, and parameters associated with circular design in addition to a structural-typological analysis of industrial facilities and an analysis of these results based on regional considerations (Fig. 1).

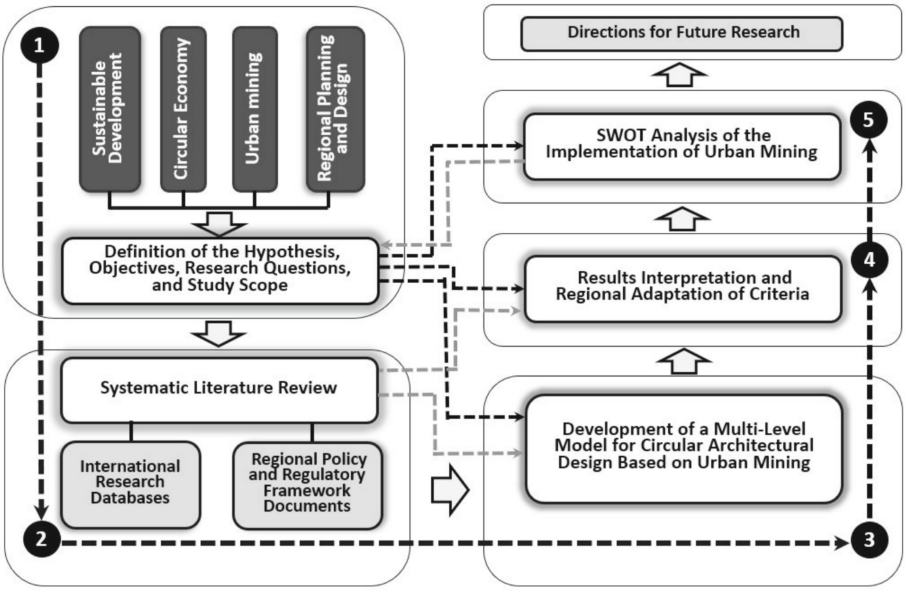

Fig. 1. Research methodology.

The study comprises five main stages:

In the first stage, the development of a conceptual framework will include the clarification of terminology, the formulation of a hypothesis, and the formulation of research questions;

The second step is to conduct a systematic review of sources based on keyword searches and analysis of scientific literature primarily from the Scopus database. A search for the following keywords and combinations thereof was conducted: “circular economy”, “urban mining”, “industrial sites”, “industrial buildings”, “regional planning”, “regional design”, “sustainable local development”, “architectural design”, “local architectural design”. According to the most comprehensive search, only two articles were found, requiring the search parameters to be widened. Duplicate articles and those with relevant titles were excluded from the various combinations. After analyzing the abstracts, a selection was made. A total of 30 articles were selected for publication.

Almost all restricted access publications were excluded from the search; these were analyzed at the abstract stage. The keywords and context of strategic regional documents were also analyzed;

At the third stage, a multilevel model of circular architectural design is developed based on urban mining within territorial boundaries and the typological boundaries of "city" and "building". In this study, approaches, indicators, metrics, and their application in regional design were identified through a deductive-inductive analysis of existing research.

In the fourth stage, the model's structure was substantiated by interpreting the results in relation to the regional context. The strategy's key levels and the mechanism for implementing it should be defined, the main indicators of regeneration should be described, conduct a pilot study and conclusions drawn re-garding the feasibility of implementing these indicators in Kazakhstan's regional context.

Stage five consists of a SWOT analysis, which enables the integration of urban mining into industrial area planning, as well as providing a structured framework to plan strategies for the sustainable development of the city.

The empirical subject of the study is the city of Ust-Kamenogorsk - one of the largest industrial centers in Eastern Kazakhstan, which has a historically established industrial structure. Many of the studies examined cities that are developing according to a circular planning strategy [25, 26].

While there are numerous environmental initiatives and sustainable development projects in Ust-Kamenogorsk, it has not yet achieved full status as a circular industrial city model. Despite the existing resources, industrial buildings, and well-developed engineering infrastructure, as well as the presence of active stakeholders as scientists and technicians at the technical university, the implementation of an international educational project on the circular economy provides the opportunity for the implementation of sustainable strategies, such as the use of urban mining principles, the reuse of building materials, and the regeneration of waste sites.

Almost all sectors of industry are represented in the city, including mining, metallurgy, mechanical engineering, manufacturing, light industry, food industry, energy production, and large-scale infrastructure. In the region, industrial production accounts for 77.2% of the total production. In 2024, the industrial output is expected to increase from 2.1 trillion tenge to 3.9 trillion tenge by 2030, as a result of reconstruction and modernization, along with the opening of new workshops and factories [5].

Ust-Kamenogorsk has a substantial industrial base, which was established primarily between 1940 and 1980. In the 1990s, the economic crisis led to the closure of production facilities, resulting in the emergence of abandoned production facilities. In the subsequent years, some sites have been partially restored or repurposed for other purposes, while others remain vacant. Presently, industrial areas constitute 13.2% of the total area of the city, of which approximately 2.3% is unused (Fig. 2). As the mining and metallurgical sector - the city's main economic driver - declines, significant modernisation and regeneration of industrial areas may also be necessary.

The legacy of these buildings has resulted in the development of an extensive stock of industrial buildings, ancillary structures, standardized warehouses, and mass-produced

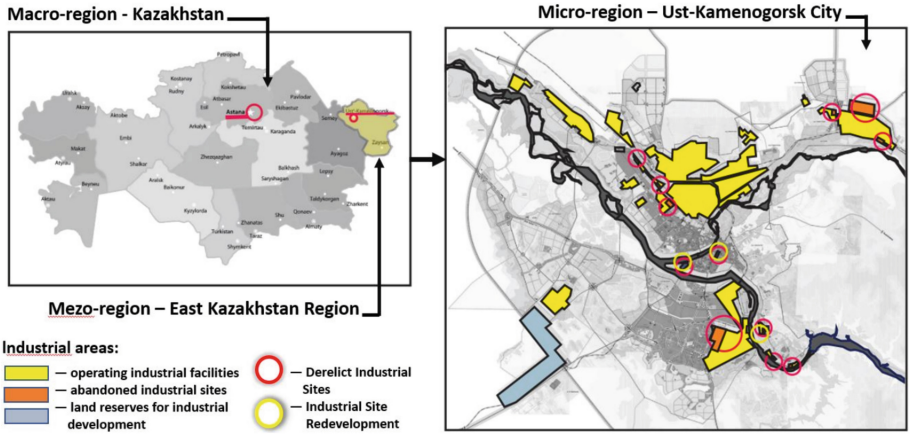

Fig. 2. Regional context of the study.

prefabricated structures. An analysis of the local building stock revealed that the structures of the region's industrial buildings consist of structural systems consisting of reinforced concrete frames, steel frames, panels, floor slabs, and flat systems made of bricks and blocks, in addition to long-span structures which form the spatial framework of industrial architecture.

4 Urban Mining as a Tool for Circular Regeneration of Industrial Sites

Recent years have seen an active development in the assessment of urban mining potential within the context of circular economies and sustainable construction. This assessment is based on analysis of the material flow, the life cycle of buildings, and the spatial characteristics of urban environments. According to a number of current studies, cities continue to be viewed as anthropogenic deposits of secondary resources, primarily building materials. Accordingly, propose an assessment of the potential of urban mining by analyzing the material stock of existing buildings, the proportion of recoverable components, the level of reuse, and the environmental effects of preserving structures in order to assess the potential for urban mining [27]. In their paper, the authors emphasize the importance of quantifying circularity indicators, such as the proportion of dismantled and reused items, the content of secondary materials, as well as the reduction of carbon footprints. As Antonella Luciano and co-authors [28] demonstrate, the preservation of existing structures can significantly reduce the consumption of primary resources and greenhouse gas emissions. Contemporary research is also taking into account the spatial aspect of urban mining by analyzing the density of material stocks in the built environment, the spatial distribution of resources, and the possibility of reusing building components in circular renovations. It is important to note that urban mining is implemented at several spatial levels, from the city as a whole to the individual building, as part of the circular transformation of industrial areas. Depending on the scale, the nature

of resources managed, the decisions made, and the design tools used will differ. Figure 3 illustrates the application of urban mining at the city and building levels.

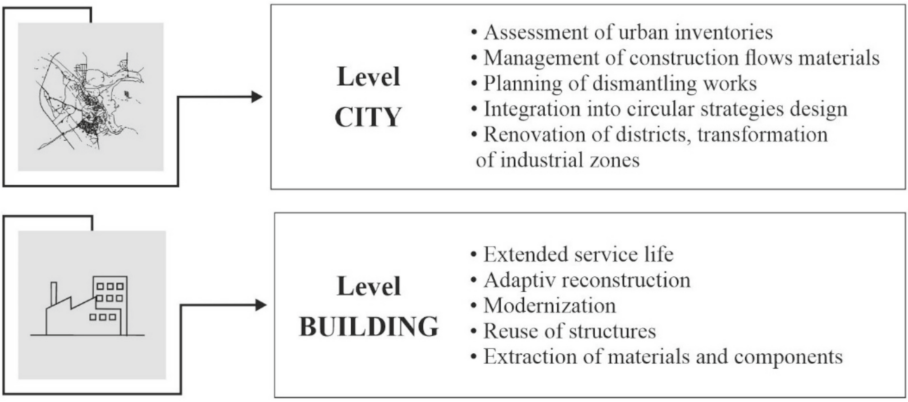

Fig. 3. Implementation of urban mining principles at the territorial and building levels of the circular transformation of the industrial environment.

In this model, the territorial material resources of an industrial city are transformed into specific architectural solutions by means of reconstruction, modernization, and reusing existing structures in order to reflect the multi-scale nature of urban mining. The findings confirm the possibility of transitioning from the assessment of urban material reserves to the design of circular buildings. The diagram serves as a basis for further examination of the levels of resources in the industrial environment.

Territories associated with industrial activity may be viewed as an interconnected system of material and spatial resources. Model is proposed, which allows the industrial environment to be viewed as a combination of territorial, structural, material, and infrastructure resources that can be integrated into circular architectural design strategies (Fig. 4).

A transition can be achieved from a regional level of materials awareness to a project level of architectural decisions through the proposed model structure. This model allows architects to interpret industrial sites as structured material banks despite the absence of quantitative GIS assessments. The system also allows architects to integrate the logic of urban resource extraction into early design stages. Research conducted by Panizza R. O. and Nik-Bakht M [27, 32]. indicates that the key resource for circular transformation is not only recoverable materials, but also the existing building stock as a whole. This requires a shift in focus from assessing individual waste streams to analyzing buildings as carriers of spatial, structural, and material resources.

In this context, the circular approach involves three stages: auditing the industrial building stock, assessing their life cycle, and modeling modernization scenarios.

Stage 1 - Identification of clusters of resources. Secondary materials may be found in industrial zones and infrastructure corridors. Structures such as reinforced concrete frames, steel trusses, prefabricated panels, and modular structural grids can be reused rather than being discarded.

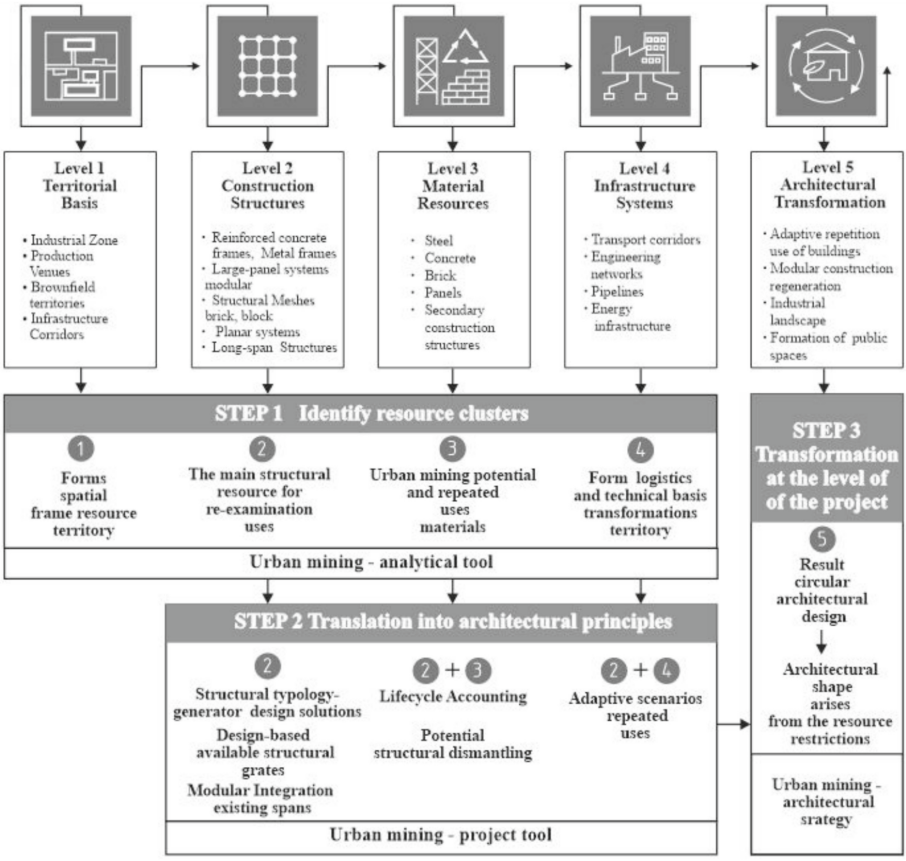

Fig. 4. A multi-level model of circular architectural design based on urban mining.

Stage 2 - Translation of the design into architectural principles.

The following principles can be applied based on the identified structural typologies:

1. Considering the life cycle of the building. The design of a structure should consider both its operation as well as its dismantling and planning for possible stages of repair, replacement, and demolition. 2. Using the available structural grids for design

3. Integrating existing spans in a modular manner. 4. Potential for structural dismantling. Identifying elements that can be reused. 5. Scenarios for adaptive reuse.

A clear division is provided in the design between structural elements that can be retained and reused and those that must be reprocessed or disposed of.

Stage 3 - Transformation at the project level. Using an industrial frame with a 6 by 6 m structural grid, for example, can serve as the spatial basis for a modular public building, where the existing load-bearing system is retained, supplemented and repurposed. The architectural form in this scenario is not the result of abstract composition, but rather of physical constraints.

A five-level hierarchy describes the spatial and material resources within an industrial site, while a three-stage circular approach outlines the sequence of how these resources are transformed into architectural solutions. It has evolved from a tool for assessing material stocks into a design generation process that integrates territorial, structural, material, and infrastructure systems into a unified framework for developing architectural solutions.

Researchers have examined a variety of approaches to reducing the resource burden associated with the construction sector, including sustainable architecture, circular design, and urban mining. Although both concepts address sustainable development, their research and practical orientations differ. Table 1 presents a comparative analysis of these approaches' key characteristics.

Table 1. Comparative analysis of circular design, urban mining and sustainable architecture.

Criterion	Circular design	Urban mining	Sustainable architecture
Main objective	Establishing closed-loop cycles for the use of materials and resources	Extracting secondary resources from the existing urban environment	Reducing the environmental impact of architecture
Scope of application	Architectural project, building systems	Urban area, material flows	Buildings, neighbourhoods, urban environments
Main subject of analysis	Life cycle of building materials and building elements	Material stocks in existing buildings and infrastructure	Energy efficiency and environmental performance of buildings
Application stage	Architectural design stage	Demolition and material recycling stage	Building design and operation stage
Architectural aspect	Integration of circular principles into architectural form and construction	Architectural aspect is considered in a limited way	Architecture is regarded as an environmentally responsible system
Type of resources used	Primary and secondary building materials	Recycled materials from existing buildings	Eco-friendly and energy-efficient materials
Main constraints	Need changes to design methods	The complexity of tracking material flows	Focus primarily on building energy
Potential for integration	Urban mining and sustainable design	Can serve as a source of secondary materials	Forms the environmental basis of architectural design

Based on the results of the analysis conducted, urban mining is primarily concerned with extracting secondary resources from existing buildings. The concept of circular design, on the other hand, involves integrating closed material cycles directly into the design of buildings. In addition, sustainable architecture is concerned with reducing the impact of buildings on the environment by implementing energy-efficient measures.

Based on the findings of this study, it is proposed that region-oriented architectural design be viewed as an integrative approach that combines spatial analysis of building resources with the principles of circular design.

This study proposes a system of interrelationships between urban mining potential, design principles, and design tools. It is used to translate spatial-resource analysis results into architectural design solutions. In Table 2, each principle is accompanied by assessment indicators that provide a quantitative assessment of how well circular strategies are implemented during the architectural design process.

Table 2. Integration of urban mining principles into region-oriented architectural design

Potential of urban mining	Architectural principle	Design tool	Evaluation indicator
Presence of existing load-bearing structures	Adaptive reuse of buildings	The design solution adapts to the standard dimensions of elements and structural characteristics Retention of the framework. Repurposing of industrial buildings, functional transformation	IO 1 Proportion of building volume retained (%)
Reuse of structural elements	Design based on available structural resources	Use of existing frames, spans, and column grids. Changes the logic of form-making: from 'form-material' to 'resource-architectural system'	IO 2 Percentage of reused structures (%)
Availability of modular industrial structures	Modularity of architectural solutions	Modular design, transformable spaces, flexible spatial configuration, standard structural steps	IO 3 Structural modularity coefficient
Possibility of dismantling structures	Dismantlability and reversibility of connections	Dismantling, reassembly. Bolted connections, dry assembly systems.	IO 4 Proportion of structures suitable for dismantling (%)
Availability of secondary building materials	Circularity of the material cycle	Use of recycled materials (bricks, metal, panels)	IO 5 Proportion of recycled materials (%)
Geographical concentration of industrial facilities	Locality of resource flows	Use of materials from regional sources Spatial connectivity of resource flows	IO 6 Average distance of material transport
Availability of large industrial sites	Renovation of the industrial landscape	Adaptable planning structures. Comprehensive redevelopment of the area.	IO 7 Functional flexibility coefficient
Potential for future dismantling	Design for demountability.	Designing with future dismantling in mind, reversible connections.	IO 8 Proportion of materials suitable for reuse (%)

A typological analysis of 16 unused industrial buildings in Ust-Kamenogorsk identified seven types of buildings and structural configurations, making it possible to assess

their potential for integration into circular architectural design strategies. The correspondence between typologies and possible architectural applications is presented in Table 3.

Table 3. – The relationship between the typologies of the industrial building stock in Ust-Kamenogorsk and circular architectural strategies.

Typology of industrial building stock		Circular design strategy	Application in architecture
RC frame industrial buildings	Long-span structures (12 m, 18 m, 24 m)	Structural reuse/ Spatial transformation with retained frame	Public buildings and multifunctional public spaces
	Buildings with regular structural grid (3x3 m; 6x6 m)		
Steel-framed industrial buildings	Long-span structures	Demolition and reuse of structural elements/ Spatial transformation with preservation of the structural frame	Modular pavilions/ temporary public structures
	Buildings with regular structural grid		
Large-panel reinforced concrete industrial buildings	Cellular structure	Reuse of construction elements	Low-rise public buildings, administrative facilities and office buildings, residential buildings
Framed-wall industrial and warehouse buildings	Structures with a regular load-bearing grid	Adaptive spatial transformation of buildings with frame preservation	Low-rise public and residential buildings, administrative facilities
Brick industrial and warehouse buildings	Structural system with longitudinal load-bearing walls	Deconstruction and reuse of structural elements/ Adaptive reuse	Office spaces, educational campuses, administrative facilities, residential buildings

Interactions defined above include: identify structural modules and elements of industrial buildings suitable for reuse; transform them into architectural design solutions; demonstrate the possibility of creating innovative spaces based on industrial structural systems.

To illustrate the application of the proposed model, a site in Ust-Kamenogorsk, situated in the coastal zone of the Irtysh River, is examined. The industrial building built in 1956 and situated in the central part of the city, has been successively adapted for various functions (furniture production, warehouse, sales area, fitness center), reflecting its typological flexibility. The building's spatial and layout parameters, along with its structural scheme—a partial frame with two rows of brick columns supported by load-bearing exterior walls—allow the structure to be interpreted as a carrier of sustainable

spatial and structural resources within the framework of urban mining. Visually, at the level of the overall model, the stages of project development are shown in Fig. 5.

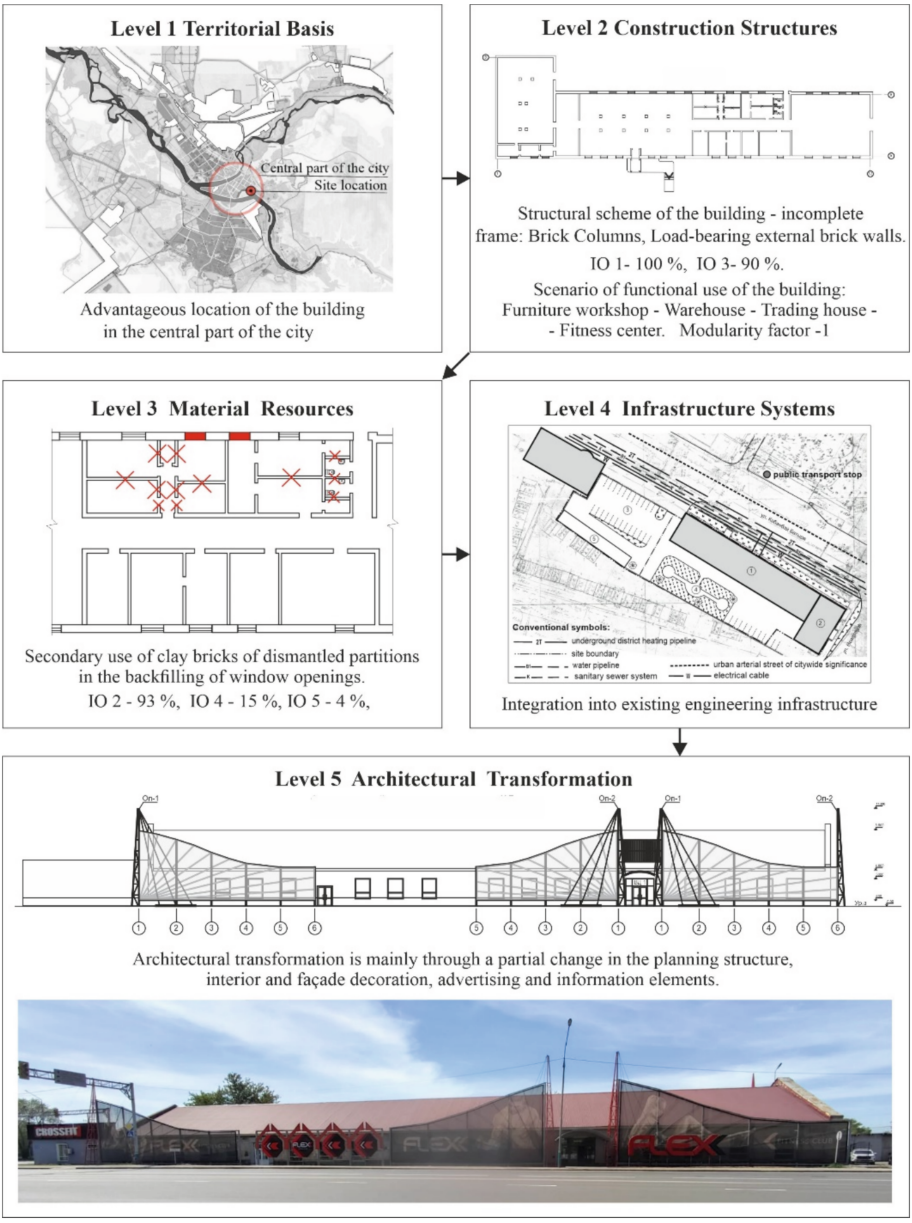

Fig. 5. Testing the model using an industrial building in Ust-Kamenogorsk as an example

Current seismic zoning requirements for the city (7 on the Richter scale) have necessitated the reinforcement of load-bearing elements with steel frames and beams. The change in the building's function and its reintegration into the urban fabric contribute to the enhancement of the waterfront area and ensure the transformation of the former industrial site into a more vibrant and community-oriented element of the urban environment. By extending the traditional understanding of urban mining and contributing to the implementation of circular economy principles, the study confirms that the industrial building stock can be regarded as a spatial-structural resource for architectural design.

In the construction industry, the challenges of the circular economy are not only technological, but also political, economic, organizational, managerial, technical, and environmental. It should be noted that the transition to circular architecture requires the transformation of the entire sector, rather than the transformation of individual technologies [18, 29]. A key characteristic of circular architecture is adaptive reuse, which significantly reduces the impact on the environment, creates new cultural values, and preserves buildings' industrial identity [30]. The post-industrial heritage can be used as a resource for sustainable development.

Taking into account the specific characteristics of post-Soviet cities, regional strengths, obstacles, and limitations, as well as opportunities for applying urban mining in Ust-Kamenogorsk, were identified (Table 4).

The advanced state of disrepair of many buildings makes mechanical demolition the predominant practice, which significantly limits the potential for urban mining. However, the city has a sufficient stock of industrial buildings constructed in more recent years that can be adapted for new uses while retaining their original volume.

5 Results

As a result of the study, a conceptual framework for region-oriented architectural design of industrial areas, based on urban mining, has been developed.

The data obtained show that the existing industrial building stock constitutes a multi-level system of spatial and material resources, comprising territorial, structural, infrastructural and material components. These elements can be interpreted as secondary structural systems capable of enabling the creation of new architectural functions without completely dismantling the original structure.

In order to facilitate the transition from material stock assessment to design solutions, a multi-scalar model of circular architectural renovation has been developed. The model operates across several spatial levels – from the urban area to the individual buildings – and comprises three sequential stages: identifying resource clusters within the industrial landscape; transforming identified resources into architectural design principles; and implementing solutions at the level of a specific building.

The greatest potential for circular renovation lies in the preservation of existing structural systems. Typical industrial architecture features include reinforced concrete and steel frames, long spans, and regular column grids. Through the preservation of these elements, it is possible to significantly reduce the consumption of primary resources, minimize the volume of construction waste, as well as guarantee flexibility in the functional renovation of buildings.

Table 4. SWOT analysis of the implementation of urban mining in Ust-Kamenogorsk.

(S) Strengths	**(W) Weaknesses (limitations)**
High concentration of industrial facilities suitable for adaptation and modernization.	Low selectivity of dismantling. Predominantly old, standardized design solutions that do not take future dismantling into account.
Large stock of obsolete Soviet-era buildings (source of secondary materials)	Lack of infrastructure for processing construction waste.
Availability of a metallurgical base and recycling facilities	Lack of a regulatory framework and standards for urban mining
High density of utility networks (potential for resource extraction)	Lack of markets for secondary materials
Availability of scientific and educational potential (school of architecture)	Lack of expertise among designers and contractors
Opportunity for integration into environmental regeneration programmer	Lack of digitalization of the building stock (no BIM passports)
(O) Opportunities	**(T) Threats**
Transition to a circular economy at national level	Economic instability and dependence on the raw materials sector
Reduced pressure on landfill sites	Low demand for recycled materials
Creation of new jobs in the 'green' economy	Risk of pollution if dismantling is carried out incorrectly
Reducing the carbon footprint of construction	Complexity of certifying recycled materials
Improving the sustainability of the urban environment	Negative public perception of 'recycled materials'

An integrated system is proposed for integrating urban mining potential into architectural design. The system establishes a relationship between the characteristics of the building stock, architectural principles, design tools, and quantitative assessments. Circular strategies may be incorporated as early as the design stage instead of only at the demolition or recycling stage.

A spatial analysis of Ust-Kamenogorsk identified industrial sites with potential for adaptive reuse. Using the example of the renovation of one such facility, this paper demonstrates a strategy for implementing a multiscale model in real-world conditions. As a result of the pilot study, the typological and structural characteristics of the facility were identified, the feasibility of preserving existing spatial and structural systems while changing the building's functional purpose was confirmed, and the impact of the facility's transformation on the urban context was determined.

The SWOT analysis revealed significant resource and logistical advantages, while also identifying critical constraints, including the lack of a regulatory framework, infrastructure for processing construction waste, digital building inventories, and a developed market for recycled materials.

6 Discussion

The findings indicate the need to rethink the role of urban mining in architectural design as a methodological foundation for developing architectural solutions.

The proposed approach facilitates a transition from the linear "demolition–new construction" model to a model of adaptive regeneration of the industrial environment. In this model, existing sites and buildings serve as vessels for spatial, structural and cultural potential. This is particularly relevant to post-Soviet industrial cities that possess a significant stock of standardized structures with high load-bearing capacities.

Compared with international practices, western and European countries support circular construction through regulatory instruments, digital material registers, and secondary resource markets. Circular strategies are at an early stage of development in the region under study, which significantly limits their implementation. As well as this, post-Soviet industrial cities possess a unique combination of factors: a high concentration of material stocks, a compact urban structure, and close proximity to industrial and storage areas. The presence of these characteristics creates favorable conditions for the flow of circular resources within the local area, provided appropriate planning and management tools are in place.

It is also confirmed by the research findings that adaptive reuse is the most environmentally effective strategy for circular architecture. A significant reduction in embodied carbon, a reduction in waste volume, and preservation of the industrial identity of the area can all be achieved by retaining existing structures.

The study is limited by the lack of detailed quantitative data on the material composition of the building stock. The primary focus has been placed on typological analysis without conducting full-scale life-cycle assessments. A promising area of future research is the development of digital building registers, GIS mapping of material stocks, and quantitative models for assessing circular efficiencies.

7 Conclusions

The study confirms that industrial sites can be reimagined as spatially organized material resources for circular development. The developed multi-level model facilitates the transition from analyzing a site's resource potential to formulating specific design solutions, linking urban planning and architectural levels. There is significant potential for functional renovation of industrial buildings from the Soviet era that does not require a significant investment of resources.

Testing of the model confirmed its suitability for analyzing the transformation of industrial sites, preserving their spatial and structural systems, reusing materials, and assessing the impact on the urban context.

In industrial regions, circular strategies are constrained by institutional, economic, and technological barriers. In order to achieve this, it is necessary to develop a specialised regulatory framework as well as tools for managing digital resources.

Thus, the integration of urban mining principles into region-oriented architectural design opens up prospects for the sustainable renewal of industrial cities.

It is imperative that further research be conducted in the following areas to ensure the practical implementation of the strategy in the region:

- develop a classification of enterprises into groups, taking into account technological processes and environmental impact, to determine the extent of potential renovation of industrial sites or their demolition;
- to study the typology of industrial buildings based on structural and spatial characteristics, taking into account the differences in the life cycle of its main elements, which determine the possibilities and scenarios for regeneration;
- to identify standardized and unified structural systems and their components used in the region under study and to determine the future nature of their use (within the structure of the preserved facility, partial or complete replacement, or the reprocessing of replaceable components);
- to develop methodological recommendations on the application of circular economy principles in the architectural design of the industrial cluster in the East Kazakhstan Region.
- This study contributes to the development of the circular economy in the region, in accordance with the institutional documents of the Republic of Kazakhstan.

This study contributes to the development of the circular economy in the region, in accordance with the institutional documents of the Republic of Kazakhstan.

Disclosure of Interests The authors declare no competing interests.

Acknowledgments. The authors would like to thank M.V. Kozlov, chief engineer of the project to renovate the building into a fitness club, for providing materials and consulting support.

References

1. Kazakhstan-2050 Strategy: New Political Course of the Established State https://adilet.zan.kz/rus/docs/K1200002050
2. Concept of Kazakhstan's transition to a green economy: (2013–2050). https://policy.asiapacificenergy.org/sites/default/files/Concept%20on%20Transition%20towards%20Green%20Economy%20until%202050%20%28RU%29.pdf?utm_source=chatgpt.com
3. Environmental Code of the Republic of Kazakhstan (Code of the Republic of Kazakhstan dated 02.01.2021 № 400VI).: https://adilet.zan.kz/rus/docs/K2100000400?utm_source=chatgpt.com
4. Kazakhstan's National Contribution to Reduce Emissions (Paris Agreement): https://adilet.zan.kz/rus/docs/P2300000313
5. Development plan of the city of Ust-Kamenogorsk for 2026-2030: Annex to the decision of the Ust-Kamenogorsk city maslikhat № 42/4-VIII (2026). https://www.gov.kz/memleket/entities/vko-oskemen/documents/details/942079
6. Finamore, M., Oltean-Dumbrava, C.: Circular economy in construction — findings from a literature review. Heliyon **10**(15), e34647 (2024). https://doi.org/10.1016/j.heliyon.2024.e34647
7. Mrad, S., Ribeiro, F.: A review of Europe's circular economy in the building sector. Sustainability. **14**(7), 4177 (2022). https://doi.org/10.3390/su14074177
8. Saavedra, Y.M.B., Iritani, D.R., Pavan, A.L.R., Ometto, A.R.: Theoretical contribution of industrial ecology to circular economy. J. Clean. Prod. **170**, 1514–1522 (2018). https://doi.org/10.1016/j.jclepro.2017.09.260

9. Abdrassilova, G.S., Aukhadiyeva, L.M.: Priority directions for the development of regional architecture in Kazakhstan in the 21st century. Vestnik KazGASA. **3**(93) (2024). https://doi.org/10.51488/1680-080X/2024.3-01
10. Arsova, S., Genovese, A., Ketikidis, P.: Barriers and drivers for circular economy implementation in the construction industry. J. Clean. Prod. **354**, 131680 (2022). https://doi.org/10.1016/j.jclepro.2022.131680
11. Tapia, C., Bianchi, M., Pallaske, G., Bassi, A.M.: Towards a territorial approach for circular economy: the role of urban mining. Res. Conserv. Recycl. **164**, 105188 (2021). https://doi.org/10.1016/j.resconrec.2020.105188
12. Guerra, B.C., Leite, F., Faust, K.M., Ardito, L.: Urban mining: a systematic literature review. Res. Conserv. Recycl. **189**, 106727 (2023). https://doi.org/10.1016/j.resconrec.2022.106727
13. Timm, M., Wiedenhofer, D., Fishman, T., Krausmann, F.: Urban mining potentials of global building stocks. J. Ind. Ecol. **27**(2), 438–453 (2023). https://doi.org/10.1111/jiec.13339
14. Zhang, C., et al.: Urban mining of construction and demolition waste: a review of technologies and practices. Res. Conserv. Recycl. **188**, 106643 (2023). https://doi.org/10.1016/j.resconrec.2022.106643
15. Hasibuan, G.S., Purwanto, W.W., Hidayatno, A.: Circular economy in construction and demolition waste management: a bibliometric analysis. Sustainability. **17**(1), 512 (2025). https://doi.org/10.3390/su17010512
16. Dokter, G., Thuvander, L., Rahe, U.: Circular construction and adaptive reuse of buildings: investigating potentials and barriers. Sustainability **13**(8), 4261 (2021). https://doi.org/10.3390/su13084261
17. Kanters, J.: Circular renovation of buildings: decision support tools for the early design phase. Sustainability **12**(14), 5759 (2020). https://doi.org/10.3390/su12145759
18. Van Uden, M., Wamelink, H., Van Bueren, E., Heurkens, E.: Governance of circular area development: challenges and opportunities for urban regeneration. J. Clean. Prod. **449**, 141602 (2025). https://doi.org/10.1016/j.jclepro.2024.141602
19. Pierzchlewicz, D., Woźniak, A., Widera, B.: Circular economy in construction and demolition waste management in Europe: current state and future directions. Res. Conserv. Recycl. Adv. **21**, 200232 (2025). https://doi.org/10.1016/j.rcradv.2024.200232
20. Gamidullaeva, L., Shmeleva, N., Tolstykh, T., Shmatko, A.: Transformation of industrial territories as a factor of sustainable urban development. Sustainability **14**(9), 5258 (2022). https://doi.org/10.3390/su14095258
21. Cudečka-Puriņa, N., Atstāja, D., Koval, V., Purviņš, M., Nesenenko, P., Tkach, O.: Circular economy development in the construction sector: theoretical and practical aspects. Sustainability **14**(12), 7267 (2022). https://doi.org/10.3390/su14127267
22. Bressanelli, G., Visintin, F., Saccani, N.: Circular economy in the built environment: a systematic literature review. J. Clean. Prod. **312**, 127798 (2021). https://doi.org/10.1016/j.jclepro.2021.127798
23. Torgautov, B., Zhanabayev, A., Tleuken, A., Turkyilmaz, A., Mustafa, M., Karaca, F.: Circular economy: a pathway to sustainable development in the construction industry of Kazakhstan. Sustainability **13**(16), 9143 (2021). https://doi.org/10.3390/su13169143
24. Zhidebekkyzy, A., Temerbulatova, Z., Amangeldiyeva, B.A., Sakhariyeva, A.: Sustainable development of the construction sector in Kazakhstan: challenges and prospects. Sustainability **15**(7), 5934 (2023). https://doi.org/10.3390/su15075934
25. Friant, M.C., Reid, K., Boesler, P., Vermeulen, W.J.V., Salomone, R.: Sustainable circular economy discourses: a typology and critical analysis. J. Clean. Prod. **394**, 136333 (2023). https://doi.org/10.1016/j.jclepro.2023.136333
26. Fratini, C.F., Georg, S., Jørgensen, M.S.: Exploring circular economy imaginaries in European cities: a research agenda for the built environment. J. Ind. Ecol. **23**(4), 944–956 (2019). https://doi.org/10.1111/jiec.12899

27. Orenga Panizza, R., Nik-Bakht, M.: Urban mining and adaptive reuse of buildings: integrating circular economy principles into the built environment. Buildings **14**(2), 356 (2024). https://doi.org/10.3390/buildings14020356
28. Luciano, A., Cutaia, L., Barberio, G., Sbaffoni, S., Mancini, L.: Urban mining as a circular strategy for the built environment: opportunities and challenges. Res. Conserv. Recycl. **190**, 106806 (2023). https://doi.org/10.1016/j.resconrec.2022.106806
29. Charef, R., Morel, J.-C., Rakhshan, K.: Barriers to implementing the circular economy in the construction industry: a critical review. Sustainability **13**(23), 12938 (2021). https://doi.org/10.3390/su132312938
30. Jabłoński, W., Banachowska, E., Patyna, K.: Circular economy indicators for the construction sector: assessment and implementation challenges. Sustainability **17**(2), 845 (2025). https://doi.org/10.3390/su17020845
31. Panizza, R.O., Jalaei, F., Nik-Bakht, M.: Towards circular construction: material and component stock assessment in Montréal's residential buildings. Designs **9**(6), 129 (2025). https://doi.org/10.3390/designs9060129
32. Republic of Kazakhstan: Industrial buildings. (Code No. 3.02–127-2013) Astana (2013)
33. Republic of Kazakhstan: Industrial plant structures (Code No. 3.02–128-2012) Astana (2012)
34. Republic of Kazakhstan: Master Plans for Industrial Plants (Code No. 3.01–103-2012) Astana (2012)
35. Republic of Kazakhstan: Regulations on the reuse of used building materials, products and structures. (Code No. 1.04–108-201) Astana (2012)

Creative Destruction and Regional Resilience: The Case of the Volvo Investment in the Košice Region

Silvia Ručinská(✉), Lenka Jakubócová, and Ondrej Mitaľ

Department of Public Policy and Theory of Public Administration, Faculty of Public Administration, Pavol Jozef Šafárik University in Košice, 04011 Košice, Slovak Republic
{silvia.rucinska,ondrej.mital}@upjs.sk,
lenka.jakubocova@student.upjs.sk

Abstract. The paper focused on the concept of creative destruction and its significance for regional development in the context of the contemporary economy. Creative destruction represents a process of continuous transformation, in which old structures are replaced by new ones through innovation. Based on this, the paper analyses the current development impulse in the Košice region, specifically the construction of the Volvo car plant in the strategic park in Valaliky. This investment represents a significant opportunity for economic growth, employment, and innovation-driven regional development. At the same time, it also highlights potential risks associated with concentrating economic activity in a single dominant sector. Experience from Oulu in Finland shows that a sudden decline of a major industry can lead to serious regional problems. Therefore, the paper emphasises the importance of building ecosystem sustainability and resilience in the Košice region through three main pillars: knowledge creation, support for entrepreneurship, and the development of community spirit. Advancing these pillars can help the region effectively adapt to changes, minimise negative impacts, and ensure long-term sustainable development. The study's key benefit is its contribution to the ongoing debate, highlighting the importance of creative resilience as an integral part of the regional ecosystem.

Keywords: Creative Destruction · Regional Development · Sustainability · Regional Resilience · Creative Resilience · Regional Ecosystem · Košice Region

1 Introduction

Every region goes through certain stages of development. A region's development depends not only on its industries and investment drivers, but also on the quality of relationships among the stakeholders within that region. Therefore, it is necessary to focus not only on the region's economic orientation but also on cooperation within the regional ecosystem.

Creative destruction is one of the fundamental mechanisms of economic development, which also has a significant impact on regional development (Huang, Dong and

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 437–452, 2027.
https://doi.org/10.1007/978-3-032-30518-3_28

Li, 2025; Gilbert, 2012; Feld, Schnellenbach and Baskaran, 2012). Creative destruction is usually described as an endogenous growth model in which growth is driven by vertical innovations involving creative destruction (Aghion and Howitt, 2023; Aghion and Howitt, 1992). It is a process through which outdated production methods, technologies and entire industries are gradually replaced by innovations, thereby creating opportunities for new economic activities to evolve (Hospers, 2005). Creative destruction is a counterpart to creative creation, and both phenomena and their development can be predicted with a certain degree of probability (Irandoust, 2024).

Over the past decade, creativity has emerged, alongside innovation, as a key strategic characteristic of growth-oriented regional economies. Moreover, its importance will grow even more significantly in the coming decades, as creativity is becoming an indispensable prerequisite for finding solutions to complex and multidimensional global challenges, such as climate change and population ageing, making it a central theme not only for researchers but also for policymakers (Andrew and Spoehr, 2011; Florida, 2002). This trend is also reflected in the concept of the creative economy, which is based on the integration of knowledge, technology, and creativity with the aim of generating new markets and job opportunities, thereby marking a shift away from the traditional model of economic growth toward innovation-driven development (Sung, 2015).

A region's economic focus can and often does hit its development limits or face development risks, and it is therefore necessary to also consider regional resilience. Regional resilience is most frequently understood as the readiness of the region and its stakeholders to respond to and recover from potential shocks or unexpected events that disrupt the normal state of balance. In light of recurring industrial cycles and technological progress, building what is known as the creative resilience of regions is becoming increasingly important (Uctu, Tuluce and Aykac, 2024). Creative resilience of regions is based on the ability to be prepared and effectively absorb external shocks, adapt to new conditions, and at the same time actively create new development trajectories (Capello and Dellisanti, 2023). The current phase of development, also referred to as the Fourth Logistics Revolution, is fundamentally transforming the economic and social structure of regions and leading to the emergence of a so-called creative society, dominated by innovative urban regions with strong global connections and interdependencies (Andersson and Andersson, 2015).

In this context, the concept of the "creative class" also plays an important role, because according to this concept, regional growth is driven by a specific group of people who prefer an open, tolerant, and diverse environment, and successful regions combine talent, technology, and tolerance (Florida, 2002). Building awareness among such a group of people within the region can be seen as a responsibility of the private and public sectors (Pittz and Pittz, 2024; De Bernard et al., 2022). Furthermore, experience from Romania shows that the creative class can help absorb economic shocks and thus improve the region's economic resilience (Pavelea et al., 2023). When significant foreign investments are made, it is therefore essential not only to utilise their development potential but also to simultaneously build a diversified and resilient regional economy in the above-mentioned contexts.

The Košice Region currently serves as an example of a region facing such a transformative challenge due to the arrival of Volvo at the strategic park in Valaliky. This

investment has the potential to become a significant driver of economic growth, employment, and technological progress, particularly in the context of the development of the automotive industry and related sectors with suppliers. At the same time, however, experience from other post-industrial regions suggests that an overreliance on a single dominant sector can lead to increased vulnerability in the event of its decline (Andersson and Andersson, 2015). Therefore, this investment must also be viewed through the perspective of creative resilience, and we must support the development of knowledge, entrepreneurship, and innovation across multiple sectors to ensure the region's long-term sustainable and balanced development.

2 Theoretical Framework

2.1 Creative Destruction in Regional Development

According to Schumpeter, the economy is constantly changing and evolving, thus old firms and technologies disappear and are replaced by new ones. This process, which he called creative destruction, involves the continuous replacement of the old with the new. Innovation, a combination of new ideas, products, or production methods, plays a fundamental role in this transformation. In this regard, innovation is not random; rather, it builds on existing positive and mutual elements of the economy and is the result of the deliberate efforts of entrepreneurs and other stakeholders who transform and disrupt established structures. Schumpeter describes this mechanism as wave-like movements of innovative activity, which are essential for economic transformation and the introduction of new solutions (Hospers, 2005). Innovation is the main driving force behind economic transformation. Innovation manifests itself in five ways: (1) through new production processes, (2) new products, (3) new materials or resources, (4) expansion into new markets, and (5) the creation of new organisational forms (Schumpeter & Swedberg, 2021; Schumpeter, 1934).

Concepts such as creativity and innovation are now an integral part of regional economic development, as addressing complex challenges, such as climate change or an ageing population, requires creative approaches (Andrew and Spoehr, 2011). Creative and innovative people who prefer a diverse, tolerant environment that is open to new ideas are also important, usually included in these innovations as the creative class. The success of regions is determined by a combination of three factors: (1) a tolerant environment, (2) access to talent, and (3) technological advancement (Florida, 2002).

The emergence and development of the creative economy represent a new model of economic growth, in which businesses foster creativity, support the fusion of knowledge, and facilitate the use of modern technologies with the aim of creating new markets and jobs. Such an economy enables regions to remain competitive in a global environment and supports long-term growth, particularly in the context of both developed and developing countries (Sung, 2015). Post-industrial development is leading to the emergence of a creative society, in which creative and innovative urban regions with global connections are gaining an advantage. These regions or areas are hubs of new technologies, talent, and capital, while less adaptable regions are undergoing a process of creative destruction as traditional industries give way to innovations. The result is growing inequality between regions, as economic development is concentrated in a relatively small number of highly

rated creative regions (Andersson and Andersson, 2015). The competitiveness of cities is influenced by the phenomenon of the "local trap," in which global actors can override local interests. Urban areas are thus a site of intense interactions between the local and global environments.

Globalization brings new investments and business opportunities, but it can also lead to wage stagnation, rising inequality, and uncertainty. Effective development therefore requires a combination of local resource utilization and a global perspective, with participatory decision-making mechanisms ensuring benefits for the entire community (Zientara, 2009). Regional and urban areas, exposed to interactions between the local and global environments, must navigate the tension between preserving local autonomy and utilising the opportunities of globalisation. The phenomenon of the "local trap" highlights that an excessive focus on global pressures can weaken decision-making at the local level and limit the interests of residents. While globalisation can lead to wage stagnation, deepening regional inequalities, and increased uncertainty, it also offers investment and business opportunities that can support the development of local economies. Innovation is the key driver of regional competitiveness and development. Based on existing empirical evidence, various classifications define types of innovation, but within the scope of this study, we can categorise innovations as:

- imitative innovations, which improve existing products and target new segments of the market;
- sustainable (incremental) innovations, which bring technologically less demanding improvements and maintain performance in existing markets;
- and disruptive innovations, which radically change products, processes, or services, target new markets, and generate a process of creative destruction.

The success of innovation activities is closely linked to the proximity of regional actors and institutions, which fosters interaction, the formation of strong clusters, and agglomeration effects that contribute to regional growth. The goal is to support sustainable productivity growth: process innovations reduce costs, while product innovations enable higher prices and create a competitive advantage through novelty or differentiation. In this way, regions are able to respond to global pressures, support local entrepreneurship, and simultaneously create an environment for economic development based on creative and innovative solutions (Talmaciu, 2012).

Schumpeter viewed the entrepreneur as an agent of creative destruction who transforms industries and enhances their competitiveness. In contemporary models, a firm encompasses both types of expertise, so that its workforce includes both inventors and innovators, which makes it possible to link the creation of new discoveries with their market application. Regional entrepreneurship affects the productivity of established firms, particularly in the service sector; while it may initially have a short-term negative effect, this is followed by a positive impact on firm performance (Andersson, Braunerhjelm and Thulin, 2012). The entry and exit of firms are also shaped by firm characteristics, industrial linkages, and regional institutional conditions, while decentralisation, globalisation, and privatisation lead to spatial and temporal disparities (Zhou, He and Zhu, 2017).

Local creative destruction occurs when the departure of existing firms frees up resources and creates space for new businesses, thereby supporting regional industrial

renewal. The process is predominantly local and depends on local knowledge and social networks, with new technologies often building on existing know-how. Inertia can prevent older firms from adapting, so their departure allows for the recycling of resources, the acquisition of experience, and the promotion of innovation. Regional policies should sensitively support these processes, focus resources on specific sectors, preserve diversity, and avoid excessive protection of failing firms. The exit of young firms often does not lead to renewal, whereas the exit of mature, underperforming firms fosters regional economic growth and innovation through the entry of new players (Pe'er and Vertinskij, 2008). Successful regional economies combine close cooperation among businesses, a skilled workforce, flexible work structures, a robust infrastructure that supports innovation, an innovative culture, and active regional governments (Bosma, Stam and Schutjens, 2011). Equally important for the sustainability of regional economies is the implementation of circular economy principles, which requires changes in management approaches, investments in technologies and infrastructure, as well as close cooperation between businesses, governments, and communities, with public support playing a key role (Dielini et al., 2024).

Regional characteristics that influence creative destruction depend heavily on the firms operating in the region, as these firms and their industries determine the region's technical capacity. A high concentration of firms in the same industry creates cluster regions that host a large number of established players and support the industry's operations, but new technologies often emerge outside these clusters. The more concentrated an industry is, the more significant it is for the local economy and the more difficult it is to implement changes that would radically transform existing competencies and capabilities, because the accumulated regional knowledge of the dominant paradigm obstructs the adaptation of other actors within the system (Gilbert 2012). According to the theory of creative destruction, economic cycles are driven by waves of innovation, with market disruption causing significant economic impacts on major business groups (Neufeld, 2021).

Technological waves are gradually transforming the economy and society, as well as regional and local economic systems. The first industrial revolution boosted production rates with steam engines and sparked urbanisation. The second wave, driven by electricity, improved industrial productivity, transportation, and communication networks, with railroads and steel playing a key role. The third wave introduced construction lines and the automotive industry; the fourth wave brought computers and internet-driven digitisation and automation. The fifth wave, supported by the internet and artificial intelligence, is transforming global trade and innovation, while the emerging sixth wave, focused on sustainability, is promoting renewable energy, eco-friendly practices, and the circular economy. The penetration of creative destruction and artificial intelligence during the sixth wave is creating new sectors, transforming existing ones, and influencing economic policies that promote growth through innovation (Uctu, Tuluce and Aykac, 2024). In the context of the latest technological wave focused on sustainability and the circular economy, theoretical approaches that link economic growth with environmental aspects are coming to the forefront, including industrial ecology. Industrial ecology is one of the modern theoretical approaches to urban and regional development that responds to the need to link economic growth with environmental sustainability. This theory highlights

that industrial systems should function in a way similar to natural ecosystems. It promotes a comprehensive (systemic) approach focused on the efficient management of material flows throughout their entire cycle – from raw material extraction, through production and use, to end-of-life processing. It is based on three principles: technological innovation, interdisciplinary collaboration, and a systems approach (Sumra et al., 2024). In this context, digitalisation plays a key enabling role. It has a non-linear impact on the circular economy, while it may initially slow down the transition, at higher levels, it significantly supports it. Digital technologies help reduce costs, improve resource efficiency, and enable the development of circular business models, while also strengthening cooperation among stakeholders and the active role of consumers (Nham and Ha, 2022).

Regional economic resilience, with creative resilience as a specific approach to regional resilience, where creative factories strengthen a region's ability to respond. The concept of regional economic resilience has become a well-developed concept over the past decade (Sutton et al., 2023). However, resilience is gradually taking on new dimensions in light of dynamic societal developments and the economic cycle, dimensions that are primarily oriented toward collective capability (D'Ambrosio and Migheli, 2025), non-economic regional development drivers (Sutton et al., 2025) and also innovation and creativity (Garcia et al., 2025; Viana et al., 2023; Boschma and Pinto, 2015).

There are two main approaches to the concept of regional economic resilience. The first approach views it as a region's ability to withstand a shock and quickly return to its original state. The second, adaptive approach emphasises the region's ability to change and adapt, noting that following a shock, there may not be a return to the original state, but rather the creation of a new economic structure (Bristow and Healy, 2014). System resilience cannot be understood merely as the ability to return to its original state after a shock, but rather as the ability to change and adapt. As opposed to stability, which emphasises immutability, resilience points to dynamics, variability, and the possibility of transition between different states. Systems thus need not have just one equilibrium, but rather multiple possible states between which they can shift. This is why being resilient means changing and adapting rather than everything just going back to how it was before (Reggiani, De Graaff and Nijkamp 2002).

The case of the Oulu region in Finland shows that creative destruction can cause serious problems for regional development, particularly when dominant and long-established industries collapse. This case also highlights the distinction between regional resilience, which generally focuses on a region's ability to recover from a shock, and creative resilience, a specific concept that takes into account creativity, innovation, and community capacities for adaptation. Oulu's experience served as the basis for creating a framework for creative resilience, which consists of three pillars: knowledge creation, entrepreneurial initiatives, and a strong community spirit (Simonen, Herala and Svento, 2020). This framework shows that, to successfully manage shocks, recovery alone is not enough; it is necessary to actively create the conditions for the region's transformation and further development.

3 Methodology

The study aims to systematise theories of regional development, creative destruction and regional resilience, and to verify their applicability by examining specific examples of investment impulses at various stages of investment cycles. These impulses will be contextualised within the specific conditions of the Košice region ecosystem, which is currently facing change as a result of the Volvo investment and the start of production scheduled for early 2027. More specifically, the study will examine the impact of the investment on regional development, creative destruction and regional resilience in its early stage in the selected regional ecosystem, reflecting trends identified by existing empirical evidence.

The object of study is the Košice region, which was chosen through an intentional selection process based on its suitability for a case study. The key selection criterium was the implementation of a major foreign direct investment in the Košice region. Volvo's investment of 1.2 billion € is aimed at producing 250,000 vehicles per year, which is one of the largest investments in the automotive sector in Slovakia and also the most significant project in eastern Slovakia, creating a total of 3,300 jobs and providing a boost to the wider region (Ministry of Economy of the Slovak Republic, 2024).

This descriptive case study of the Košice Region goes beyond a simple description of the existing reality, and there are several reasons for this approach. Firstly, the study examines Volvo's investment in the Košice Region and its impact on regional development, creative destruction, and the resilience of the regional ecosystem within a specific context. This context can only be described, not influenced or modified, during the time period covered by our research. Another reason for this approach is that the study provides a comprehensive explanation of the situation in the selected Košice region, linking it to theories of regional development, creative destruction, regional ecosystem sustainability, and regional resilience. At the same time, this approach enables us, in the initial phases of our Volvo investment case study, to reflect on current empirical study trends in other regions, identify the basis for future research questions, and subsequently focus on evaluating and confirming these questions. However, this ambition will only be possible based on data collected once the impacts of the expected development begin to take effect. In this regard, it will be possible to confirm or deny the trends identified in other regions and in the studies mentioned in the theoretical section in 2027 and later.

The Košice Region is one of the eight regions of the Slovak Republic classified as NUTS 3 units. The Košice Region is located in the east of the Slovak Republic and borders Ukraine and Hungary. As part of Eastern Slovakia (NUTS 2), the Košice Region, given its location on the eastern border of the European Union, receives direct support from the Government of the Slovak Republic in attracting foreign investment. In response to this fact, the Košice Region has long been supported through the government's structural policy, specifically in the form of support for the development of industrial parks. In total, there are 78 industrial parks in Slovakia, and 9 parks are currently functioning in the Košice Region; thus, the Valaliky Industrial Park, together with Volvo's investment, represents another such concentration of economic activity (MH Invest, 2025). Košice Region has a total population of nearly 800,000, with the regional capital city, Košice, having a population of approximately 230,000. This is a region that has seen significant

investment over the past twenty years in the form of new companies in the IT and automotive sectors. Currently, the arrival of a new investor is providing fresh momentum for the region's development, with expectations of increased employment and production. In this case study, we aim to answer the following research questions:

- RQ1: What impact will the opening of the Volvo factory have on the economy and development of the Košice Region?
- RQ2: What opportunities and risks does the Volvo factory present for the regional development of Košice?
- RQ3: How can the launch of production at the Volvo factory in Valaliky support the regional development of the Košice Region through knowledge transfer, business support, and the development of a sense of community?
- RQ4: How will the Volvo factory influence the cooperation within the regional ecosystem and regional resilience?

The theoretical part of this study was conducted using an analytical-synthetic method and a systematic literature review, with data primarily sourced from the Web of Science and SCOPUS databases. The theoretical findings will then be applied to the conditions of the Košice Region in the analysis of the automotive industry. The empirical part is based on the analysis of statistical and secondary data from different sources, such as the Statistical Office of the Slovak Republic, Finstat, Slovak Investment and Trade Development Agency, ministries, and the official website of the Valaliky Industrial Park. Finally, the findings will be systematized and summarised in order to formulate conclusions regarding the development, structure, and perspectives of industry in the Košice Region.

4 Case Study: Opportunities and Risks of the Volvo in Eastern Slovakia

The relevance of this study is in the connection between the theory of creative destruction and the real development of the Košice Region. The region serves as an example of transformation driven by investment and allows an analysis of the impact of a new investor's arrival on its future development. The Košice Region comprises eight districts: Košice I through IV, Košice-okolie, Michalovce, Spišská Nová Ves, Trebišov, Rožňava, Gelnica, and Sobrance. The Košice region is one of Slovakia's most important industrial regions, with the following sectors and major companies:

- Metallurgy: U. S. Steel Košice (Košice II) – production of steel and steel products, Kovohuty Krompachy – refined copper production.
- Mechanical engineering: VSS Košice (Engineering Works Košice) – production of special vehicles and machinery; Embraco Slovakia (Spišská Nová Ves) – compressors for refrigeration and air-conditioning systems; Prako-Príruby (Prakovce) – metal components and engineering production.
- Automotive industry: Marelli (formerly Magneti Marelli).
- Electrical engineering industry: BSH Drives and Pumps (Michalovce), SEZ Krompachy (Slovak Electrical Works) – electrical equipment and switchgear; Panasonic (formerly Matsushita Audio Video Slovakia, Krompachy) – consumer electronics manufacturing.

- Chemical industry and mineral processing: Chemko Strážske – production of chemical products; ŽELBA (Spišská Nová Ves) – mineral processing and metal ores treatment.
- Food and consumer goods industry: Frucona Košice – food production; Ryba Košice – fish processing; Zempmilk (Michalovce) – dairy production; MPC (Spišská Nová Ves) – bakery and confectionery products; Nový domov Holding – furniture manufacturing. (Statistical Office of the Slovak Republic, 2026; Vavrinčík, 2004).

The most important industries in the region are metallurgy, chemicals, and electrical engineering. Metallurgy is the most significant (60% of production, 50% of exports), with U.S. Steel Košice operating there. The chemical industry is represented by Chemko Strážske, and the electrical engineering sector is dominated by foreign companies, such as BSH Drives and Pumps Michalovce (Slovak Investment and Trade Development Agency, 2026a). Investment aid in Slovakia is aimed at supporting industry and job creation, with a particular focus on less developed regions. The Košice Region is one such area and is therefore eligible for a higher aid intensity, which increases its attractiveness to investors (Slovak Investment and Trade Development Agency 2026b).

The automotive industry is the most important sector of the Slovak economy and a major source of investment, employment, and innovation, accounting for approximately 10.4% of GDP and 46.5% of industrial output, and employing around 255,000 people. Slovakia ranks among the world leaders in per capita automobile production, with an output of approximately one million vehicles per year and the presence of several global manufacturers. The importance of the automotive industry is gradually expanding into the Košice Region, where the supplier sector and investment are growing (Slovak Investment and Trade Development Agency, 2025).

The Regional Innovation Strategy of the Košice Region 2021–2030 indicates that the automotive industry is one of the most important sectors of industrial production in the region and is closely linked to metallurgy, mechanical engineering, and electrical engineering, primarily functioning as a subcontractor. In terms of regional distribution, it is concentrated primarily in the districts of Košice and Michalovce, with the majority of revenue coming from Košice. At the same time, the strategy highlights the sector's declining trend, as both the number of companies and employees has been declining in recent years. Given the automotive industry's national significance, its development in the Košice Region takes on particular importance. According to the Regional Innovation Strategy of the Košice Region 2021–2030, this sector has so far been characterised primarily by a focus on subcontracting without the presence of final automobile production; therefore, its further development takes on new significance in connection with the planned launch of automobile production after 2027 (Košice Self-Governing Region, 2021).

The Valaliky Industrial Park, with the Volvo Car Košice project, represents a major strategic investment in eastern Slovakia focused on the production of electric vehicles. The factory is expected to reach a capacity of up to 500,000 cars per year in two phases and will become one of the largest electric vehicle manufacturing centres in Europe. The production process encompasses the entire automotive supply chain, from melting and stamping parts, through body manufacturing and painting, to the assembly of batteries

and electric motors, and the final assembly of vehicles. Connections with external suppliers and logistics networks play a significant role, reflecting the global nature of the automotive industry. Through this project, Slovakia is moving toward modern electromobility, leveraging its strong tradition in the automotive industry, skilled workforce, and strategic location in Europe. For the Košice Region, the project holds significant development potential, particularly in the form of the creation of approximately 4,000 to 5,000 jobs and additional indirect employment opportunities in services, logistics, and subcontracting firms. The investment will also support infrastructure development and may attract additional investors, thereby contributing to the region's transformation from a traditional metallurgical and engineering hub into a modern industrial hub focused on new technologies. On the other hand, the project also entails certain risks, particularly environmental and infrastructure-related ones. Increased pressure on air quality, water resources, and the transportation network is expected, with the impact on air quality assessed as slightly negative. The project may also create pressure on the housing market, drive up real estate prices, and place increased demands on public infrastructure. In the long term, it also poses a risk of the region's dependence on a single dominant investor and on the development of the global electric vehicle market. The Valaliky Industrial Park project with the Volvo plant is a combination of private investment and state support, and has been officially designated a significant investment by the Government of the Slovak Republic. The state plays an important role, particularly in building infrastructure and supporting the development of the entire strategic park (Ministry of the Environment of the Slovak Republic, 2024) (Table 1).

Table 1. Technological waves in the Košice Region

Technological waves	1st wave	2nd wave	3rd wave	4th wave	5th wave	6th wave
Focus	Heavy industry (steam)	Engineering and Electrical Engineering	Information Technology	Automation and globalization	Digitalization and Startups	Green Economy and AI
Innovation waves	Hydropower Textiles Iron	Steam Rails Steel	Electricity Chemicals Internal combustion engines	Petrochemical industry Electronics Aviation	Digital Network Software New Media	Digitalization (AI, IoT, AV, robots, and drones) Clean technologies
Time period	60 years	55 years	50 years	40 years	30 years	25 years
Major companies	East Slovak Steelworks (U. S. Steel Košice)	ZTS Košice	ESET	Siemens	Deutsche Telekom IT Solutions Slovakia (Košice)	Volvo
	Košice brickyard (does not exist)		Asseco Central Europe		Ness Digital Engineering (Košice)	

Source: modified by authors based on (Finstat, 2026; Uctu, Tuluce and Aykac, 2024)

The region's transformation from a traditional centre of metallurgy and engineering into a modern industrial hub focused on electromobility and innovative technologies is creating new drivers of development, but at the same time increasing its vulnerability to external factors. For this reason, the main opportunities and risks associated with this development can be identified in aspects included in Fig. 1.

	Opportunities		Risks
+	Creation of new jobs (number of university graduates and high school students)	-	The region's high dependence on strategic foreign investors
+	Development of electromobility and new technologies	-	The automotive industry's impact on economic cycles as demand for vehicles, which is closely tied to the economic situation in global markets.
+	Multiplier effect of investments, specifically through the arrival of subcontractors and logistics companies	-	Environmental impact
+	Improvement of infrastructure, specifically through the construction of an industrial park	-	Shortage of skilled labor
+	Integration into global value chains, specifically through automotive exports	-	Regional disparities

Source: authors

Fig. 1. Opportunities and risks associated with development in Košice Region

The entry of a strategic investor aims to strengthen the regional resilience of the Košice Region, particularly through job creation, an influx of investment, and integration into global value chains. At the same time, however, risks associated with dependence on a single sector and external factors persist, highlighting the limitations of regional resilience itself. Therefore, the need to develop creative resilience comes to the fore, enabling the region's transformation toward a more diversified, innovative, and long-term sustainable economy.

5 Discussion

Creative resilience represents an extension of the concept of regional resilience, as it emphasises not only a region's ability to recover from shocks, but also its capacity to generate new development trajectories through innovation and entrepreneurship. The Volvo factory in Valaliky is owned by a global entity, its infrastructure is used by city residents and regulated by public administration institutions, thereby creating public-private partnerships.

As a result, urban infrastructure becomes "locally unbranched but globally interconnected," offering benefits to all parties, with predictable returns for investors and infrastructure use for residents. This case shows that local and global scales are not necessarily at odds with one another and can lead to "win-win" outcomes (Zientara, 2009). In connection with the three pillars of creative resilience (Simonen, Herala and

Svento, 2020), when Volvo begins production, it is necessary to provide support, mainly knowledge creation, support for entrepreneurship and development of community spirit.

Knowledge creation is important because technology transfer, training, and collaboration with local educational institutions bring new knowledge and skills. Among the educational institutions is the Technical University of Košice, which mainly offers study programs focused on mechanical engineering, electrical engineering, mechatronics, automotive manufacturing, and industrial engineering. Pavol Jozef Šafárik University in Košice also plays a significant role, particularly in the fields of natural sciences and IT. At the secondary education level, there are secondary schools focusing on technical fields such as mechatronics, electrical engineering, mechanical engineering, and automotive repair, which are gradually being integrated with practical training through dual education programs.

Support for entrepreneurship is provided through the development of supply chains, investment opportunities, and a favourable business environment that attracts foreign direct investment and supports local entrepreneurship.

Several companies linked to the automotive industry are already operating in the region, particularly in the Kechnec Industrial Park, which is home to companies such as Getrag Ford Transmissions Slovakia (gearbox manufacturing), Magneti Marelli (automotive components), and Molex (electronic components for the automotive industry). These companies form the foundation for the development of subcontractor networks, and Volvo's presence may further expand this ecosystem to include new businesses in the fields of parts manufacturing, battery systems, logistics, and services.

Community spirit is fostered through job creation, social networks, and the plant's participatory activities in the region, which strengthen social capital and a sense of community belonging. The investment helps reduce the outflow of young people from the region and encourages the return of skilled workers.

This framework reflects the need for local authorities and businesses to act both locally and globally. Locally because of critical assets such as social capital, human resources, and infrastructure, and globally, because globalization offers new investment and business opportunities that boost the attractiveness of the region. Experience from Romania shows that low trust and weak social capital can slow down the development of innovation clusters (Talmaciu, 2012). This is why it is important in the Košice Region to implement mechanisms of participatory democracy and careful decision-making. These mechanisms will ensure that development projects, such as the Volvo project, serve the interests of the entire community.

In summary, although it might initially seem that regional development will be driven primarily by the automotive industry, based on the available empirical evidence mentioned, it should be predicted that sustainability will not depend solely on the Volvo car plant itself. Based on the trends identified in the theoretical part of the study, it is currently possible to assume that, in the case of the Košice region under examination, overall innovation potential will be supported and that mutual relationships and regional resilience within the ecosystem will be strengthened.

In this sense, Volvo investment serves as an important example and confirmation of models of regional collaboration, where a regional ecosystem operating on a triple-helix

or quadruple-helix model supports the growth of the entire regional ecosystem (Mitaľ, Ručinská and Ručinský, 2025).

Finally, it should be noted that Volvo's investment is only just beginning to be realised with the start of production in early 2027. This study describes the current phase of the investment's implementation. Within a few years, this descriptive comprehensive overview can be strengthened with additional quantitative and qualitative economic and social data.

Based on these findings, it can be concluded that the further development of the region should focus on strengthening creative resilience, which enables the transition from passive adaptation to active transformation of the economy, particularly in the context of the transition from traditional industrial sectors to digitalisation, artificial intelligence and green economy within the current fifth and sixth technological waves.

6 Conclusion

At the regional level, creative destruction manifests itself as a process in which long-established firms and industries give way to new technologies, products, and business models, leading to economic shocks and pressure to adapt(Neufeld, 2021; Hospers, 2005; Schumpeter, 1934). Contrary to traditional regional resilience, which focuses on returning to the original state, creative resilience is a specific approach that emphasises a region's ability to transform itself and create new economic structures in response to change (Bristow and Healy, 2014; Reggiani, De Graaff and Nijkamp, 2002). The case of the Oulu region in Finland shows that effectively managing these shocks requires three pillars of creative resilience: knowledge creation, entrepreneurial initiatives, and a strong community spirit (Simonen, Herala and Svento, 2020).

The key benefit of the study is therefore its contribution to the ongoing debate, highlighting the importance of creative resilience. Regional resilience and creative resilience represent distinct yet complementary approaches to regional development. While regional resilience focuses on the ability to adapt and recover after a shock, creative resilience emphasises active transformation through innovation and entrepreneurship. In the context of current technological developments and ongoing technological waves, as well as within the conditions of the Košice Region, it therefore appears essential not only to maintain regional resilience but also to systematically develop the principles of creative resilience as a prerequisite for the region's long-term and sustainable development.

References

Aghion, P., Howitt, P.: A model of growth through creative destruction. Econometrica **60**(2), 323–351 (1992)

Aghion, P., Howitt, P.: The creative destruction approach to growth economics. Eur. Rev. **31**(4), 312–325 (2023). https://doi.org/10.1017/S1062798723000212

Andersson, Å.E., Andersson, D.E.: Creative cities and the new global hierarchy. App. Spat. Anal. **8**, 181–198 (2015). https://doi.org/10.1007/s12061-015-9141-7

Andersson, M., Braunerhjelm, P., Thulin, P.: Creative destruction and productivity: entrepreneurship by type, sector and sequence. J. Entrepreneurship Public Publicy **1**(2), 125–146 (2012). https://doi.org/10.1108/20452101211261417

Andrew, A., Spoehr, J.: Beyond the creative quick fix conceptualising Creativity's role in a regional economy. Creativity and solving global challenges. In: Drivers of Innovation, Entrepreneurship and Regional Dynamics, pp. 369–385. Springer, Berlin (2011). https://doi.org/10.1007/978-3-642-17940-2_18

Boschma, R., Pinto, H.: Introduction Resilient territories. In: Pinto, H. (ed.) Resilient Territories: Innovation and Creativity for New Models of Regional Development, pp. 1–10. Cambridge Scholars Publishing, Newcastle upon Type (2015)

Bosma, N., Stam, E., Schutjens, V.: Creative destruction and regional productivity growth: evidence from the Dutch manufacturing and services industries. Small Bus. Econ. Dordrecht **36**(4), 401–418 (2011). https://doi.org/10.1007/s11187-009-9257-8

Bristow, G., Healy, A.: Building resilient regions: complex adaptive systems and the role of policy intervention. Raumforschung und Raumordnung-Spat. Res. Plann. **72**(2), 93–102 (2014). https://doi.org/10.1007/s13147-014-0280-0

Capello, R., Dellisanti, R.: Creativity in cultural and creative industries as a source of regional adaptive resilience. In: Pascariu, G.C., Ţigănaşu, R., Kourtit, K., Nijkamp, P. (eds.) Resilience and Regional Development, pp. 10–29. Edward Elgar Publishing, Cheltenham (2023). https://doi.org/10.4337/9781035314058.00008

D'Ambrosio, A., Migheli, M.: Regional resilience through the lenses of the capability approach. J. Evol. Econ. **35**, 833–353 (2025). https://doi.org/10.1007/s00191-025-00915-8

De Bernard, M., Comunian, R., Gross, J.: Cultural and creative ecosystems: a review of theories and methods, towards a new research agenda. Cultural Trends **31**(4), 332–353 (2022). https://doi.org/10.1080/09548963.2021.2004073

Dielini, M., et al.: Security of sustainable development of the agricultural sector of the region based on digitalization and circular economy: a case study of Ukraine. J. Disciplinary Res. **14(2)-XLIV**, 104–108 (2024) https://evnuir.vnu.edu.ua/handle/123456789/29412

Feld, L.P., Schnellenbach, J., Baskaran, T.: Creative destruction and fiscal institutions: a long-run case study of three regions. J. Evol. Econ. **22**, 563–583 (2012). https://doi.org/10.1007/s00191-011-0264-y

Finstat: Database of Financial Results of Slovak Companies. (2026). https://www.finstat.sk/databaza-financnych-udajov?Activity=automobilový+priemysel&Region=košický&SalesFrom=&Employee=&Years=2009&Years=2008&PerPage=20&Sort=&Tab=

Florida, R.: The Rise of the Creative Class. Basic Books, New York (2002) http://www.washingtonmonthly.com/features/2001/0205.florida.html

Garcia, O., Antonio, J., Gras, A., Nicole, E., Moya Tudela, P.: Creativity as a driver to resilience in cities: the more creative the more resilient. Street Art Urban Creativity Sci. J. **11**(5), 67–89 (2025). https://doi.org/10.62161/sauc.v11.5836

Gilbert, B.A.: Creative destruction: identifying its geographic origins. Res. Policy. **41**(4), 734–742 (2012). https://doi.org/10.1016/j.respol.2011.11.005

Hospers, G.-J.: Joseph Schumpeter and his legacy in innovation studies. Knowl. Technol. Policy 18, 20–37 (2005). https://doi.org/10.1007/s12130-005-1003-1https://onlinelibrary.wiley.com/doi/full/10.1111/j.1468-0270.2009.01895.x

Huang, X., Dong, J., Li, X.: Fintech, technological innovation and regional economic growth: theoretical modeling and empirical evidence. China Econ. Rev. **91** (2025). https://doi.org/10.1016/j.chieco.2025.102397

Irandoust, M.: Employment and technology: creative creation or creative destruction? Asymmetric Anal. Aust. Econ. Pap. **63**(2), 201–219 (2024). https://doi.org/10.1111/1467-8454.12327

Košice Self-Governing Region: Regional Innovation Strategy of the Košice Region 2021–2030. Košice Self-Governing Region, Košice (2021)

MH Invest: Comprehensive Information on PZ/PP as of June, 2025 (2025). https://www.priemyselneparkyslovenska.sk/sk/#

Ministry of Economy of the Slovak Republic: The European Commission today officially approved investment aid for Volvo, which will create 3,300 jobs in eastern Slovakia (2024). https://www.mhsr.sk/press/europska-komisia-dnes-oficialne-schvalila-poskytnutie-investicnej-pomoci-pre-volvo-ktora-na-vychodne-slovensko-prinesie-3-300-pracovnych-miest

Ministry of the Environment of the Slovak Republic: Final Statement – Volvo Car Košice (Valaliky Strategic Park). Č. 5653/2024–11.1.1, 19. 04. (2024)

Mitaľ, O., Ručinská, S., Ručinský, M.: Innovation potential of the Košice region ecosystem. In: Hintea, C., Balogh, M. (eds.) Fostering Local and Regional Innovation through Partnerships - the Role of Universities, pp. 174–191. Accent, Cluj-Napoca (2025)

Neufeld, D., Ma, J.: Long waves: The history of innovation cycles: Visual capitalist. **30** (2021) https://www.visualcapitalist.com/the-history-of-innovation-cycles/

Nham, N.T.H., Ha, L.T.: Making the circular economy digital or the digital economy circular? Empirical evidence from the European region. Technol. Soc. **70** (2022). https://doi.org/10.1016/j.techsoc.2022.102023

Pavelea, A.M., Pavel, A., Neamţu, B., Kourtit, K.: Is the creative class a shock absorber against an economic recession? Spatial labour market resilience in Romania. Reg. Stud. **57**(12), 2457–2472 (2023). https://doi.org/10.1080/00343404.2023.2194314

Pe'er, A., Vertinskij, I.: Firm exits as a determinant of new entry: is there evidence of local creative destruction? J. Bus. Ventur. **23**(3), 280–306 (2008). https://doi.org/10.1016/j.jbusvent.2007.02.002

Pittz, S.F., Pittz, T.G.: Towards an ethical awareness of entrepreneurs: a Nietzschean perspective on creative destruction. Entrep. Reg. Dev. **36**(5–6), 707–721 (2024). https://doi.org/10.1080/08985626.2024.2305193

Reggiani, A., De Graaff, T., Nijkamp, P.: Resilience: an evolutionary approach to spatial economic systems. Netw. Spat. Econ. **2**, 211–229 (2002). https://doi.org/10.1023/A:1015377515690

Schumpeter, J.A., Swedberg, R.: The Theory of Economic Development. Routledge (2021)

Schumpeter, J.A.: The Theory of Economic Development. Oxford University Press, London (1934)

Simonen, J., Herala, J., Svento, R.: Creative destruction and creative resilience: restructuring of the Nokia dominated high-tech sector in the Oulu region. Reg. Sci. Policy Pract. **12**(5), 931–954 (2020). https://doi.org/10.1111/rsp3.12267

Slovak Investment and Trade Development Agency: Investment Incentives in Slovakia. (2026b)

Slovak Investment and Trade Development Agency: Košice Region. (2026a). https://www.sario.sk/sk/kosicky-region-1

Slovak Investment and Trade Development Agency: I Automotive Sector in SLOVAKIA. (2025)

Statistical Office of the Slovak Republic: Košice Region – Characteristics of the Region. (2026). https://slovak.statistics.sk/wps/portal/ext/themes/regional/kosicky%20kraj/about/!ut/p/z1/jZHLboMwEEW_JQu29hjzMN05VDGuaArhEepNRSJKqHhEhIbfL02yqdTSzG6kc49Gc7HCGVZtfq7KfKi6Nq-n_VVZb6kfsOWScBB2TECu11GSspA-ewRvL4AruGfYPgDzhQmSe8nGCSkFTrG6Jw9_DIf78jOAmtdvsbogoS2vBmbaMBmiBCCNRBgY8wCPzRuwigKdO4ZwHzcvK5Cxq7PIt3QAcgPmnvTfmU9YVbsGjfsGAbKYwyxCqUnBYbZufbfE2x1lJVZ98V70RY8%2D%2D6m8wzAcTw8aaDCOIyq7rqwLtO8aDX6LHLrTgLOfJD42SZJBJT_M-uzzxeILrqPYlQ!!/dz/d5/L2dBISEvZ0FBIS9nQSEh/

Sumra, K.B., Siddique, H., Afzal, S., Qazi, A.: Sustainable smart cities: promotion of circular economy in urban GCC regions. J. Sci. Technol. Policy Manag. **17**(1), 60–99 (2024). https://doi.org/10.1108/JSTPM-01-2024-0025

Sung, T.K.: The creative economy in global competition. Technol. Forecast. Soc. Chang. **96**, 89–91 (2015). https://doi.org/10.1016/j.techfore.2015.04.003

Sutton, J., Arcidiacono, A., Torrisi, G., Arku, R.N.: Regional economic resilience: A scoping review. Prog. Hum. Geogr. **47**(4), 500–532 (2023). https://doi.org/10.1177/03091325231174183

Sutton, J., Mendez, P., Arku, G.: Regional economic resilience and post-crisis growth paths: hysteresis and homeorhesis. J. Econ. Geogr. **lbaf034** (2025). https://doi.org/10.1093/jeg/lbaf034

Talmaciu, M.: Considerations regarding the development of Romanian regional economies through innovation and entrepreneurship. Procedia Econ. Finan. **3**, 914–920 (2012). https://doi.org/10.1016/S2212-5671(12)00250-X

Uctu, R., Tuluce, N.S.H., a Aykac, M.: Creative destruction and artificial intelligence: the transformation of industries during the sixth wave. J. Econ. Technol. **2**, 296–309 (2024). https://doi.org/10.1016/j.ject.2024.09.004

Vavrinčík, P.: The possibilities of industry development in Košice region. Econ. J. **52**(7), 878–890 (2004) https://www.sav.sk/journals/uploads/051609112004_7_Vavrincik.pdf

Viana, L.F.C., Hoffmann, V.E., Miranda Junior, N.D.S.: Regional resilience and innovation: paper profiles and research agenda. Innovation Manag. Rev. **20**(2), 119–131 (2023). https://doi.org/10.1108/INMR-11-2021-0225

Zhou, Y., He, C., Zhu, S.: Does creative destruction work for Chinese regions? Growth Chang. **48**(3), 274–296 (2017). https://doi.org/10.1111/grow.12168

Zientara, P.: A few critical remarks on globalisation, democracy and spatiality. Econ. Aff. **29**(2), 56–61 (2009). https://doi.org/10.1111/j.1468-0270.2009.01895.x

Circular Ecological Heritage Communities: Regeneration and Evaluation Strategies for Historic Villages, Landscapes and Rural Territories (CHEC 2026)

Regenerating Historical, Landscape, and Environmental Heritage Through the Municipal Territorial Plan in Lombardy

Dina Jovanovic(✉), Martina Mulinacci, Davide Cerati, and Daniela Oreni

Politecnico di Milano, 20133 Milan, Italy
dina.jovanovic@polimi.it

Abstract. Heritage-induced urban and landscape regenerations are imperative in highly urbanised territories. Integrated spatio-temporal historical knowledge needs to directly inform the legal documentation of municipal planning. Nevertheless, the methodological and application gap exists, separating academic heritage research from operational governmental practices. This paper demonstrates a GIS-based methodology used to develop a variant of the Territorial Plan (PGT) for Desio, a medium-sized municipality in Lombardy. The collaboration between university research teams and municipal planning offices resulted in the structured implementation of historical cadastral data and geospatial information, spanning from the 18th century until the present, to trigger a synchronic reading of landscape and urban evolution. The outcome was a design of the Municipal Ecological Network (REC), in the form of two multifunctional green belts. Their role is to reconnect the dispersed farmsteads and religious buildings, fragmented historical woodlands, and rural road networks through contextualised actions within the official planning documentation. The ecosystem enhancement targets 1,538,000 m^2 of currently undeveloped free areas, improving carbon storage, water regulation, and socio-cultural services. The urban regeneration relies on reconnecting these areas with available resources supported by the historical and cultural values of the local territory. The scalability and replicability of the approach are reflected in the applied case study, where historically informed GIS spatial databases can be intersected with regulatory planning frameworks on various governmental levels, aligning regeneration actions with EU policies and objectives.

Keywords: Historical Cartography · Urban Regeneration · Municipal Ecological Network · Municipal Territorial Plan (PGT) · Heritage-induced Planning · Geographic Information System (GIS)

1 Introduction

Historical data and spatio-temporal analysis of the territory are imperative for the all-encompassing and contextualised development of urban planning regulations and actions. Alongside the state of the built and natural heritage in the current setting, data-informed analyses recover the negative impacts of previous urbanisation decisions and

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 455–470, 2027.
https://doi.org/10.1007/978-3-032-30518-3_29

uncontrolled city growth. Present urban planning practices demonstrate that most of the integrated historical research is static, resulting in uneven and under-documented data on heritage in official records, especially in smaller municipalities [1, 2]. Nowadays, digital technologies such as Geographic Information Systems (GIS) are widely used to build databases and manage government urban plans, making the incorporation of structured and systematised heritage information much easier. Therefore, the role of historical information and spatial-temporal data describing existing or lost territorial heritage not only informs the urban planning decisions but should actively influence them. With the creation of a multilayered database containing both contemporary and historical data, planning documents can also address other underlying issues, such as social context, cultural identity, sustainable mobility and landscape regeneration [3].

The paper presents collaboration between the municipal planning office and university researchers as a fundamental step in creating analytical frameworks. This includes structured repositories of historical and contemporary data, further used in the creation of the *Piano di Governo del Territorio* (PGT)[1], the most important Territorial Plan on the municipal level in Lombardy. The research covers studies carried out by several research teams to inform knowledge-based territorial investigations and formulate strategies and actions for the general variant of the PGT of the Municipality of Desio. This plan is implemented in accordance with Regional Law no. 31/2014, 'Provisions for the reduction of land consumption and for the redevelopment of degraded land'. Currently in force is a partial variant of the PGT, made in 2021[2]Currently in force is a partial variant of the, which will be replaced by the ongoing general variation to the PGT developed in 2025, together with Politecnico di Milano and external working groups. The additional value of this collaboration is that the university provided the municipal planning office with experts in the field, especially regarding the intricate, layered past of the urban and rural tissue, making the theory of urban studies and regeneration methodology feasible and well-documented.

Desio is a municipality in the north of Milan, located in the province of Monza and Brianza. Until the second half of the 1800s, the town held an important position, being at the intersection of major roads, created predominantly by agricultural and religious activities [4]. By the end of the century, and with the arrival of the 1900s, heavy industrialisation took hold, leading to the abandonment of the historical centre and fragmentation of the landscape. The result was an incoherent town organisation with fragile heritage characteristics that urgently needed reevaluation and regeneration.

For these declining areas, both in value and use, often located on the outskirts of big hub towns in northern Italy, the implementation of urban regeneration actions is crucial [5]. The planning regulations target not only the enhancement of the present state of natural and built environment through infrastructure, cultural assets, and overall sustainability, but also involve improvements in economic and social conditions. Moreover, urban and environmental regeneration is closely related to the 'Agenda 2030', developed by the United Nations General Assembly in 2015. Specifically, Sustainable

[1] PGT is a territorial planning instrument introduced through the regional law 12/2005, and it is specific to the regulatory framework for all Lombard municipalities.

[2] Consultable through the geoportal of the Lombardy region that provides access to territorial documentation coming from the public administration: https://www.multiplan.servizirl.it.

Development Goal (SDG) 11 relates to the concepts of urban development, planning, and improvement of sustainability in the cities [6]. The SDG goals require the creation of the Integrated Urban Development Model (IUDM). Amongst a collection of approaches toward urban planning, the model also bridges the gap between various governmental levels of territorial management and their resources. Applied urban models often require complex systems and data organisation, but with the development of GIS technologies, the proposed approaches can be managed across scales [7].

Subsequently, the European Union investigated multi-level governance frameworks to bring urban regeneration strategies as a part of local planning activities [8]. They often highlight the involvement of stakeholders, openness to the community, and integration of multiple data typologies into project preparation [9]. With the real case studies, the urban regeneration strategies have shown that they can significantly alter the urban tissue, both on micro and macro scales, shifting the change from only built to broader cultural impact [10]. Further in the European context, the strategic policies are being developed to focus on establishing good practises for revitalisation and reuse of undervalued space and communities [11]. The European Framework for Action on Cultural Heritage, established in 2018, uses concepts, approaches, and methodologies in common with the 'Agenda 2030', linking the goals for creating sustainable, inclusive and resilient European cities to the cultural heritage. This framework aims to apply regeneration through heritage, both in urban and rural contexts [12]. By following these policies and frameworks, the paper investigates the possibility of achieving 'urban renaissance' through the Municipal Territorial Plan in Desio.

2 Methodology for Data Implementation

Mentioned recommendations and reports conclude that cultural heritage behaves like a catalyst for urban environmental regeneration and development. It is one of the most important elements for redeveloping cities which were previously consumed by rapid industrial growth, as in the case of the Municipality of Desio. The application of these concepts directly in the official planning documentation can help recover the fragmented spatial qualities, lost cultural identity, and a sense of community [13]. While there are many academic discussions and research advancements on the topics of urban innovation in Italy, the public sector still shows financial weakness and the private one is often discouraged due to the long-term crisis. Therefore, the planners, governments, and third sector involved in urban regeneration would rather turn to urban reuse, readapting existing structures and areas within a newly established network [14].

Activating urban regeneration policies for territories takes into account the large number of stratified, identity-related, environmental, cultural, and social values embodied within. The centuries of different spatial uses greatly impacted the relationship between historical development, architectural components, and landscape features. To untangle the intricate and layered contexts in which these elements reside, a collection of various historical and contemporary data is needed. Historical cartography, together with archival documents and available indirect sources, enables a profound reinterpretation and critical analysis of the evolution of territories over time. This is fundamental for their enhancement and sustainable regeneration, as it allows the rediscovery of values

and strengths, reactivating relationships that have often been erased by the urban planning policies of the recent past, which were almost exclusively aimed at meeting housing and productive needs. Frequently, there is a methodological gap between historical analysis supporting the heritage and contemporary planning [3], but the research's methodology strives to establish standardised practices integrated into public documentation.

The research paper investigates established practices of GIS implementation in urban studies, using it for historical data systematisation, linkage to existing databases, and programming environmental actions. The applied procedure is built on the extensive heritage GIS literature, demonstrating that this geospatial technology is imperative for collecting, reading, and managing cultural heritage across the territory [15–17]. Nevertheless, the employed methodology further advances through its use in the real case study, shaping the dialogue with governmental administrations and their requirements. The georeferenced maps, vectorised historical cadasters, and multi-layered database connecting historical and contemporary data are not just the research outputs. They are embedded into a legally binding PGT as a tool available to the planning office. Unstructured and dispersed historical sources became an operational normative, digitised and linked with present standards of spatial representations and cataloguing. Moreover, the valorisation and quantification of landscape features, such as the rural road network, historical production and land usage, toponyms, and woodland boundaries, transform them into tangible targets in municipal planning.

This methodological approach formed the basis of the knowledge-building investigations carried out for the Municipality of Desio. It provided useful indications for outlining the new PGT, while at the same time, following the supra-municipal regulations. The activities of historical investigation and the construction of the innovative digital tool for consulting the collected data began with the georeferencing and digitisation of all available historical cartography for the Desian territory. The organisation of historic databases and the construction of information layers for their archiving and consultation today represent a fundamental tool [18]. This approach allows the reading of historical maps alongside current plans and cadasters, rendering the relationship between form and use, and between the past and present of a territory [19]. Historical documentation spans from the early decades of the eighteenth century to the present day, including cadastral maps, geographic maps, and orthophotos, across different scales. For example, the reading of the historical evolution of land use over the centuries revealed the most significant data for defining strategies for 'green regeneration' interventions. Processed historical data facilitated: identifying the system of farmsteads forming a belt around the built-up core of Desio, tracing rural paths connecting the various agricultural settlements, today still partly recognisable in the local road network, and marking the portions of green space historically devoted to deciduous woodlands [20].

The research investigated different environmental frameworks, prioritising PESTEL, to evaluate external factors influencing the specific study context. Its legal component has a direct influence on the application and performance of the methodology through policies and planning regulations, separating the environmental dimension into an independent methodological category. These micro-environmental factors, as one of the crucial drivers of the workflow and the standards used, were implemented throughout the project development phases in coordination with governmental entities and research

teams. When dealing with rigorous official municipal planning documentation, it is necessary to follow and validate all regulatory factors. Within the research, the PESTEL framework aided in managing and mitigating risks arising from political, economic or social spheres, while considering important technological effects on planning practices and administrative operations through the use of GIS and open geospatial data (Fig. 1).

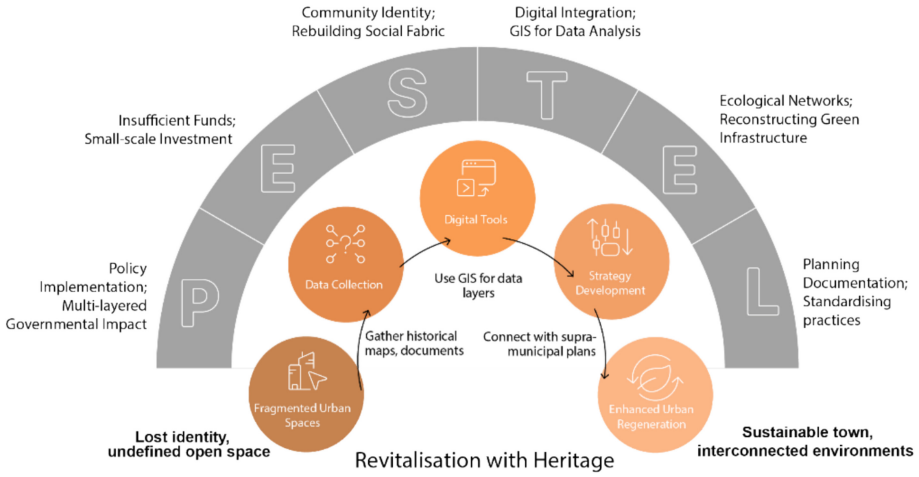

Fig. 1. Heritage-led urban regeneration methodological workflow analysed through the PESTEL strategic framework. External factors, such as planning documentation and governmental regulations, strongly influence the research methodology and its effective use.

Through the GIS digital environment, it is possible to construct dynamic two-dimensional representations that are fundamental for a clearer reading of the landscape in highly anthropised and urbanised territories, where historical features and the associated environmental values are often recognisable only in toponyms or in surviving fragments of landscape and built environment [21]. Comparison of historical continuities provided recognisable elements such as toponyms, agricultural and forest cultivation patterns, and the historical rural road network. These data formed an interoperable database, and together with contemporary information, codification, and vector shapefiles, can be visualised and analysed in parallel. Structured geographic metadata allowed implementation of strategies and actions within the Territorial Plan aimed at reconstructing the ecological and municipal green networks. On the one hand, strategies connect with supra-municipal plans and programs, and on the other, they locally enhance the historical, landscape, and environmental specificities of the territory of Desio. Subsequently, the research outcomes enabled the quantification and assessment of ecosystem benefits related to climate-environmental regulation, support and provisioning services, as well as socio-cultural benefits.

Systematising, georeferencing and vectorising historical cartography is the crucial part of preparatory and preliminary analyses for the formulation of urban planning actions. Vectorised historical maps were then linked with digitised land register data through attribute tables and cross-referenced with current information to create thematic

analysis. The synchronic reading of the historical-morphological evolution of the territory, together with analyses of agricultural green areas, aided the reconstruction of the peripheral landscape of Desio. Moreover, it can shape two multifunctional green belts of high ecosystem value for the whole territory [22]. These are interconnected through linear green infrastructures, reassembled along historical axes of the agricultural and forest landscape, redesigning the relationships between historical assets and rural landscapes.

Fragmentary elements of the 'green belt' extend from north to south on the west side of the Municipality of Desio. Many of these remaining areas are historical remains of the forests and orchards, and they were once part of a continuous green system characterising a larger territory. Today, a heterogeneous landscape is present with sparse farmsteads, often abandoned or demolished. In the recent past, regional and provincial interventions were developed for the revitalisation and reconnection of natural heritage with built points of interest, merging them into a unique infrastructure. In 2019, *Parco Branza Centrale* and *Parco Grugnotorto Villoresi* merged into GruBrìa. This local park, under Lombardy's Local Parks of Supra-Municipal Interest system, abbreviated PLIS, contributes to the Regional Ecological Network. It supports planning and ecological connectivity, while preserving natural, historical, and cultural heritage [23]. Within this merged park, Desio is the largest contributing municipality of a total of ten. It covers approximately 28.46% of the park (587 ha) [24]. Park GruBrìa benefits from a dense network of infrastructure: pedestrian paths, cycling routes, and rail connections, enhancing accessibility and connectivity with other regional parks (Fig. 2).

Additionally, the project focuses on redeveloping peripheral municipal areas with historical forests and cultural assets such as rural noble residences, churches, and farmsteads. According to the SIRBeC, an information system for cultural heritage in Lombardy, data obtained from the Geoportal of Lombardy (*Geoportale della Lombardia*) show that in the Municipality of Desio, there are forty-one buildings listed as cultural heritage[3]. Attention is given to the built heritage on the outskirts, from which nine belong to the 'religious and rural architecture' category and seven to the 'rural architecture' category. More specifically, SIRBeC reports seven completely or partially remaining farmsteads (i.e., typology: *cascina*) on the territory of Desio (Fig. 3).

3 Proposed Regeneration Actions and Results Implemented in the PGT

The territory of the Municipality of Desio was largely devoted to agricultural use until the eighteenth century, except for certain areas designated as deciduous woodlands, located mainly in the western sector toward the municipalities of Cesano Maderno, Binzago, and Bovisio. This situation remained largely unchanged until the end of the nineteenth century, when infrastructural and industrial development led to a progressive and continuous transformation of land use and the replacement of green areas with built-up spaces. This resulted in a gradual saturation of vacant and open spaces, occurring continuously from the 1950s through to the 2000s. These phenomena are clearly visible

[3] Source: https://www.lombardiabeniculturali.it/architetture/luoghi/560.763/ [last accessed 04/05/2026].

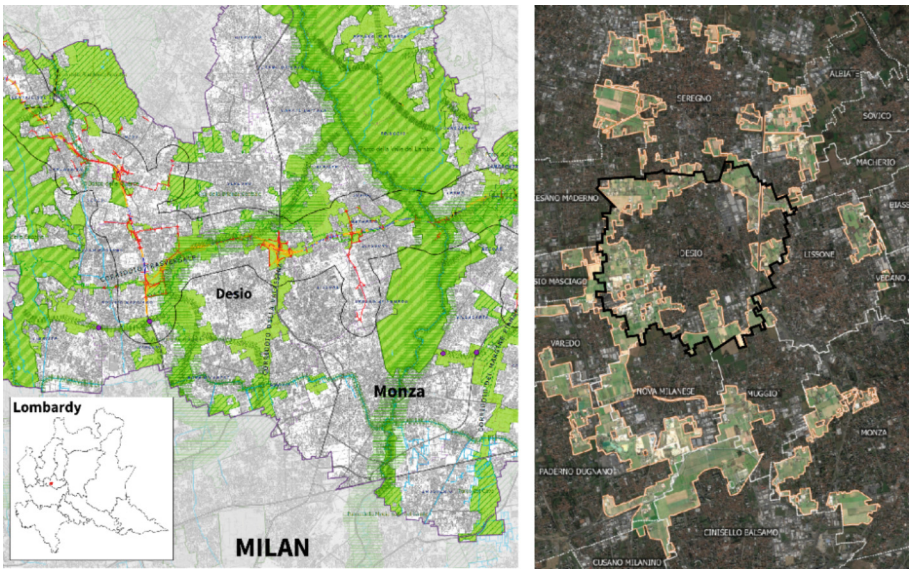

Fig. 2. On the left: Provincial Coordination Territorial Plan (PTCP). Table 6a: Landscape Protection and Enhancement Project (approved in 2013) (Source: https://www.provincia.mb.it/Temi/pianificazione-territoriale-e-parchi/pianificazione-territoriale/piano-territoriale-di-coordinamento/documentazione/ [last accessed 16/03/2026]. On the right: Map of PLIS GruBrìa Park showing Desio and adjacent municipalities (2021), Source: https://www.comune.desio.mb.it/it/page/grubria-fusione-dei-parchi-grugnotorto-villoresi-e- [last accessed 10/03/2026]).

through the overlay of georeferenced historical maps and development plans within the current system (Fig. 4). The results illustrate both directions and modes of land use consumption and landscape transformation, and they can identify the residual traces and enduring elements of buildings, green spaces, and historical infrastructures that have characterised this territory for centuries.

From the previously illustrated georeferenced analyses of Desio's historical territorial development, combined with the survey of the original vocations of the lands, it has been possible to reconstruct a cognitive framework of the historical agricultural and forestry functions of the Desio territory. Today, its readability is highly compromised by the intense urbanisation of recent decades, and it can impact the creation of a suitable framework for its investigation. Based on this multilayered territory reading, the Municipal Ecological Network (*Rete Ecologica Comunale* - REC) was constructed. Its structure originates in a landscape, recreational, and ecosystemic interpretation involving an informative reading of historical maps. Conducted analyses have highlighted that many of the agricultural areas recorded in the Lombardo-Veneto cadastre exhibited a forestry function linked to wood production, an element now perceptible only in restricted portions of the Desio territory and in its toponymy.

The strategies and actions underlying the implementation of the REC have primarily adopted the guidelines contained in the superordinate plans, the Regional Ecological Network (RER) and the Provincial Ecological Network (REP). On the other hand, when

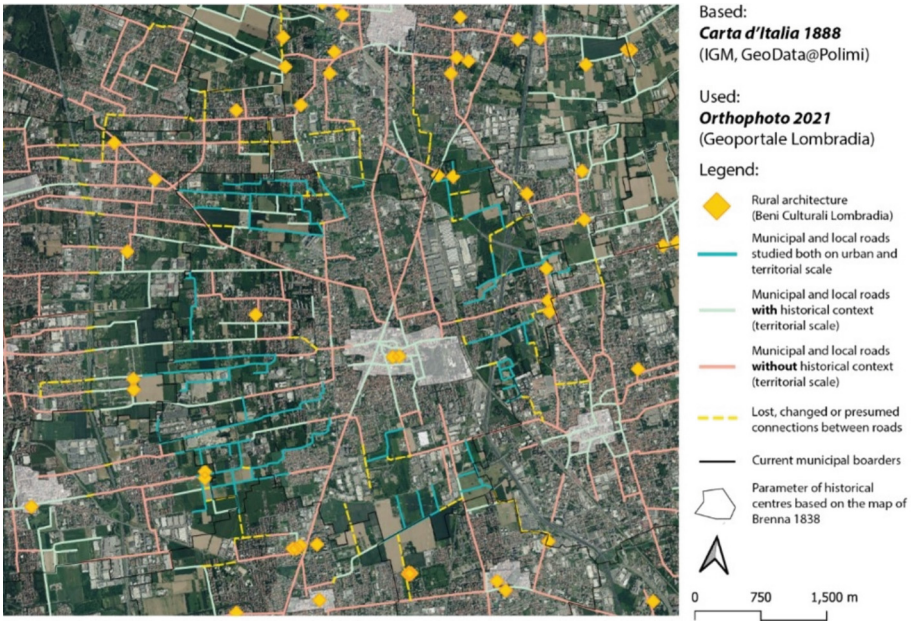

Fig. 3. Representation of the local rural roads and farmsteads (*cascine*) in the green belt around the Municipality of Desio. Road network was mapped using as a reference *Carta d'Italia, Istituto Geografico Militare* (1888), on the contemporary street network from *Geoportale della Lombardia* (2021).

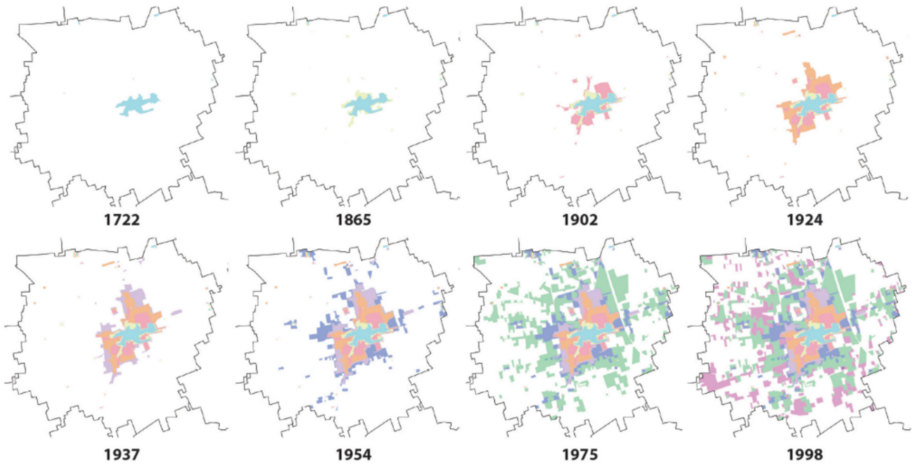

Fig. 4. Historical evolution of built areas from 1721 to 1998 of the Municipality of Desio. Elaborations are based on the historical cadaster (*Archivio di Stato di Milano*), topographic maps (*Istituto Geografico Militare*) and satellite imagery (*Geoportale della Lombardia*).

implementing them at a large scale, it is necessary to highlight existing discontinuities and thus identify possible local integration, implementation and final interventions. This involves a system composed of areal elements interconnected by linear elements, such as tree rows or cycle-pedestrian paths that have both ecological and environmental value. The green structure is articulated through two main systems. The outermost, so-called 'woods belt' consists of the existing and planned wooded areas, which develop mainly in the north-western portion of the municipal territory, in close connection with the Provincial Ecological Network. The implementation of this system pursues the dual objective of reconstructing the historical wooded landscape that characterised this part of the territory and enhancing the landscape, also in terms of providing ecosystem services.

On the contrary, the innermost ring is characterised by the so-called 'parks belt' that already partly surrounds and permeates the consolidated urban fabric. The implementation of this landscape-environmental system also occurs through environmental compensation arising from municipal territory transformation projects, particularly those affecting the Transformation Areas identified by the General Variant of the PGT and those related to new road infrastructures. Within this framework, the criticalities generated by the construction of the Lombard Pedemontana Highway in the north of the municipality of Desio have been treated as opportunities, and mitigation and environmental compensation projects have been identified as a strategic node of the Municipal Ecological Network (Fig. 5).

To better specify the composition of the two green systems, the former, nominated as the 'woods belt', is formed by areas already wooded or partially wooded, located predominantly along the western boundary of the municipal territory. This system is implemented and integrated through adjacent publicly owned areas that can be transformed into new woodland via the execution of compensation projects. To this are added compensation actions linked to the construction of the Pedemontana Highway and the identification of new strategic agricultural or green areas, privately owned, adjacent to the wooded belts. The latter system, referred to as 'parks belt', is formed by existing parks within the urban fabric, to which municipally and privately owned areas are added, capable of ensuring continuity to the city's landscape-environmental system (Fig. 6). Within the belt, buildings considered to be the morphological and typological features that characterise the historic built landscape have been mapped. Urban planning and building regulations aimed at their preservation are currently being drafted for each historical period to which they belong.

4 Expected Long-Term Impacts

The large green infrastructure generated by the implementation of the REC, designed by recovering the historical landscape features and by integrating the existing rural road network[4], and in all, represents a multifunctional belt capable of significantly increasing

[4] *Strade vicinali* (or *vie agrarie*) are private agricultural roads of local interest outside built-up areas. While they remain private property according to the Streets Code, strade vicinali grant public access through '*servitù di passaggio*' (source: Cereseto, G. B. (1894). *Le strade vicinali*. Unione Tipografico-Editrice).

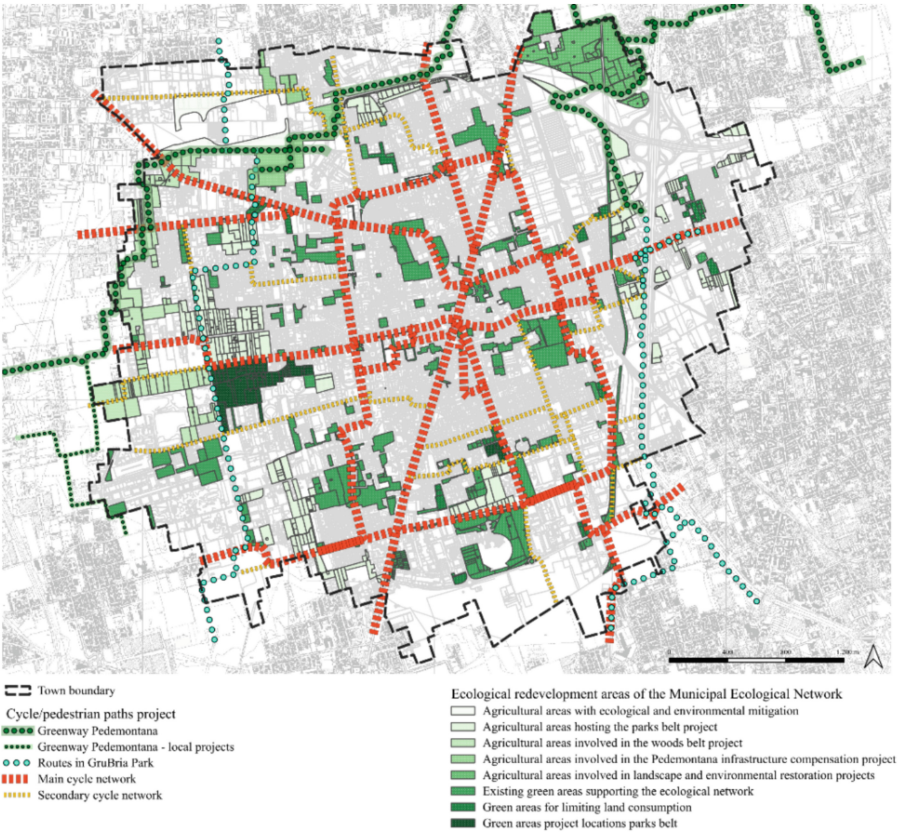

Fig. 5. Overlay of the Municipal Ecological Network (REC) with project cycle lanes identified by the General Variant of the PGT of Desio (edited by the authors).

the provision of ecosystem services[5] [25] of a regulating and socio-cultural nature for the territory and population of Desio. The application of the green systems of the new REC allows for a considerable increase in the ecological-environmental value of the underdeveloped free soils in the municipal territory, particularly by enhancing carbon storage capacity in the context of climate change mitigation. The new areas constituting the REC, located along the edges of the urban fabric, act as systems capable of absorbing, filtering, and managing water naturally and efficiently [26].

The transformation of free soils, which today are mostly areas deprived of vegetation, into forested systems or parks and gardens, with substantial tree and shrub cover, enables the delay of meteoric water infiltration. This is done through the interception by the foliage of tree canopies and an increase in the soil's absorption capacity, also

[5] According to the definition provided by the Millennium Ecosystem Assessment, ecosystem services provide multiple benefits, including regulating services (such as climate, water, and natural hazard regulation) and cultural services (such as recreational, aesthetic, and identity-related).

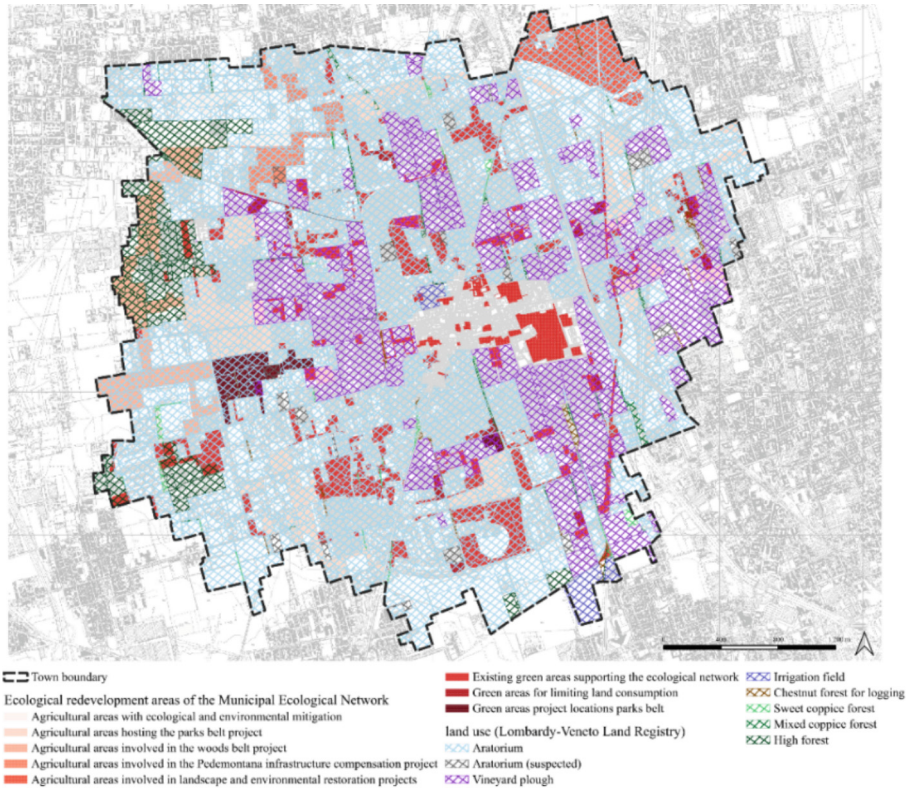

Fig. 6. Overlay of the Municipal Ecological Network (REC) with land use data extracted from the Lombardo-Veneto land cadastre (edited by the authors).

by constructing so-called natural systems for the infiltration and retention of meteoric waters (such as infiltration bio basins). These actions, measurable upon submission of the individual compensation projects, concern the litres of water intercepted by individual tree-shrub elements and the quantity of water infiltrated and retained by the bio basins. Furthermore, the creation of new forested surfaces, urban parks and gardens in the REC areas increases, both in number and extent, the surfaces capable of providing this ecosystem service of regulating soil temperature [27].

In total, the measures envisaged by the PGT variant provide for the enhancement, in ecosystem terms, of 1,538,000 m^2 of currently undeveloped free areas. This can be done by executing interventions of urban afforestation and reforestation, pursued through specific objectives and measurable targets directly correlated with the emerged vulnerabilities. These data are yet to be turned into actionable tactics for a long-term transformation of the industry-fragmented town of Desio. Building a resilient, eco-sustainable and green environment using the existing resources of the territory, not only targets the international and national recommendations and regulations but sets up the basis for future urban developments and positive impacts (Table 1).

Table 1. Summary of targets and performance for the implementation and enhancement of the Municipal Ecological Network[6] (edited by the authors).

Target	Performance indicators
Improvement of the ecological and environmental quality of the municipal area	- Potential green areas that contribute to improving ecological and environmental quality through urban afforestation and reforestation: 1,538,000 square metres, of which 409,200 square metres contribute to the 'forest belt', 912,996 square metres contribute to the 'park belt', 164,300 square metres of Pedemontana compensation areas and 96,816 square metres for interventions related to reducing land consumption - Existing green areas supporting the ecological network: 1,113,591 square meters - Increase: 142%
Preservation and reconstruction of the historical and naturalistic elements of the landscape	100% new wooded areas compared to areas historically designated as woodland (1,358,200 square metres of wooded areas of various types, cf. Extra-urban land use in georeferenced historical maps from the Lombardo-Veneto Land Registry)
Improving citizens' well-being in the use of green areas	- 3.5 km of existing municipal cycle and footpaths and 7.4 km of planned municipal cycle and footpaths crossing the compensation and support elements of the Municipal Ecological Network (REC) - Considering also the external cycle connections (existing and planned) to the REC, approximately 80% of the areas belonging to the parks and forest belt will be accessible
Improvement of the system for enjoying the landscape/environment/cultural heritage	100% of historical and cultural assets connected by the municipal cycle network (max. Distance 100 m)

The expected results of environmental performance remain contingent upon the implementation of the Plan, which is currently being drafted. After its approval, it will be possible to carry out the projects set out in the urban planning framework and begin a phase of thorough monitoring of the landscape and environmental actions. However, a further constraint of their implementation in the PGT is linked to the construction of grey public infrastructure, and urban and building transformation of territorial areas. They are currently open spaces on which land is being planned by private developers.

[6] The presented calculations were carried out using information from national and regional databases, with the shapefiles compiled in accordance with the rules of the Lombardy Region.

This approach to urban planning, which focuses on improving the environmental quality of free space, also, through its historical interpretation, remains linked to actions involving significant human pressure, such as new infrastructure (Pedemontana) and land consumption (areas of transformation).

5 Conclusions

The study has shown that integrating systematic studies of historical cartography and multitemporal territorial data within the PGT can effectively reconnect fragmented historical, environmental, and landscape values in a highly urbanised context, such as the Municipality of Desio. Collaboration between municipal planning offices and university research through GIS technology allowed the proposed methodology to transform heritage-informed analyses into concrete planning actions, following the regulations and frameworks on various governmental levels. One of these straightforward applications was the implementation and design of the Municipal Ecological Network and the definition of multifunctional green belts. The outcomes obtained from layering georeferenced historical cadastres, satellite imageries and contemporary data demonstrated that rural infrastructures and cultural assets can be reframed as active drivers of urban regeneration. They enhance ecosystem services and provide context-informed ecological continuities binding the whole territory.

Heritage-led regeneration can be operationalised within legal means from a planning perspective, like in the case of PGT. It allows various stakeholders to align local requirements with wider frameworks such as the REP, RER, the European frameworks for cultural heritage, and the SDG objectives. With tangible and intangible relationships between elements of the city, researchers were searching for true meanings and connections to create systems that rely on the gathered and processed historical knowledge. The linkage was obtained through toponyms, local present and past use, private and public property management, and reuse of assets already available within the territory to create a sustainable system for regeneration, development and movement. Remaining forested areas, *cascine*, *strade vicinali*, and tree rows are all part of the present heterogeneous systems, which were connected through this collaboration. These are all quantifiable benefits in the form of reduced land consumption, revived both inner and outer green belts, and improved conditions for socio-cultural services at multiple scales.

Nevertheless, the research indicates persistent challenges reflected in data gaps, the necessity of trained GIS municipality staff, constant data gathering, management and maintenance of the areas, both in person and through the geospatial databases. If these requirements are met, ecosystem regeneration can thrive in the long term, being beneficial for built and natural heritage, and the overall well-being of the citizens. The future advancements of this research will track and evaluate the applied methods in the field. There will be a need to refine the indicators for monitoring green services and infrastructure and prepare the database to receive updated information for processing.

The Desio case study confirms the efforts to use a GIS-supported historically informed planning tool as a scalable instrument to foster urban regeneration that can be environmentally, socially, economically and culturally inclusive. This specific, practically applied case study demonstrated the implementation of academic research into

governmental planning procedures. It also shows the replicability and feasibility of the methodology created using legal, socio-economic, and environmental regulations through a structured spatial-temporal GIS database. The approaches used in the research demonstrate the need to follow the documentational rigidity, while also integrating social, cultural, and environmental intangible values, which should be recognised and valorised before any planning actions.

Several challenges need to be addressed in the future. In most medium to small-sized Italian municipalities, there is a lack of standardised GIS practices and a permanent expert. The consequence is a partial database, lost or even unnecessarily repeated data. Additionally, the historical material is held in various public, private and religious holdings, mostly as part of larger collections. As such, they need to be collected, digitised, systematised, and handed to the municipal administrations for their use. There is an increasing risk that geospatial data will remain unmaintained and static for years to come, without constant information input. The solution would be a constant research urge of the public administration, investment in expert training, often collaborations with external study groups, and open, interoperable data formats. The real barrier emerges in replicating this methodology to other Italian regions, where municipal planning documents have a different form, or to other European cities with diverse planning models.

Acknowledgments. This research paper derives from the collaboration of the Municipality of Desio, represented by the Technical Area, headed by Arch. Floriana Petracco and the Politecnico di Milano to produce the variant of Municipal Territorial Plan (PGT). The university acted as a scientific partner, involving professors and researchers from two departments, the Department of Architecture, Built Environment, and Construction Engineering (DABC) and the Department of Architecture and Urban Studies (DAStU). We thank the scientific directors: Elena Mussinelli, Andrea Tartaglia (UniCT), Marco Scaioni, and Daniela Oreni, and the whole research team that worked and coordinated this project.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chen, D.: Examining urban design strategies in historic districts: a systematic literature review. Proc. Inst. Civ. Eng.-Urban Des. Plan. **177**, 232–239 (2024). https://doi.org/10.1680/jurdp.24.00046
2. Pellegrini, P., Micelli, E.: Paradoxes of the Italian historic centres between underutilisation and planning policies for sustainability. Sustainability **11**, 2614 (2019). https://doi.org/10.3390/su11092614
3. van Mil, Y., Hein, C., Baptist, V.: The potential of geospatial technologies and open data in planning history. Plan. Perspect., 1–18 (2025). https://doi.org/10.1080/02665433.2025.2561881
4. Jovanovic, D., Oreni, D.: Sifting through historical maps. methodology for the implementation of cartographic content in urban planning: case study of desio. In: Presented at the International Conference on Computational Science and Its Applications (2022). https://doi.org/10.1007/978-3-031-10545-6_12

5. Fanzini, D., Tartaglia, A., Riva, R.: Project challenges: sustainable development and urban resilience. Maggioli Editore (2019)
6. United Nations: Transforming our world: the 2030 Agenda for Sustainable Development. A/RES/70/1 (2015). https://docs.un.org/en/A/RES/70/1. Accessed 12 Feb 2026
7. Mille/r, E.J.: Integrated urban modeling: past, present, and future. J. Transp. Land Use. 11, 387–399 (2018). https://doi.org/10.5198/jtlu.2018.1273
8. European Commission: Urban Agenda for the EU Multi-Level Governance in Action. European Union Brussels, Belgium (2019). https://ec.europa.eu/regional_policy/sources/brochure/urban_agenda_eu_en.pdf. Accessed 21 Apr 2026
9. Ángeles Huete Garcia, M., Merinero Rodríguez, R., Muñoz Moreno, R.: Urban regeneration policy from the integrated urban development model in the European Union: an analytical approach based on the study of Spanish cities. Local Gov. Stud. **42**, 267–286 (2016). https://doi.org/10.1080/03003930.2015.1110520
10. Acierno, A.: Integrated urban development and culture-led regeneration in the EU. Territ. Della Ric. Su Insediamenti E Ambiente Riv. Internazionale Cult. Urban. 7–14 (2017). https://doi.org/10.6092/2281-4574/5324
11. Nocca, F.: The role of cultural heritage in sustainable development: multidimensional indicators as decision-making tool. Sustainability. **9**, 1882 (2017). https://doi.org/10.3390/su9101882
12. European Commission: Directorate-General for Education, Youth, Sport and Culture: European Framework for Action on Cultural Heritage. Luxembourg: Publications Office of the European Union (2019). https://data.europa.eu/doi/https://doi.org/10.2766/949707
13. Boeri, A., Longo, D., Gianfrate, V., Roversi, R.: Cultural heritage-led initiatives for urban regeneration. Pilot implementation actions in Bologna public spaces. In: DIVERSE city, pp. 463–472. CNAPPC (2019)
14. Baraldi, S.B., Salone, C.: Building on decay: Urban regeneration and social entrepreneurship in Italy through culture and the arts. Eur. Plan. Stud. **30**, 2102–2121 (2022). https://doi.org/10.1080/09654313.2022.2030675
15. Rozmus, S., Kiedrowicz, M.: Maintenance of chronology of facts in the historical geographic information system (HGIS). In: 26th Geographic Information Systems Conference and Exhibition "GIS ODYSSEY 2019", pp. 256–265. Croatian Information Technology Society – GIS Forum, Zagreb, Croatia, Slit, Croatia (2019)
16. Zerbi, A.: Gli HGIS catastali strumenti di rappresentazione aumentata del paesaggio urbano. Il caso di Parma fra XVIII e XX secolo/Cadastral HGIS as tools for augmented representation of the urban landscape. The case of Parma between the XVIII and XX centuries. In: Delli Aspetti de Paesi. Vecchi e nuovi Media per l'immagine del Paesaggio. Tomo II–Rappresentazione, memoria, conservazione, pp. 107–116. CIRICE (2016)
17. Knowles, A.K., Hillier, A. (eds.): Placing History: How Maps, Spatial Data, and GIS are Changing Historical Scholarship GIS and History. ESRI Inc, Redlands, California (2008)
18. Brigante, R., Radicioni, F.: Georeferencing of historical maps: GIS technology for urban analysis. Geogr. Tech. **9**, 10–19 (2014)
19. Jovanovic, D., Oreni, D.: The methodology to systematise, present and use historical cartography: potentials and limits to analyse and enhance widespread historical centres in Northern Italy. In: The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, pp. 339–346. Copernicus Publications, Beijing, China (2021). https://doi.org/10.5194/isprs-archives-XLVI-M-1-2021-339-2021
20. Jovanovic, D., Oreni, D.: Reevaluating historical road connections through historical cartography. In: Presented at the VL 2024 International Conference on Visualizing Landscape, pp. 738–753. PVBLICA, Alghero (2024)
21. Lelo, K.: A GIS approach to urban history: Rome in the 18th century. ISPRS Int. J. Geo-Inf. **3**, 1293–1316 (2014). https://doi.org/10.3390/ijgi3041293

22. Cittadino, A., Garnero, G., Guerreschi, P., Inzerillo, M.: Strumenti e metodi per la rappresentazione e l'interpretazione delle dinamiche territoriali/Tools and methods for the representation and interpretation of territorial dynamics. Urban. Doss., 121–129 (2023)
23. Regione Lombardia: Rete Ecologica Regionale e programmazione territoriale degli enti locali / Regional Ecological Network and spatial planning of local authorities. (2010). https://www.regione.lombardia.it/content/dam/rl/canali-tematici-servizi/06-ambiente-e-territorio/05-biodiversita/rete-ecologica-regionale/allegati/rete-ecologica-regionale.pdf. Accessed 02 May 2026
24. Parco GruBrìa: Caratteristiche e dati, https://www.grubria.it/cgrubria/zf/index.php/servizi-aggiuntivi/index/index/idtesto/6. Accessed 18 Feb 2026
25. MEA Millennium Ecosystem Assessment Board: Ecosystems and Human Well-Being. Island press Washington, DC (2005)
26. Antonini, E., Tucci, F.: Architettura, Citta' e Territorio Verso la Green Economy. La costruzione di un Manifesto della Green Economy per l'Architettura e la Città del Futuro / Architecture, City and Territory Towards a Green Economy. Building a Manifesto of the Green Economy for the Architecture and the City of the Future. Edizioni Ambiente (2017)
27. Mussinelli, E., Tartaglia, A., Bisogni, G.L., Malcevschi, S.: Il ruolo delle nature-based solutions nel progetto architettonico e urbano / The role of nature-based solutions in architectural and urban design. Techne. **15**, 116–123 (2018). https://doi.org/10.13128/Techne-22112

Co-creating Circular Tourism Action Plans in Remote Heritage Sites: Evidence from Three Scottish Castles

Antonia Gravagnuolo[1(✉)] and Patrizia Riganti[2]

[1] Institute of Heritage Science, CNR ISPC, Naples, Italy
antonia.gravagnuolo@cnr.it
[2] School of Interdisciplinary Studies, University of Glasgow, Dumfries, UK

Abstract. The transition from conventional tourism models toward more regenerative and place-based approaches has increased interest in the relationship between circular tourism and responsible tourism, particularly in remote heritage sites that, despite their significant cultural assets, often remain outside major tourism flows and struggle to convert heritage value into broader territorial benefits. In this perspective, circular tourism is understood not only in terms of resource efficiency, but as a broader approach linking heritage conservation, local economies, environmental stewardship, community well-being, and long-term resilience.

Against this background, the paper analyses three Scottish castles—Doune, Caerlaverock, and Blackness—as case studies for testing a participatory and place-sensitive approach to the co-creation of Circular Tourism Action Plans. Methodologically, the study adopts a collaborative multicriteria evaluation approach combining documentary analysis, site-based evidence, and stakeholder engagement to identify challenges, prioritise intervention fields, and support locally grounded strategies.

The findings show that a shared evaluative framework can be applied across different heritage contexts while generating differentiated and site-specific priorities. The process proved effective not only in defining action fields, but also in supporting co-evaluation and collective learning among local actors. The paper argues that Circular Tourism Action Plans can emerge as transferable yet context-sensitive outputs, adaptable to other heritage destinations facing similar challenges of conservation, accessibility, local benefit, and resilience.

Keywords: Circular cultural tourism · Responsible tourism · Participatory evaluation · Peripheral heritage destinations · Heritage regeneration

1 Introduction

Heritage destinations are increasingly expected to reconcile conservation, visitor management, local benefit and climate responsibility within the same strategic framework. This challenge is particularly urgent in peripheral and less-known areas, where tourism can produce uneven territorial effects and where heritage management requires tools capable of combining place sensitivity, community involvement and long-term resilience.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 471–487, 2027.
https://doi.org/10.1007/978-3-032-30518-3_30

The growing convergence between heritage-led development, climate-sensitive territorial planning, and circular economy thinking has opened new avenues for reinterpreting cultural tourism beyond conventional growth-oriented models [1–3]. In this evolving field, circular cultural tourism offers a particularly fertile conceptual and operational lens. It shifts attention away from linear tourism dynamics based on attraction, consumption, and leakage, and toward regenerative and place-based processes in which heritage, landscape, local economies, environmental resources, and community capacities are mutually reinforcing [4–12]. Within this perspective, tourism is no longer understood as an external force acting upon a heritage site, but as a potential catalyst for more balanced territorial development, capable of retaining value locally, reducing environmental pressures, strengthening identity, and fostering longer-term resilience [4]. The European-funded Horizon 2020 research and innovation project Be.CULTOUR (Beyond cultural tourism) explicitly articulated this ambition as a move beyond conventional cultural tourism toward a more human-centred and circular tourism economy grounded in heritage and landscape values [13].

This shift becomes clearer when circular tourism is contrasted with conventional linear tourism. Although tourism is usually classified as a service sector, it relies on large flows of energy, water, food, mobility, land, and cultural and natural resources, and may therefore reproduce a "take-make-dispose" logic with significant environmental, social, and cultural externalities [14]. Earlier circular tourism scholarship has argued that tourism needs to move beyond extractive and volatile development models and toward more sustainable and regenerative strategies capable of generating multidimensional benefits while reducing environmental and social costs [2]. Thus, circular tourism is not only about mitigating impacts, but about rethinking tourism development and management in ways that retain value locally, reduce pressure on resources, and strengthen the long-term resilience of destinations and communities [15]. Rather than limiting circularity to resource efficiency alone, the Be.CULTOUR framework is built around three interrelated domains: environmental regeneration, human-centred approach, and cultural experience. It therefore includes not only the reduction of waste, emissions, and resource consumption, but also human capital enhancement, community wellbeing, fairness and inclusion, local entrepreneurship, cooperation and collaboration, cultural heritage appreciation, place authenticity, landscape integrity, and transformative travel experience. Circularity is thus linked to the regeneration of territorial capacities and to the enhancement of less-known, underused, and peripheral cultural and natural resources.

This conceptual orientation is especially significant in peripheral, rural, or lesser-known heritage contexts. In such areas, tourism often occupies an ambiguous position. On the one hand, it may support place recognition, local entrepreneurship, and heritage valorisation [16–20]. On the other hand, it may exacerbate infrastructural fragility, generate uneven benefits, produce conservation pressures, and deepen the gap between visitors' experience of place and residents' lived realities [16, 21–23]. A circular approach suggests that tourism should be assessed not only through visitor flows or economic performance, but through its capacity to support conservation, low-carbon accessibility, social inclusion, local value recirculation, community wellbeing, and adaptive territorial governance [15].

This orientation resonates strongly with the Responsible Tourism Framework developed by Historic Environment Scotland (HES) [24]. The framework defines responsible tourism through a set of interrelated principles - environmental stewardship, sustainable resource use, community involvement, maximising local economic benefits, and equality, diversity, and inclusion - and situates tourism within a wider agenda of climate action, net zero transition, wellbeing economy ambitions, accessibility, community wealth building, and place-based development. It also makes a crucial methodological point: progress in responsible tourism cannot be understood through a one-size-fits-all logic. Instead, it requires differentiated baselines, site-sensitive indicators, and evaluation models capable of reflecting local conditions, capacities, and aspirations. HES further argues that traditional indicators such as visitor numbers and commercial revenue capture only part of the picture, and that broader qualitative and multidimensional measures - such as authenticity of place, community wellbeing, satisfaction, and quality of experience - must also be considered. This makes the framework especially relevant for heritage-led territorial regeneration, where tourism must be aligned with conservation, local wellbeing, accessibility, and long-term place-based resilience.

In this paper, circular tourism is not used as a synonym for responsible tourism or regenerative tourism, but as a distinct analytical lens. Responsible tourism is understood primarily as a normative and policy-oriented framework, centred on principles such as environmental stewardship, community benefit, inclusion, and sustainable resource use. Circular tourism, by contrast, is used here as a multidimensional evaluative and operational framework that focuses on how value, resources, benefits, and territorial capacities can be retained, recirculated, and strengthened over time within heritage destinations. In this sense, circularity adds a specific emphasis on interrelations between heritage conservation, local economies, accessibility, environmental pressures, and governance conditions, allowing tourism to be analysed not only in terms of impacts, but also in terms of systemic territorial regeneration. Regenerative and place-based approaches are closely related to this perspective, but they remain broader and less directly operational in the context of this study. The concept of circular tourism is therefore preferred because it provides a more explicit framework for structuring the evaluation process and for translating it into locally grounded action priorities.

This paper addresses these issues through a collaborative experimentation undertaken in Scotland by the University of Glasgow, the National Research Council of Italy (CNR-ISPC), and Historic Environment Scotland. The study applied circular cultural tourism criteria and indicators developed within Be.CULTOUR to three heritage case studies (Doune Castle, Caerlaverock Castle, and Blackness Castle) and used them to support the co-creation of Local Action Plans through participatory processes involving students, experts, stakeholders, and institutional actors. The research was informed simultaneously by the Be.CULTOUR circular cultural tourism approach and by the HES Responsible Tourism Framework, thereby allowing the study to connect heritage-based circularity with responsible tourism objectives such as low-carbon mobility, accessibility, local economic benefits, community engagement, and place-sensitive governance.

Despite the growing policy and academic attention devoted to regenerative, circular and responsible tourism, there is still a limited understanding of how these principles can be operationalised in concrete heritage settings, especially in peripheral destinations

where conservation needs, visitor pressures, local benefit and climate-related concerns must be addressed simultaneously. This gap is particularly relevant at a time when heritage managers and public authorities are increasingly required to align tourism development with broader goals of territorial resilience, community involvement and long-term stewardship. Against this background, the paper investigates how a participatory and multicriteria evaluation approach can support the co-creation of Local Action Plans in three Scottish heritage sites. More specifically, it addresses the following research questions: (1) How can circular and responsible tourism principles be translated into a place-sensitive evaluative process for heritage sites? (2) How can a shared methodological framework generate differentiated priorities across distinct local contexts? (3) To what extent can such an approach be considered transferable to other heritage destinations?

The originality of the paper lies not in proposing a wholly new circular tourism framework, but in making the Be.CULTOUR circular cultural tourism perspective and the Historic Environment Scotland Responsible Tourism Framework operational through their application in three remote Scottish heritage sites. More specifically, the study shows how these theoretical frameworks can be translated into a shared evaluative process that combines indicators, contextual evidence, and stakeholder engagement in order to generate place-sensitive Local Action Plans. The paper's distinctive contribution therefore lies in the experimental application of these frameworks in real heritage contexts and in demonstrating their usefulness for producing context-sensitive and potentially transferable action priorities in peripheral destinations.

1.1 Study Areas

The research focused on three Scottish heritage sites managed by Historic Environment Scotland: Doune Castle, Caerlaverock Castle, and Blackness Castle. The three cases were selected because they allowed a common methodological framework to be tested across markedly different territorial and tourism configurations.

Doune Castle, located in Stirlingshire, is one of Scotland's most recognisable medieval monuments. Its significance lies both in its historical and architectural importance and in its strong visibility within screen tourism circuits. Official HES documentation shows a marked growth in visitation over the last decade, with substantial increases associated with its role as a filming location for productions such as Monty Python and the Holy Grail, Game of Thrones, and Outlander. This success, however, has also intensified pressure on the landscape setting and on the site's carrying conditions. The Doune Castle Landscape Conservation Management Plan identifies deterioration linked to footfall and vehicle movement, including erosion of earthworks, the exposure of below-ground archaeology, parking problems, and a lack of coordinated landscape management responses.

Caerlaverock Castle presents a different territorial condition. While its iconic triangular form and moat make it a highly distinctive monument, its heritage significance is closely tied to a broader environmental and territorial setting composed of coastal landscapes, nature reserves, local cultural assets, nearby attractions, and slow-mobility opportunities. The relevance of this case lies in the fact that tourism performance cannot be meaningfully interpreted at the scale of the monument alone. Instead, it demands a destination-scale perspective in which the castle operates as part of a wider system of

access routes, landscape experiences, local economic opportunities, and regional collaboration. It is therefore particularly useful for analysing how circularity depends on the integration of heritage, landscape, mobility, and dispersed local experiences.

Blackness Castle adds a third and complementary condition. Here, the most pressing challenge is not simply site pressure or destination definition, but the interaction between heritage tourism, infrastructure, accessibility, local liveability, and the redistribution of economic benefits. Blackness combines strong scenic value, coastal location, historical depth, and links with screen tourism, yet it also faces serious limitations related to public transport, parking, interpretation, accessibility, and local value recirculation. As such, it is particularly useful for examining the infrastructural and socio-spatial preconditions of circular tourism.

Taken together, the three cases create a robust comparative framework. Doune foregrounds the problem of managing visibility and pressure; Caerlaverock foregrounds territorial integration and destination logic; Blackness foregrounds access, infrastructure, and community benefit. This diversity makes it possible to test whether a common analytical structure can remain methodologically coherent while yielding differentiated and context-sensitive priorities.

2 Methodology

The study adopted a participatory evaluation design combining documentary analysis, indicator-based assessment, contextual site reading, and stakeholder-oriented co-creation. The objective was not to produce a fixed ranking of the three sites, nor to apply a standardised sustainability index, but to develop an evaluative process capable of translating broad principles of circularity and responsibility into locally meaningful strategic orientations.

The first conceptual pillar of the research was the circular cultural tourism framework developed within Be.CULTOUR [13, 19]. In operational terms, this framework was translated into four interrelated dimensions (social and human capital, cultural capital, economic capital, and environmental capital) used as a common analytical structure across the three Scottish case studies. These dimensions made it possible to interpret tourism as a multidimensional process affecting not only heritage use and visitor experience, but also local entrepreneurship, accessibility, environmental quality, cooperation, wellbeing, and place authenticity.

The second conceptual pillar was the HES Responsible Tourism Framework, which provided the institutional and policy-oriented basis for reinterpreting circularity within the context of Scottish heritage management. Through its emphasis on environmental stewardship, sustainable resource use, community involvement, local economic benefits, and inclusion, the HES framework allowed the research to situate circular tourism within a broader agenda of climate responsibility, accessibility, community benefit, and place-based governance.

Beyond the indicator structure, the methodological design was also informed by the Be.CULTOUR protocol for human-centred innovation in sustainable and circular cultural tourism. In this paper, the co-creation of Local Action Plans was conceived as a structured, place-based process grounded in stakeholder engagement, collaborative interpretation of evidence, and the definition of locally relevant strategic priorities.

The overall co-design logic followed the early stages of the Be.CULTOUR methodology, particularly the exploration and Action Plans co-design phases. Accordingly, the participatory multicriteria evaluation was organised as a sequential process in which the Be.CULTOUR and HES frameworks were combined into a shared evaluative structure based on four dimensions—social and human capital, cultural capital, economic capital, and environmental capital—and applied to the three case studies through quantitative and qualitative indicators, together with documentary, contextual, and site-based evidence. This preliminary assessment was then discussed and refined through participatory co-evaluation workshops involving institutional actors, heritage managers, researchers, students, and local stakeholders, and finally synthesised into the identification of key issues, context-sensitive priorities, and the resulting Local Action Plans (Fig. 1).

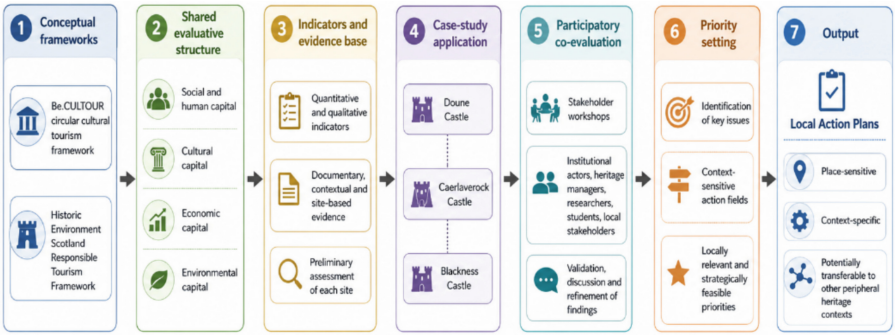

Fig. 1. Participatory workflow for co-creating Circular Tourism Action Plans.

The participatory process was organised through a series of case-based workshops designed as structured co-evaluation and co-creation sessions. Each workshop involved a mixed group of participants, including institutional actors, heritage managers, researchers, students, and local stakeholders with knowledge of the site and its wider territorial context. The workshops were based on the discussion of preliminary findings derived from documentary, contextual, and site-based analysis, and were designed to compare analytical evidence, local perceptions, and strategic expectations. Rather than aiming at consensus from the outset, the process allowed different viewpoints to emerge around issues such as accessibility, visitor pressure, local benefit, and conservation priorities. These exchanges were then used to refine the interpretation of each case and to identify the action fields that were considered both locally relevant and strategically feasible.

The three cases were analysed through the same evaluative architecture in order to test its replicability across different heritage contexts while allowing for place-specific outcomes (Table 1).

The evidence base of the study combined official published sources used to define the institutional, policy, and territorial context of the three sites, including the HES Responsible Tourism Framework, related policy and strategy documents, and official site-level publications; and site-specific contextual data used to operationalise the indicator framework, including information on visitor dynamics, accessibility, transport conditions,

Table 1. Methodological design of the comparative study.

Case study	Main territorial/heritage profile	Main sources and evidence	Participatory process	Output
Doune Castle	High-profile heritage attraction affected by visitor concentration and seasonal pressure	Documentary sources, site-based evidence, workshop discussions	Stakeholder workshop aimed at identifying critical pressures and action priorities	Local Action Plan focused on visitor management, conservation compatibility and seasonal redistribution
Caerlaverock Castle	Heritage site embedded in a wider landscape and destination context	Documentary sources, site-based evidence, workshop discussions	Stakeholder workshop aimed at connecting site management with destination-scale opportunities	Local Action Plan focused on destination integration, mobility, local experiences and territorial connectivity
Blackness Castle	Heritage site affected by accessibility issues, infrastructural constraints and community tensions	Documentary sources, site-based evidence, workshop discussions	Stakeholder workshop aimed at identifying barriers, local concerns and collaborative opportunities	Local Action Plan focused on accessibility, decentralisation of visitor spending and regional collaboration

site significance, conservation issues, interpretative provision, and local infrastructure (Table 2).

Also, the study relied on evidence generated within the research process through indicator application, contextual interpretation, and participatory co-creation. This evidence formed the empirical basis for the identification of action priorities in each case study. The Local Action Plans can also be interpreted as the output of a co-evaluative decision-support process. This is particularly relevant for evaluation studies, as it shows how participatory planning can be structured around indicators without reducing the process either to technocratic assessment or to unstructured consultation.

Table 2. Comparative synthesis of case-study problems and observed variables.

Case study	Main problems identified	Indicators / variables observed
Doune Castle	Visitor concentration, seasonality, landscape fragility, conservation-management tensions	Visitor pressure, seasonal patterns, access conditions, landscape degradation, local benefit potential, interpretative orientation
Caerlaverock Castle	Remoteness, fragmented visitor offer, weak integration between site and wider destination	Accessibility, connectivity, destination integration, local services, mobility opportunities, visitor experience diversification
Blackness Castle	Accessibility barriers, infrastructural deficits, weak local capture of spending, community tensions	Public transport conditions, accessibility, interpretative provision, local spending retention, regional collaboration potential

3 Results

3.1 Doune Castle: Managing Visibility, Pressure, and Landscape Fragility

The Doune Castle case demonstrated the usefulness of the integrated framework for analysing a heritage site characterised by strong visibility and increasing pressure. A first and decisive result concerns the identification of a mismatch between the site's cultural attractiveness and the fragility of its territorial support conditions. Official HES evidence shows substantial visitor growth over the last decade, associated not only with the monument's intrinsic significance but also with the impact of film and television exposure. This growth has increased the site's appeal and strategic relevance, but it has also amplified pressures on the wider landscape setting and on the site's capacity to absorb visitation without compromising heritage value.

This condition is strongly reinforced by the Doune Castle Landscape Conservation Management Plan, which identifies significant deterioration linked to rising footfall and increased vehicle movement [25]. The plan points to erosion affecting earthworks, exposure of below-ground archaeology, parking-related disturbances, and a piecemeal approach to landscape maintenance and presentation. It also stresses that the management of the site's broader landscape setting has often remained secondary to the conservation of the built fabric itself, despite the fact that landscape quality is fundamental both to heritage significance and to the visitor experience.

A second result concerns the value of the circular indicator framework in revealing the interdependence of issues that might otherwise be treated separately. At Doune, visitor concentration, seasonality, access patterns, interpretative orientation, and local benefit could not be meaningfully separated from landscape conservation and long-term site resilience. The analysis therefore highlighted that site resilience depends not only on conservation measures, but also on how tourism is distributed, interpreted, and territorially connected.

A third result lies in the action-oriented nature of the participatory process. The Doune case showed that the pathway toward a more circular tourism model could not be reduced to managing flows alone. Instead, it pointed toward a more articulated strategy involving de-seasonalisation, sustainable access, strengthened local economic linkages, more inclusive and diversified interpretation, and visitor management approaches compatible with heritage conservation.

The resulting action fields for Doune converged around five priorities. The first was visitor redistribution and de-seasonalisation, with the aim of reducing concentration peaks and creating a more balanced annual pattern of use. The second was sustainable mobility and access, particularly through a stronger emphasis on low-carbon visitor routes and access options. The third was community-based value creation, aimed at strengthening local economic recirculation rather than limiting the benefits of tourism to the site itself. The fourth was interpretative diversification, broadening the experiential offer beyond narrowly screen-tourism-oriented narratives. The fifth was conservation-compatible visitor management, recognising that any tourism strategy for Doune must remain aligned with the protection of landscape, archaeology, and long-term site resilience.

3.2 Caerlaverock Castle: From Site-Based Heritage to Shared Destination Logic

If Doune highlighted the challenge of managing visibility and pressure, Caerlaverock shifted the evaluative scale toward a wider territorial and destination logic. One of the most significant results of this case was the recognition that circular tourism at Caerlaverock could not be meaningfully framed at the level of the monument alone. The castle emerged instead as an anchor within a larger territorial constellation composed of coastal landscapes, natural assets, nearby attractions, cultural experiences, routes of access and mobility, nature-based experiences, and local service networks.

This change of scale is analytically important. It suggests that in some heritage contexts, circularity is less about improving the efficiency or sustainability of a single site than about constructing a more integrated destination ecology. In the Caerlaverock case, the relationship between monument, landscape, local services, nearby experiences, and slow-tourism opportunities became central to the evaluative process.

A second result concerns the reinterpretation of seasonality and relative remoteness not only as constraints, but as opportunities. Rather than approaching lower shoulder-season visitation simply as a deficit, the case pointed toward the possibility of developing more immersive, slower, and higher-quality tourism experiences, less dependent on volume growth and more aligned with place-sensitive development. This supports an important point in circular cultural tourism: peripheral heritage destinations need not imitate high-intensity visitor economies. Their strength may lie precisely in their capacity to offer different temporalities, richer territorial immersion, and a stronger balance between heritage, nature, and community.

A third result concerns the identification of structural barriers to this transition. These included access and mobility challenges, limited facilities, fragmented promotion, limited integration between the castle and surrounding attractions, and insufficient support for longer stays and broader visitor itineraries. The evaluative process showed that these

constraints reduce both visitor dwell time and the possibility of distributing benefits across the wider destination.

The action fields emerging from the Caerlaverock case therefore reflected a clearly territorial logic. They included destination identity and positioning, aimed at strengthening the site's uniqueness within a wider offer; sustainable access and wayfinding, especially through walking, cycling, and integrated access solutions; integrated visitor experience, designed to encourage longer stays and multi-site exploration; inclusive and digital accessibility, particularly through interpretative enhancement; community engagement and local sourcing, linking tourism more directly to local products, education, and capacity building; and environmental and circular infrastructure improvements, aimed at reinforcing responsible destination management.

The Caerlaverock case thus demonstrated that participatory evaluation can effectively transform the object of planning itself. Rather than focusing solely on a heritage asset, the process helped construct a shared destination perspective, in which site value is inseparable from territorial relationships.

3.3 Blackness Castle: Infrastructure, Accessibility, and the Redistribution of Value

The Blackness case brought another dimension sharply into focus: the infrastructural and socio-spatial conditions required for circular tourism to produce broad-based local benefits. The site combines a strong scenic coastal setting, historical depth, regional centrality, and screen-tourism potential, yet these strengths coexist with serious limitations in access, interpretation, local service provision, and community relations.

A first result concerns the coexistence of high territorial potential and severe infrastructural constraint. Blackness appears well placed, both geographically and symbolically, to serve as a node within a wider heritage and coastal tourism network. Yet the evaluative process showed that this potential is constrained by poor public transport, limited parking, weak digital presence, and accessibility barriers both to the site and within it. In this case, circularity could not be advanced meaningfully without addressing infrastructure.

A second result concerns the centrality of economic recirculation as an evaluative issue. The Blackness case made visible the problem of short, site-centred visits and limited local spending. Tourism, in this context, risked remaining extractive even at modest scale: visitors arrived, consumed the site, and left, while the surrounding area captured only a limited share of the benefits. This finding is important because it shows that circular tourism cannot be understood only through environmental and conservation indicators. It must also address the geography of spending, length of stay, and the distribution of economic opportunity.

A third result concerns the relationship between tourism and local liveability. At Blackness, tensions linked to transport and parking made it clear that tourism cannot be assumed to benefit local communities automatically. A responsible and circular approach requires that burdens and benefits be jointly assessed. Tourism that generates inconvenience, congestion, and limited local return may undermine its own legitimacy.

The action fields identified for Blackness therefore reflected an explicitly systemic transition agenda. These included transport and responsible access, with a focus on

public transport enhancement, shared mobility, and parking solutions; inclusive access and interpretation, combining physical accessibility improvements with digital interpretation; experience diversification and off-season programming, aimed at broadening the site's role and reducing narrow patterns of visitation; decentralisation of spending, through stronger local service linkages and longer-stay opportunities; regional collaboration, positioning Blackness within a wider network of nearby attractions; and green operational improvements, linking site management to circular economy and climate-responsible practices.

Compared with Doune and Caerlaverock, Blackness foregrounded most clearly the fact that circular tourism depends on the repair and strengthening of relationships: between the site and the community, between visitors and the regional economy, and between tourism attractiveness and local infrastructural capacity.

3.4 Cross-Case Synthesis

Taken together, the three case studies confirm that the integrated methodological framework is both replicable and non-standardising. The same analytical structure—based on circular cultural tourism dimensions and responsible tourism principles—was maintained across the three sites, yet the priorities that emerged were clearly differentiated (Table 3).

At Doune, the dominant issue was how to manage strong visibility and visitor pressure while safeguarding landscape integrity and strengthening local recirculation.

At Caerlaverock, the dominant issue was how to develop a shared destination model integrating monument, landscape, mobility, and quality territorial experiences.

At Blackness, the dominant issue was how to address infrastructural, accessibility, and community-related constraints in order to decentralise value and create broader territorial benefits.

Workshop discussions broadly converged on pressure and conservation issues at Doune, destination integration at Caerlaverock, and accessibility and local benefit at Blackness, while also revealing different expectations regarding the scale and feasibility of future interventions.

The methodological framework was replicated across different heritage sites while generating differentiated and context-sensitive priorities, thereby supporting both comparability and local specificity. Cross-case comparison highlights both the replicability and the adaptability of the proposed evaluative process. While all three case studies shared a common methodological structure, the resulting priorities differed according to their specific combinations of visitor pressure, environmental sensitivity, infrastructural conditions, local economic integration and governance dynamics. This confirms that the value of the approach lies not in producing standardised outputs, but in enabling structured comparison while preserving place-based differentiation.

This variation is methodologically important. It confirms that the usefulness of circular tourism indicators does not lie in producing uniform rankings or generic sustainability labels. Their real value lies in structuring a place-sensitive and participatory evaluative process capable of revealing differentiated pathways of transition while maintaining a coherent evaluative structure across diverse heritage contexts.

Table 3. Synthesis of workshop results.

Site	Main challenges identified	Action priorities	Predominant scale of intervention	Evaluative insights
Doune Castle	Visitor pressure, seasonality, landscape fragility, conservation-management tensions	Visitor redistribution, de-seasonalisation, sustainable access, conservation-compatible management	Site and surrounding landscape	The main challenge lies in balancing attractiveness with resilience and conservation needs
Caerlaverock Castle	Remoteness, fragmented visitor offer, weak integration between site and wider destination	Destination integration, improved connectivity, multi-site experience design, local value creation	Territorial / destination scale	The key shift is from site-based management to a broader destination perspective
Blackness Castle	Accessibility barriers, infrastructural deficits, weak local capture of spending, community tensions	Accessibility improvements, transport solutions, decentralised spending, stronger regional collaboration	Site and regional network	The main issue is linking heritage tourism to local liveability and shared territorial benefits

4 Discussion

The three Scottish cases offer several broader insights relevant to heritage planning, environmental appraisal, and multicriteria evaluation. The originality of this experimentation lies not only in the application of indicator frameworks, but in the way indicators were mobilised within a participatory and human-centred co-creation process. In this study, indicators did not operate merely as technical measurement devices; rather, they served as structuring tools for dialogue, interpretation, and action prioritisation. The Local Action Plans can therefore be understood as the outcome of a shared evaluative process aimed at aligning circularity, responsibility, and place-based tourism development.

Moreover, the research confirms the analytical value of integrating circular cultural tourism with responsible tourism. The Be.CULTOUR framework offered a strong multidimensional structure through the four capital dimensions, while the HES Responsible Tourism Framework provided a policy-oriented articulation of those dimensions through community involvement, low-carbon access, local economic benefits, accessibility, and environmental stewardship. This integration prevented circularity from being reduced to

material circularity alone. Instead, it enabled the study to approach circular tourism as a broader condition involving social inclusion, interpretative quality, local recirculation, and territorial collaboration.

A further theoretical implication concerns the kind of circularity adopted in tourism planning. More radical and culture-led approaches are needed if circularity is to support strong sustainability, community wellbeing, place-making, and regenerative territorial development. From this perspective, the present study aligns with an interpretation of circular cultural tourism that treats tourism not as an end in itself, but as an instrument for strengthening heritage conservation, local entrepreneurship, ecological responsibility, and the social and cultural resilience of destinations. The paper therefore adopts a regenerative and culture-led interpretation of circular tourism, rather than a narrowly technical or efficiency-based one.

This evidence formed the empirical basis for the identification of action priorities in each case study. Through co-creation, analytical evidence was translated into strategic fields of action that reflected both expert understanding and local judgement. The process moved from assessment to prioritisation, and from prioritisation to action orientation. This is particularly relevant for evaluation studies, because it shows how participatory planning can be structured around indicators without reducing the process to technocratic assessment or, conversely, to unstructured consultation.

Fourth, the study highlights the importance of territorial scale in circular heritage tourism. In all three cases, but especially at Caerlaverock and Blackness, the circularity of tourism could not be assessed at the scale of the monument alone. It depended on relations with mobility networks, village services, local businesses, regional itineraries, surrounding landscapes, and nearby attractions. This suggests that circular cultural tourism should be approached not only as a site-management issue, but as a territorial governance issue.

Fifth, the research reveals the significance of data asymmetries and data gaps. Some dimensions, such as official site significance, conservation pressures, policy priorities, and some visitor figures, were relatively well documented. Others, such as transport behaviour, local spending patterns, waste monitoring, and community perceptions of tourism benefit, were more unevenly evidenced. This has methodological consequences. What cannot be measured or observed consistently tends to remain less visible in decision-making. Strengthening monitoring systems therefore becomes part of the circular transition itself, not merely a technical add-on.

Finally, the cases collectively suggest that peripheral heritage sites should not be interpreted as marginal or secondary versions of major tourism destinations. Rather, they can act as laboratories for more responsible and circular models of tourism development. Their relative remoteness, lower density, and strong territorial embeddedness may enable forms of experimentation less feasible in high-pressure mass-tourism settings. Yet this potential can only be realised if evaluation frameworks are capable of linking heritage value, community aspirations, infrastructural realities, and long-term local development pathways.

More broadly, the study suggests that circular tourism should not be reduced to a weak sustainability perspective focused only on efficiency gains, waste reduction, or technological fixes. In line with the Be.CULTOUR framework, circularity in cultural

tourism also entails local entrepreneurship, community wellbeing, cooperation, cultural heritage appreciation, landscape integrity, and place authenticity. Tourism is therefore not treated as an end in itself, but as an instrument within a wider territorial development perspective, whose value depends on its capacity to strengthen rather than erode the ecological, cultural, and social conditions of the destination.

5 Conclusions

This paper has examined a collaborative experimentation aimed at co-creating Local Action Plans for circular cultural tourism in three Scottish heritage sites (Doune Castle, Caerlaverock Castle, and Blackness Castle) through the integration of circular cultural tourism indicators, the HES Responsible Tourism Framework, and a participatory multicriteria evaluation process.

Four main conclusions emerge. First, circular tourism indicators are most useful when employed as decision-support tools rather than as static descriptive metrics. Their relevance lies in their ability to structure a multidimensional understanding of tourism and to reveal interactions between conservation, access, mobility, interpretation, local economy, and community wellbeing.

Second, the integration between the circular cultural tourism framework and the HES Responsible Tourism Framework proved particularly productive. It enabled the study to connect a multidimensional understanding of circularity with institutional priorities concerning climate action, low-carbon visitor experiences, accessibility, local economic benefits, and community engagement. In doing so, it created a stronger bridge between research-oriented innovation and public-policy relevance.

Third, participatory co-creation functioned as a form of co-evaluation. The Local Action Plans that emerged were not merely operational lists of actions, but the outcome of a process through which evidence, situated knowledge, and strategic judgement were combined. This supports the idea that participatory processes in heritage tourism planning are most valuable when they are evaluative as well as consultative.

Fourth, the three cases confirm the need for context-sensitive and place-based approaches. While the same methodological architecture was applied across the three sites, the priorities that emerged were substantially different: pressure management and landscape fragility at Doune, territorial integration and shared destination logic at Caerlaverock, and infrastructure, accessibility, and decentralised benefits at Blackness. Replicability, therefore, should not be equated with standardisation. It should be understood as methodological coherence combined with local adaptation.

In this paper, the Local Action Plans are not presented as fully developed implementation-ready plans, but as co-created strategic outputs identifying key issues, priority action fields, and context-sensitive directions for future intervention. The approach also presents some limitations. Its effectiveness depends on the availability of sufficiently robust documentary and contextual evidence, on the institutional capacity of the actors involved, and on the presence of stakeholders able and willing to engage in a meaningful co-evaluation process. For this reason, the same methodology may prove more difficult to apply in contexts characterised by fragmented governance, limited data, or weaker policy support. Its transferability should therefore be understood not as direct

replicability of outputs, but as the adaptation of a shared evaluative process to different territorial and institutional conditions.

Overall, the study contributes to current debates on heritage-led regeneration and sustainable tourism transitions by showing how circular tourism principles can be translated into locally grounded action through a participatory evaluative process. In this respect, the paper offers both a methodological contribution to participatory evaluation and a practical contribution to heritage-based tourism planning in peripheral contexts. In this sense, Local Action Plans become not only planning instruments, but also tools for aligning conservation, responsibility, community benefit, and territorial resilience within a shared development vision for heritage-led, responsible, and circular territorial regeneration.

Acknowledgments. P.R. conceptualization, methodology, case studies selection, discussion and conclusions; A.G. conceptualization, methodology, data analysis, introduction, results, discussion and conclusions, manuscript draft and editing, funding acquisition. Authors acknowledge the substantial contribution to the research by Historic Environment Scotland, UK, in particular Dr. Vanessa Glindmeier, Dr. Angela Giancola, Dr. Angus Hulbert who contributed to case studies selection, indicators identification and data collection. The authors also wish to thank the students involved in the Course on The Economics of Circular Tourism at the University of Glasgow, academic years 2025–2026, 2024–2025, 2023–2024 and related academic and workshop activities, as well as all participants in the three local workshops, whose insights and contributions were valuable in shaping the co-evaluative process and the resulting Local Action Plans.

Funding. This research was funded under the framework of European Union's Horizon 2020 research and innovation programme under grant agreement No 101004627, research project Be.CULTOUR "Beyond cultural tourism". This research was funded by Italian Ministry of University and Research (MUR), in the context of the Projects of Relevant National Interest—National Recovery and Resilience Plan (PRIN PNRR). Project Title: Italian historic regeneration through circular ecological heritage communities. Grant Number B53D23029490001.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Kourtit, K., Nijkamp, P., Suzuki, S.: Environmental and cultural tourism in heritage-led regions—performance assessment of cultural-ecological complexes using multivariate data envelopment analysis. Sustainability. **17**, 5871 (2025). https://doi.org/10.3390/su17135871
2. Manniche, J., Larsen, K., Broegaard, R.: The circular economy in tourism: transition perspectives for business and research. Scand. J. Hosp. Tour. **21**, 1–18 (2021). https://doi.org/10.1080/15022250.2021.1921020
3. Gaonkar, S., Sukthankar, S.: Measuring and evaluating the influence of cultural sustainability indicators on sustainable cultural tourism development: scale development and validation. Heliyon. **11**, e42514 (2025). https://doi.org/10.1016/j.heliyon.2025.e42514
4. De Angelis, R., Reynolds, L., Koenig-Lewis, N., Bosangit, C.: conceptualising circular tourism: taking a place-based eco-system perspective. J. Circular Econ. 3 (2025). https://doi.org/10.55845/TVUZ5672

5. Das, A., Mishra, S.: A destination-scale framework for regenerative tourism integrating circular economy and community resilience. Green Technol., Resilience, Sustain. 6 (2026). https://doi.org/10.1007/s44173-026-00027-x
6. Pisolkar, Y.: Cultural heritage management and sustainable development: major themes and research trajectories. J. Electr. Syst. **20**, 2417–2431 (2024). https://doi.org/10.52783/jes.3224
7. Tomassini, L., Cavagnaro, E.: Circular economy, circular regenerative processes, a growth and placemaking for tourism future. J. Tourism Futures **8** (2022). https://doi.org/10.1108/JTF-01-2022-0004
8. Hussein, F., Alhebsi, K.: Adaptive re-use of cultural heritage sites: a strategy for circular economy. Sustainability **17**, 6403 (2025). https://doi.org/10.3390/su17146403
9. Fusco Girard, L., Nocca, F.: From linear to circular tourism 1. AESTIMUM. **70**, 51–74 (2017). https://doi.org/10.13128/Aestimum-21081
10. Jørgensen, M., Hansen, A., Sørensen, F., Fuglsang, L., Sundbo, J., Jensen, J.: Collective tourism social entrepreneurship: a means for community mobilization and social transformation. Ann. Tour. Res. 88 (2021). https://doi.org/10.1016/j.annals.2021.103171
11. Garau, C.: Perspectives on cultural and sustainable rural tourism in a smart region: the case study of Marmilla in Sardinia (Italy). Sustainability **7**, 6412–6434 (2015). https://doi.org/10.3390/su7066412
12. Prayag, G., Hosany, S., Muskat, B., Del Chiappa, G.: Understanding the relationships between tourists' emotional experiences, perceived overall image, satisfaction, and intention to recommend. J. Travel Res. **56**, 41–54 (2017). https://doi.org/10.1177/0047287515620567
13. Gravagnuolo, A., Marasco, A., Fasciglione, M., Xydia, S., Matei, A., Garzillo, C.: Be.CULTOUR: "Beyond CULtural TOURism: human-centred innovations for sustainable and circular cultural tourism" Deliverable 3.1 Protocol/Methodology for human-centred innovation in sustainable and circular cultural tourism-Type of action: IA (Innovation action) (2022)
14. Rodriguez Robaina, C., Florido, C., Jacob, M.: Circular economy contributions to the tourism sector: a critical literature review. Sustainability **12**, 4338 (2020). https://doi.org/10.3390/su12114338
15. Manniche, J., Larsen, K.T., Brandt Broegaard, R., Holland, E.: Destination: A circular tourism economy a handbook for transitioning toward a circular economy within the tourism and hospitality sectors in the South Baltic Region (2017)
16. Lane, B., Kastenholz, E., Carneiro, M.: Rural tourism and sustainability: a special issue, review and update for the opening years of the twenty-first century. Sustainability. **14**, 6070 (2022). https://doi.org/10.3390/su14106070
17. Godovykh, M., Fyall, A., Pizam, A.: Exploring the impacts of tourism on the well-being of local communities. Sustainability **17**, 5849 (2025). https://doi.org/10.3390/su17135849
18. Kausar, D., Nishikawa, Y.: Heritage tourism in rural areas: challenges for improving socio-economic impacts. Asia Pacific J. Tourism Res. **15**, 195–213 (2010). https://doi.org/10.1080/10941661003629995
19. La Rocca, L., et al.: Towards a circular cultural tourism impact assessment framework for decision support in less-known and remote destinations. Presented at the March (2024). https://doi.org/10.1007/978-3-031-65537-1_8
20. De Martino, M., Apicerni, V., Gravagnuolo, A.: Sustainable hospitality and tourism in the Anthropocene era: the need for a more radical shift of the current circular economy models Sustainable hospitality and tourism. Int. J. Contemp. Hosp. Manag. (2024). https://doi.org/10.1108/IJCHM-06-2023-0854
21. Wang, M., Stoffelen, A., Su, M., Zhang, M., Wang, J.: The impacts of heritage conservation and tourism development on the community well-being of agricultural heritage sites from

sustainable development perspective. Sustain. Dev. **33**, 182–201 (2025). https://doi.org/10.1002/sd.70004
22. Hospers, G.-J., Reverda, N.: Managing population decline in europe's urban and rural areas (2015)
23. Bertolini, P., Pagliacci, F.: Quality of life and territorial imbalances. a focus on Italian inner and rural areas. Bio-based Appl. Econ. **6**, 183–208 (2017). https://doi.org/10.13128/BAE-18518
24. Historic Environment Scotland: Responsible Tourisme Framework: transforming heritage tourism (2023)
25. Harrison, K., et al.: Doune Castle Landscape Conservation Management Plan 2 (2021)

Circular Regeneration Pathways for Rural Heritage: The Role of Adaptive Reuse Across 18 European Villages

Piera Della Morte(✉)

Engineering Department, Pegaso Telematic University, Naples, Italy
piera.dellamorte@unipegaso.it

Abstract. The transition to a circular economy in the built environment represents a significant challenge for European territorial policies, particularly in fragile rural areas affected by depopulation, economic decline, and environmental vulnerability. Rural villages, which cover a large portion of Europe's territory, require regeneration models capable of integrating heritage conservation, ecological transition, and socioeconomic revitalization. This article analyzes the role of adaptive reuse as a key "circular invariant" within the broader regeneration process for rural heritage. To do so, a comparative framework was applied to 18 European case studies (10 Italian and 8 international villages). The methodology combines document analysis, qualitative assessment based on source interpretation, fuzzy normalization, cluster analysis, and correlation tests. Two analytical matrices were developed: a diagnostic framework that identifies structural drivers, constraints, catalysts, and triggers; and a strategic framework that examines regeneration actions related to heritage, environment, innovation, energy, and governance. The findings identify three recurring regeneration profiles: (i) transition-oriented models, driven by innovation and environmental performance; (ii) reactive models focused on heritage recovery and territorial resilience; and (iii) strategic governance models supported by institutional capacity or long-term investments. In all clusters, adaptive reuse emerges as a recurring enabling factor, albeit with different functions depending on local conditions.

The study suggests that rural circular regeneration does not follow a single model, but rather combinations of strategies based on the local context, shaped by the territorial context, social capital, and available resources. Adaptive reuse can therefore act as a multidimensional lever to transform abandoned cultural assets into resources for sustainable rural development.

Keywords: Adaptive Reuse · Circular Regeneration · Rural Heritage · Fuzzy Logic · Soil Safeguard · Territorial Resilience

1 Introduction

The transition toward a circular economy in the built environment represents a fundamental challenge for current European cohesion policies, particularly within the framework of the Green Deal. Recent European strategies, from the Green Deal to the New

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 488–505, 2027.
https://doi.org/10.1007/978-3-032-30518-3_31

European Bauhaus, have identified sustainability and circularity as two priority axes for transforming the built environment [1]. At the same time, policies dedicated to rural territories—from the Cork Declaration to the Smart Villages programmes—promote place-based approaches aimed at strengthening local resilience, quality of life, and the valorization of endogenous resources [2, 3]. In this context, the ecological transition does not concern only large urban areas, but also small rural settlements, which cover around 83% of the territory of the European Union and constitute an essential component of Europe's territorial heritage.

However, many European rural villages are currently affected by processes of systemic fragility caused by depopulation, demographic ageing, declining services, and the progressive loss of historical land-management practices. Population decline not only leads to the deterioration of the built heritage through lack of use, but also causes the interruption of ordinary activities—water regulation, terrace maintenance, forest management, and widespread land stewardship—that for centuries ensured the dynamic balance of rural landscapes [4, 5]. The resulting "anthropic void" transforms the cultural landscape into a vulnerable system, exposed to erosion, hydrogeological instability, and increasing fragility in the face of extreme climatic events [6].

In this scenario, village regeneration cannot be interpreted merely as a building restoration process. It requires an integrated approach capable of connecting heritage conservation, environmental innovation, social inclusion, and new local economies [7–11]. In rural contexts, indeed, the village is not an isolated entity, but rather the settlement node of a broader cultural landscape, understood as the outcome of the historical relationship between nature and human communities [12–15]. Architecture, agricultural patterns, ecological infrastructures, and social practices all contribute to the definition of layered territorial identities, whose resilience depends on the ability to maintain an active relationship between settlement and territory [16, 17]. From this perspective, landscape care and the reactivation of the genius loci may represent not only cultural values, but also tools for risk prevention and attractors of new social capital [18].

The Italian situation is particularly significant. The widespread presence of small historic centers undergoing progressive decline has turned the country into a privileged laboratory for experimenting with new territorial regeneration policies [19, 20]. At the same time, the debate on inner areas has highlighted the complexity of the issue: while processes of marginalisation are driven by long-term structural factors, the COVID-19 pandemic, recent energy crises, and the growing demand for more sustainable lifestyles have renewed interest in forms of rural renaissance and new residential geographies [21, 22]. The National Strategic Plan for Inner Areas 2021–2027 also highlights the need to complement anti-depopulation measures with innovative tools for shared service management, intermunicipal cooperation, and support for the most fragile territories [23, 24].

In this context, circular development assumes a central role [25]. In rural villages, circularity does not merely coincide with the optimisation of material flows or energy efficiency, but tends to acquire a territorialised meaning, encompassing the reuse of built heritage, the reactivation of local knowledge, landscape stewardship, community resilience, and the creation of new economic and social functions [7, 26]. Adaptive reuse

of the existing architectural heritage represents, in this perspective, a privileged strategy to contain land consumption, preserve biodiversity and agricultural land, enhance soil permeability, and strengthen hydrogeological resilience [27]. Recent literature also highlights how the success of regeneration processes depends on the ability to balance energy performance, soil protection, and landscape integration [28–32].

Despite growing attention to the circular economy and adaptive reuse, comparative studies dedicated to rural villages remain limited, especially regarding the ways in which these strategies combine with different territorial conditions, levels of fragility, and governance models. In particular, there is still a lack of a systematic interpretation capable of relating regeneration triggering factors to the intensity and nature of the strategies subsequently implemented.

In light of this gap, the present research analyses 18 case studies (10 Italian and 8 European) through the application of a multidimensional evaluation framework [33, 34]. The aim is to identify recurring regeneration configurations in rural villages and to further investigate the role played by adaptive reuse within broader processes of circular transformation [35, 36]. The underlying hypothesis is that effective strategies for territorial revitalization must start from a careful reading of the local status quo, since the material, environmental, and social conditions of each village often already contain the potential trajectory of an authentically place-based regeneration process.

2 Methodology

2.1 Purposeful Case Selection and Data Sources

The cases were deliberately selected to ensure geographical diversity and to represent different regeneration trajectories. Given the exploratory nature of the research and the limited sample size, the results should be interpreted as the identification of patterns rather than predictive classifications. The 18 selected cases are listed in Table 1.

Table 1. Cases study

id	Country	Village	id	Country	Village	id	Country	Village
V1	France	Cahors	V7	Romania	Piscu	V13	Italy	Fiamignano
V2	France	Ungersheim	V8	Austria	Feldkirch	V14	Italy	Calcata
V3	France	Mouans-Sartoux	V9	Italy	S. Stef. di Sessanio	V15	Italy	Torri Superiore
V4	Portugal	Monsanto	V10	Italy	Bagnoregio	V16	Italy	Solomeo
V5	Spain	Fanzara	V11	Italy	Brisighella	V17	Italy	Piaggine
V6	Hungary	Gyermely	V12	Italy	Casale Monferrato	V18	Italy	Santa Fiora

The 18 case studies were selected as representative of a range of possible regeneration strategies for fragile territories. Given the specificity of the topic, and the heterogeneity of rural contexts and achievable objectives, it was necessary to diversify the sample as much as possible, including cases attributable to different regeneration approaches and intentions. Data were collected through document analysis, including academic literature, planning documents, institutional reports, websites, and technical sources (n = 311 sources). This document-based approach made it possible to reconstruct the regeneration process of each village. The document analysis of the case studies revealed several recurring themes (Fig. 1).



Fig. 1. Network map of recurring themes emerging from the comparative analysis of 18 European rural regeneration case studies. Node size indicates thematic frequency, while links represent thematic co-occurrence.

2.2 Concept-Driven Coding Framework and Scoring Protocol: Diagnostic and Strategic Framework

In order to make the case studies comparable, it was considered appropriate to develop a concept-driven coding framework. The multidimensional analysis of the 18 case studies (10 Italian and 8 European) was therefore preceded by the construction of two complementary interpretative frameworks: a diagnostic framework, aimed at identifying the causes and factors that triggered regeneration processes [34], and a strategic framework, intended to classify the main multi-objective strategies adopted in different contexts, thereby enabling the systematic comparison of the case studies. Through an inductive process [37], moving from the particular to the general, the specific causes and the actions concretely implemented in the villages were progressively traced back to

broader categories and recurring fields of intervention considered significant for the sustainable regeneration of rural villages. This synthesis phase proved crucial for the construction of the analytical matrices subsequently employed in the comparative analyses and cluster analysis [36, 38]. The adopted approach makes it possible to organize territorial complexity by breaking down phenomena of crisis and revival into logical categories ordered according to a temporal and functional sequence. The interpretative grid integrates the territorialisation model [39] for the analysis of structural drivers, the push-pull factor dialectic typical of studies on inner areas [40], and the perspective of evolutionary resilience [41], with the aim of identifying the dynamics that caused territorial decline or fostered territorial revitalization through regeneration processes. The "diagnostic" framework classifies the determining conditions and processes into four main macro-categories (Table 2).

Table 2. Diagnostic framework

Macro-Category	Category		
STRUCTURAL DRIVERS	SOCIO-DEMOGRAPHIC DECLINE	**T1**	**-**
(Socioeconomic)	INDUSTRY DECLINE	**T2**	**-**
(-)	LOCAL ECONOMIC DEPRESSION	**T3**	**-**
STRUCTURAL DRIVERS	BUILDING STOCK DECAY	**T4**	**-**
(Landscape-Env.)	HYDROGEOLOGICAL INSTABILITY	**T5**	**+**
(-)	EROSION OF RURAL CHARACTER	**T6**	**-**
STRUCTURAL DRIVERS	EROSION OF COLLECTIVE MEMORY	**T7**	**--**
(Immaterial) (-)/(+)	RECOGNITION OF IDENTITY VALUE	**T8**	**+**
PRECIPITATING FACTORS	HEALTH / ENV. / SEISMIC EMERGENCY	**T9**	**-**
(Triggers)	SPECIFIC HISTORICAL DECOMMISSIONING	**T10**	**-**
(-)	SOCIAL CONFLICT	**T11**	**-**
	ENERGY VULNERABILITY	**T12**	**-**
CONSTRAINTS	LANDSCAPE-ENVIRONMENTAL RESTRICTIONS	**T13**	**+**
(-)/(+)	LAND SCARCITY/ SOIL QUALITY	**T14**	**-**
CATALYSTS	FINANCIAL PUBLIC RESOURCES	**T15**	**+**
(+)	FINANCIAL PRIVATE RESOURCES	**T16**	**+**
	SOCIAL AWARENESS	**T17**	**+**
	EXISTING COOPERATIVE MODEL	**T18**	**+**

Structural Drivers, divided into socio-economic components such as socio-demographic decline, industrial decline, and local economic depression [42–44], landscape-environmental components such as building stock decay, hydrogeological instability, and erosion of rural character [45–47], and intangible components such as erosion of collective memory and recognition of identity value [48], represent the long-term "background" conditions affecting rural areas, such as demographic decline or the erosion of collective memory. Precipitating Factors are interpreted as "trigger" events that accelerate processes of abandonment or require urgent intervention [41, 49, 50]. These include health, environmental, or seismic emergencies, specific historical decommissioning events, social conflict, and energy vulnerability. Constraints are factors that limit action, such as regulatory restrictions [51] or land scarcity and poor soil quality [30]. Catalysts are considered enabling factors, such as public and private financial resources [52], social awareness, and existing cooperation models [53, 54], which positively activate the regeneration process.

The "strategic" framework (Table 3) is useful for constructing a structured multidimensional matrix aimed at describing the transformation strategies through which rural villages can evolve from marginal areas into laboratories of sustainability.

Table 3. Strategic framework

PILLAR	STRATEGIES ADOPTED	
Heritage	ADAPTIVE REUSE	S1
	HERITAGE CONSERVATION	S2
	ENERGY EFFICIENCY	S3
Knowledge and Innovation	MATERIAL EXPERIMENTATION	S4
	KNOWLEDGE TRANSMISSION	S5
	DIGITALIZATION	S6
New Buildings, New Perfomance and Public Spaces	INTEGRATION OF NEW BUILDINGS	S7
	NEAR ZERO ENERGY BUILDING	S8
	PUBLIC SPACES	S9
Environment	SOIL SAFEGUARD	S10
	LANDSCAPE INTEGRATION	S11
	WATER MANAGEMENT	S12
Social/ Economic	PARTICIPATORY GOVERNANCE	S13
	BOTTOM-UP CO-DESIGN	S14
	NETWORK DEVELOPMENT	S15

Each domain defines a pillar of regeneration. The identified pillars are as follows:

Heritage, namely the historic fabric of rural villages—including buildings, roads, squares, and agricultural structures—understood as active capital. Strategies related to the regeneration of this heritage include: Heritage Conservation; Adaptive Reuse, aimed at reintegrating abandoned buildings into the economic cycle and reducing the need for new construction; and Energy Efficiency. These strategies balance the conservation of historical identity with the contemporary need to reduce the carbon footprint of the built environment.

Knowledge and Innovation, insofar as the village can become a centre for the production of intangible value linked to: Material Experimentation, involving the recovery of local construction techniques (e.g., earthen architecture) integrated with new performance standards; Knowledge Transmission, related to traditional know-how; and Digitalization. Digitalisation (Smart Villages) helps bridge geographical gaps, while the transmission of traditional knowledge ensures cultural continuity and creates new specialised employment opportunities.

New Buildings, New Performance and Public Spaces, which considers the integration of the "new" in a synergistic and functional way. This domain includes the return of new public spaces to villages. Strategies include Integration of New Buildings, such as architectural insertions or functional extensions; Near Zero Energy Buildings; and Public Spaces, where the regeneration of vacant urban areas and squares acts as a social "connector," transforming residual spaces into places of aggregation and community resilience.

Environment, related to the climate challenge and the protection of resources. Strategies include Soil Safeguard, meaning zero land consumption and the prioritisation of permeability; Landscape Integration, namely the harmonious relationship between the built environment and the agricultural/natural landscape; and Water Management, involving the implementation of water recovery and circular management systems, which are essential to counter drought in fragile rural contexts.

Social/Economic, since long-term sustainability depends on human involvement and financial solidity. Strategies include Participatory Governance; Co-design, whereby community involvement ensures that reuse responds to the real needs of inhabitants, increasing acceptance and stewardship of the project; and Network Development, meaning the creation of networks among villages and regional stakeholders (public and private), enabling circular economies to scale and attracting investments that a single municipality could not manage alone.

Each variable was assessed using a five-point scale, based on the relevance of the cause/catalyst or on the intensity of implementation of the strategies emerging from the literature review. Scores were assigned by the author through expert coding based on the documented information available for each case study. The scores were then converted into normalised fuzzy values ranging from 0 to 1 in order to improve comparability among cases (Table 4).

The matrices relating to the diagnostic framework (Table 5) and those relating to the strategic framework (Table 6) were then constructed accordingly.

Table 4. Scoring protocol

Diagnostic framework	Strategic Framework	Score
Very relevant	Fundamental pillar of the strategy (project excellence)	1
Somewhat relevant	Strategy successfully and relevantly implemented	0,75
Neither relevant nor irrelevant	Strategy present but not defining or standard	0,5
Slightly relevant	Strategy partially or with difficulty implemented	0,25
Not at all relevant	Strategy absent or with significant critical issues	0

Table 5. Diagnostic Matrix

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	V13	V14	V15	V16	V17	V18
T1	0,75	0,75	1	0,75	0,75	0	1	0,5	0,75	0,5	0,25	1	0,75	0,5	0,5	0,75	0,75	0,75
T2	0,75	0,75	0,75	0,75	0,75	0	0,75	0,75	1	0,75	0,75	0,75	1	0,75	0,75	0,75	1	0,75
T3	1	1	0,5	1	1	0,25	0,75	1	1	0,75	0,75	0,75	1	1	1	1	0,75	0,75
T4	0,5	0,5	0,75	0,5	0,5	0,5	0,5	1	0,75	1	1	0,75	0,75	1	0,75	0,5	0,5	0,5
T5	0,75	0,75	1	0,75	0,75	0,75	0,75	0,75	0,75	0,25	0,75	0,75	0,75	0,75	0,5	0,75	0,75	0,75
T6	0,75	0,75	0,75	0,75	0,5	0	0,75	0,25	0,75	0,75	0,5	0,5	0,75	0,5	0,25	0,75	0,5	0,5
T7	1	1	0,75	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T8	1	1	1	0,5	0,75	0,5	0	0,75	1	1	0,75	1	1	1	0,5	0,5	0,5	0,75
T9	0,75	1	0,75	0,75	0,75	0,75	0,75	1	0,75	0,75	0,75	1	0,75	1	1	0,75	0,75	0,75
T10	0,5	0,5	0,75	0,5	1	0,5	0,5	0,75	0,5	0,5	0,5	0,75	0,5	0,75	0,5	0,5	0	0,5
T11	1	1	0,75	0,75	0,5	0,75	1	0,75	0,75	0,75	0,75	0,5	0,75	0,75	0,75	0,5	0,75	0,75
T12	1	0,75	1	0,75	0,5	0,75	0,75	0,75	1	1	1	1	0,75	1	1	1	1	1
T13	0,75	0,75	1	0,75	0,5	0,75	0,5	1	0,5	1	0,75	0,5	0,5	1	1	0,75	0,5	0,5
T14	1	1	1	1	0,75	0,75	0,5	1	0,75	0,75	1	1	1	0,75	0,25	0,25	0,75	1
T15	0,75	0,75	1	0,75	0,75	1	1	0,75	1	0,75	0,75	0,75	0,75	1	1	1	0,25	0,75
T16	0,75	1	1	0,75	1	1	1	1	0,75	1	1	1	0,75	1	1	0,75	1	0,75
T17	1	1	1	1	1	1	0,75	1	1	0,75	1	0,75	1	0,75	1	0,75	1	0,75
T18	1	1	1	1	1	1	0,75	1	1	0,75	1	0,75	1	0,75	1	0,75	1	0,75

Table 6. Strategic Matrix

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	V13	V14	V15	V16	V17	V18
S1	1	1	0,75	0,75	1	0,75	0,75	0,75	1	1	1	0,75	1	1	1	0,75	1	1
S2	1	0,75	0,75	1	0	0,5	1	1	1	1	1	1	0,75	1	1	1	0,75	1
S3	1	1	0,75	0,5	0,5	1	0,75	1	0,75	0,25	0,75	0,5	0,75	0,25	0,75	0,75	0,75	0,75
S4	0,75	0,75	0,5	0,75	0,75	0,75	1	1	1	0,5	0,5	0,75	0,5	0,5	0,75	0,5	0,5	0,5
S5	1	1	1	1	1	1	1	1	1	0,75	1	1	0,75	1	1	1	1	0,75
S6	0,75	0,5	0,5	0,75	0,75	0,75	0,75	0,75	0,5	0,75	0,5	0,5	1	0	0,5	0,5	0,75	1
S7	0,75	0,75	0,5	0	0	0,75	1	1	0	0	0	0,25	0,25	0	0,5	0,75	0,25	0
S8	0,5	1	0,5	0,5	0	0,75	0,75	1	0,25	0	0	0,25	0,75	0	0,75	0,5	0,5	0,5

(*continued*)

Table 6. (*continued*)

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	V13	V14	V15	V16	V17	V18
S9	1	1	0,75	0,75	1	1	0,75	1	0,75	0,75	1	1	1	0,75	1	0,75	1	0,75
S10	0,75	1	1	1	0,75	1	0,75	1	1	1	0,75	1	0,75	1	1	1	0,75	0,75
S11	1	1	1	1	1	1	1	0,75	1	1	1	1	1	1	1	1	1	0,75
S12	0,75	0,75	0,75	0,75	0	0,75	0,5	0,75	0,5	0,5	0,5	0,75	0,5	0,5	0,75	0,75	0,5	1
S13	1	1	1	0,75	1	1	0,75	1	0,75	0,75	0,75	1	1	0,75	1	0,25	1	0,75
S14	0,75	1	1	0,75	1	0,75	1	0,75	0,75	0,5	0,75	0,75	0,75	1	1	0	1	0,25
S15	1	1	1	1	1	1	1	0,75	1	1	1	1	1	0,75	1	0,75	0,75	1

2.3 Data-Set Analysis

By combining the two matrices, a normalised dataset was obtained, consisting of 18 cases and 33 variables, of which 18 represent the initial territorial conditions (T) and 15 represent the implemented response strategies (S). To achieve a multidimensional understanding of the dataset, the following complementary procedures were adopted: descriptive statistics; contingency tables with Pearson's Chi-square test; and K-means cluster analysis (K = 3) (Table 7), using the PSPP software [55, 56]. This combination made it possible both to analyse individual variables and to identify overall patterns among the cases. First, descriptive statistics were calculated, including mean, standard deviation, skewness, kurtosis, minimum, maximum, and range. Subsequently, contingency tables were produced between T and S variables using Pearson's Chi-square test, with the aim of identifying possible relationships between the two groups of indicators. Associations between territorial drivers (T) and adopted strategies (S) are considered statistically significant if the p-value is lower than 0.05, and as borderline trends if the p-value is between 0.05 and 0.10.

Table 7. Village Clusters

CLUSTER 1	CLUSTER 2	CLUSTER 3
Ungersheim	*Mouans-Sartoux*	*Cahors*
Gyermely	*Monsanto*	*Piscu*
Feldkirch	*Fanzara*	*Corciano-Solomeo*
Torri Superiore	*S. Stefano di Sessanio*	*Santa Fiora*
	Bagnoregio	
	Brisighella	
	Monferrato Casalese	
	Fiamignano	
	Calcata	
	Piaggine	

Finally, the theme of adaptive reuse of heritage was further explored through the construction of binary 0/1 matrices, identifying reuse functions, together with energy efficiency measures, the use of renewable energy sources, and the use of traditional materials/bio-materials.

2.4 Limitations

Among the main methodological limitations of the study are the small sample size, which limits the generalisability of the results and the robustness of the statistical analyses; the dependence on secondary sources for the construction of the dataset; and the presence of interpretative scoring procedures, which are inevitably exposed to margins of subjectivity in the coding of indicators. Moreover, the exploratory nature of the adopted techniques suggests caution in interpreting the results. Therefore, the associations identified should be considered preliminary and exploratory evidence rather than causal relationships or generalisable models.

3 Results

The analysis of the diagnostic framework shows that, in most cases, the loss of collective memory is perceived as a central critical issue, social awareness is high, and pre-existing cooperative models constitute a concrete lever for action. The analysis of the strategic framework indicates that the examined cases prioritise landscape-integrated strategies, based on knowledge transmission, oriented toward territorial cooperation, and founded on the reuse of existing heritage.

The correlation analysis reveals a system of relationships between certain determining factors and regeneration strategies, highlighting recurring configurations rather than isolated relations. In contexts characterised by socio-demographic decline (depopulation, ageing, reduction of the resident population), the conservation of built heritage tends to emerge as a priority strategic response. A relationship close to statistical significance also emerged between demographic decline and the adaptive reuse of underused, abandoned, or no longer functional buildings. Even where not statistically strong, some relationships appear conceptually plausible: When collective memory weakens, practices of transmitting local, artisanal, and cultural knowledge become central; the presence of pre-existing cooperative models may facilitate participatory governance in regeneration processes; greater social awareness tends to favour shared planning processes and bottom-up co-design; and contexts vulnerable from an energy perspective are naturally more oriented toward efficiency and energy self-sufficiency strategies.

The cluster analysis produced the following results. Cases belonging to Cluster 1 show convergence toward maximum scores in the dimensions of Knowledge and Innovation and Environment. Cases such as Gyermely (HU) and Ungersheim (FR) stand out for the integration between local production chains and resource management (Water Management and Soil Safeguard). Cluster 2, the largest in the sample, reflects a condition of structural fragility. It shows systematic polarisation toward Adaptive Reuse and Heritage Conservation strategies (constant scores), contrasted with lower values in technological performance (Energy Efficiency). Cluster 3 groups cases representing regeneration

models driven by a strategic vision supported by significant public or private capital. The matrix shows high and balanced scores across all 15 strategies, with peaks in the Social/Economic dimensions (Network Development and Participatory Governance).

The integration of this matrix with the analysis of functions identified in the 18 case studies makes it possible to complete the analysis of adaptive reuse through the lens of functional mixité. Three main functional clusters emerge. The first configures the village as a productive and knowledge hub, where adaptive reuse supports local supply chains, energy production, food processing, and training centres, transforming the built heritage into infrastructure for a circular metabolism. The second cluster, definable as a model between dispersed hospitality and culture, includes villages that rely on hospitality, cultural functions, and craftsmanship as economic drivers, enhancing historic and landscape heritage to counter decline. The third cluster concerns the integration of advanced services, with coworking, offices, and local services, showing how strategic investments can attract urban functions and new temporary residents. In all clusters, the creation of public space and cultural centres emerges as a recurring element.

From a technological perspective, a low-tech approach prevails, consistent with the constraints of historic rural contexts. The most widespread solutions are internal insulation, high-performance windows, the use of natural thermal inertia, and natural ventilation, all of which recover bioclimatic principles of vernacular architecture. More advanced systems, such as mechanical controlled ventilation or solar greenhouses, appear only in a few pilot cases.

For energy production, strong attention is paid to landscape integration. Photovoltaic roof tiles are more widespread than traditional panels in historic villages, while geothermal energy and biomass represent relevant solutions for local heating. In some cases, innovative management models also appear, such as Renewable Energy Communities (RECs) and digital systems for monitoring consumption.

From a material perspective, regeneration is based on a combination of traditional local materials and bio-materials. Central is the use of hemp (lime-hemp, hemp-clay), alongside local timber, plant fibres, cork, sheep wool, blown cellulose, raw earth, and natural lime. This confirms that the circular regeneration of villages is based on new construction techniques that integrate vernacular knowledge, environmental compatibility, and bio-ecological innovation

4 Discussion

The comparative analysis carried out shows that the circular regeneration of rural territories cannot be reduced to a single operational model. Rather, the eighteen case studies reveal a plurality of trajectories determined by different combinations of territorial resources, institutional capacity, socio-economic structure, environmental fragilities, and cultural capital. This finding is particularly relevant when compared with part of the circular economy literature, which is still predominantly oriented toward urban and building-centred models focused on material flows, energy retrofit, and technological efficiency [57, 58]. In fragile rural contexts, circularity tends instead to acquire a territorialised meaning, encompassing the reuse of built heritage, the reactivation of local knowledge, landscape stewardship, community resilience, and the generation of new economic and social functions [59–61].

The comparative analysis between triggering drivers and the intensity of implemented strategies makes it possible to identify three recurring regeneration configurations in European villages. The circularity paradigm is not applied uniformly, but adapts to the nature of initial conditions, generating differentiated development trajectories:

Evolutionary Transition and Bio-Economic Metabolism. In this configuration, regeneration does not arise as a response to an explicit crisis, but as a proactive choice of territorial transformation. Circularity is expressed through the integration of local production chains, resource management, and new forms of self-sufficiency. Cases such as Mouans-Sartoux and Ungersheim show how food sovereignty, innovative public policies, and ecological planning can configure the village as a territorial laboratory of sustainability [61, 62]. In such contexts, adaptive reuse goes beyond the mere conservation of buildings and assumes the role of "bio-economic infrastructure", while soil protection coincides with safeguarding local natural capital.

Reactive Emergency and Identity Resilience. In this second configuration, structural fragility, depopulation, and physical vulnerability constitute the initial drivers of regeneration. Circularity takes on an adaptive function: the reuse of existing heritage becomes a necessary response to resource scarcity and to the impracticability of traditional expansionary models. The reactivation of consolidated urban voids thus becomes a form of territorial stewardship against abandonment. Santo Stefano di Sessanio, Fiamignano, Brisighella, and Civita di Bagnoregio show how landscape or geomorphological constraints can be reinterpreted as factors of cultural and identity-based attractiveness [63, 64]. In these cases, resilience derives from the ability to convert conditions of marginality into new territorial opportunities.

Strategic Vision and Institutional Governance. The third model includes processes guided by a holistic top-down vision, supported by significant public or private capital. Strategic profiles are high and relatively balanced across the entire analytical framework. In such contexts, institutional capacity or corporate patronage makes it possible to integrate conservation, innovation, territorial networks, and the quality of public space. Solomeo, Gyermely, and Feldkirch show how adaptive reuse can be elevated to an advanced standard of aesthetic and functional quality, while digitalisation and the refunctionalisation of collective spaces act as catalysts for new settlement and professional flows [65]. Soil protection and urban permeability also become components of a territorial strategy based on sustainability and quality of life.

Beyond the differences among the three identified clusters, the comparative analysis highlights several recurring cross-cutting themes present in numerous regeneration processes. These recurrences should not be interpreted as universally valid recipes, but as recurring factors whose impact depends on the institutional, social, economic, and territorial conditions in which they are activated. Their effects may therefore vary and are not without controversial aspects.

A first cross-cutting theme concerns the role of *local governance*. The most coherent and durable transformation processes are found where public administrations, cooperatives, or strong territorial actors have been able to coordinate integrated policies rather than isolated interventions. The case of Mouans-Sartoux shows how local food sovereignty, sustainable public procurement, and the *Maison d'Éducation*

à l'Alimentation Durable transformed food policy into a territorial strategy [62, 66]. Similarly, Feldkirch demonstrates that even small centres can integrate heritage, climate neutrality, and administrative innovation when supported by solid institutional capacity. However, the centrality of governance also raises controversial issues: not all territories possess the same administrative skills, planning capacity, or civic leadership, with the risk that regeneration may concentrate in already relatively strong contexts, thereby widening pre-existing territorial inequalities.

A second theme concerns the relationship between *culture, identity, and regeneration.* Several cases suggest that symbolic and social reactivation may precede economic reactivation. Fanzara transformed local conflict into a process of civic mediation through urban art, generating new territorial attractiveness [67, 68]. Piscu shows how the recovery of ceramic craftsmanship and the creation of a museum-workshop can become tools of heritage pedagogy. Calcata confirms the potential of creative communities and spontaneous reuse in countering village abandonment. These experiences show that rural regeneration depends not only on infrastructural investment, but also on the capacity to generate new shared meanings. At the same time, identity-led processes may create tensions between authenticity and local memory on the one hand, and external cultural consumption on the other, especially when symbolic valorisation prevails over the everyday needs of residents.

The recurring creation of *public spaces* following adaptive reuse interventions also suggests that circular village regeneration often begins with the reconstruction of social capital. The reuse of architectural and urban heritage does not merely produce new functional containers, but recreates the place of relationships: the square, understood in both physical and symbolic terms, as the beating heart of community life. From this perspective, reuse also appears as a process of resemanticisation, capable of enriching the village's system of intangible values in addition to its material assets. However, the mere creation of refurbished spaces does not automatically guarantee social inclusion or lasting participation: in the absence of active communities, some spaces risk remaining underused or dependent on occasional events. At the same time, the comparison invites caution against overly celebratory readings of tourism-led strategies. Some of the best-known cases reveal significant structural tensions. Civita di Bagnoregio represents an emblematic paradox, where heritage valorisation coexists with geomorphological risk, increasing tourist pressure, and the progressive reduction of residential functions [64, 69]. When the visitor economy becomes dominant, villages risk turning into consumable images rather than living communities. Tourism success therefore does not automatically coincide with sustainable territorial regeneration.

This observation is particularly relevant for the *Albergo Diffuso* model, often presented as a paradigmatic response to depopulation in minor historic centres. The case of Santo Stefano di Sessanio confirms the model's capacity to reactivate unused properties, attract economic resources, and improve the architectural quality of heritage assets. However, the comparison suggests that dispersed hospitality cannot automatically be equated with comprehensive regeneration. In the absence of permanent services, stable residency, economic diversification, and inclusive governance, such models may remain exposed to seasonality, dependence on external demand, and social fragility. The risk is that villages may be physically restored without fully reconstructing their civic fabric.

From a constructive and environmental perspective, the *low-tech approach* to energy efficiency, rooted in the principles of traditional and bioclimatic architecture, emerges as particularly coherent with rural contexts. The use of bio-materials and local materials for thermo-hygrometric insulation and envelope retrofitting contributes to soil and environmental protection through three main levers: reduction of embodied energy, improved healthiness and building permeability, and valorisation of local productive landscapes. Cahors experimented with energy retrofits using bio-based materials in historic heritage [58], while Piaggine and Vorarlberg show the connection between building reuse and new forest-based and territorial economies. In this sense, the circular regeneration of villages should be read as reintegration into material cycles. However, it must also be noted that such solutions require technical skills, active local supply chains, and compatible regulatory frameworks, which are not always present in more marginal territories.

Likewise, the *decarbonisation* of villages does not follow a single technological path. Rather, there is a tendency toward forms of discreet electrification—such as geothermal systems or integrated photovoltaic tiles—in contexts of high historical value, while collective systems such as biomass plants or Renewable Energy Communities emerge where social cohesion and public support are stronger. In all cases, energy production tends to remain confined within the perimeter of the built environment or nearby resources, avoiding new occupation of virgin land and confirming the role of the regenerated village as an energetically dense and relatively self-sufficient territorial unit. Nevertheless, open questions remain concerning economic replicability, long-term maintenance, and the managerial capacity required by such systems.

Overall, the results suggest that effective regeneration of fragile territories rarely depends on a single intervention. Rather, it emerges from the integration of adaptive reuse, governance capacity, identity activation, ecological stewardship, and economic diversification. However, despite the plurality of observed models, one recurring mechanism cuts across all identified clusters: the transformation of unused built heritage into new social, productive, cultural, or environmental functions. For this reason, although it is not possible to derive a universal model of rural regeneration, the collected evidence suggests interpreting adaptive reuse as a circular invariant, that is, a possible structural constant capable of connecting different trajectories of territorial transformation.

5 Conclusions

The research highlights that Adaptive Reuse, in synergy with Soil Safeguard, represents a recurring element in the regeneration processes of fragile rural villages, which are primarily based on the enhancement of existing resources rather than on expansionary models. The analysis nevertheless shows that its role varies according to context: in evolutionary clusters, it assumes a performative and innovative value; in critical scenarios, it performs an adaptive and conservative function; while in strategic models, it becomes a systemic lever for governance, local networks, and long-term investment.

The results also suggest that effective regeneration depends on the integration of adaptive reuse, governance, environmental innovation, public space quality, and economic diversification, while remaining exposed to limitations such as intensive tourism, administrative fragility, or dependence on external capital. Although the study does not

allow the definition of a universal model, it proposes adaptive reuse as a circular invariant, that is, a recurring constant capable of connecting different trajectories of territorial regeneration.

Funding. This research was funded by Italian Ministry of University and Research (MUR), in the context of the Projects of Relevant National Interest—National Recovery and Resilience Plan (PRIN PNRR). Project Title: Italian historic regeneration through circular ecological heritage communities. Grant Number B53D23029490001.

References

1. European Commission: New European Bauhaus: Beautiful, Sustainable, Together (2021). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52021DC0573
2. European Commission: Smart Rural 27 Project (2023)
3. European Union: Cork 2.0 Declaration "A Better Life in Rural Areas." Publications Office of the European Union, Luxembourg (2016)
4. Pettenati, G.: I paesaggi culturali Unesco in Italia (2019)
5. Agnoletti, M., Rotherham, I.D.: Landscape and biocultural diversity. Biodivers. Conserv. **24**, 3155–3165 (2015)
6. Tarolli, P., Rizzo, D., Brancucci, G.: Terraced landscapes: land abandonment, soil degradation, and suitable management. In: World Terraced Landscapes: History, Environment, Quality of Life, pp. 195–210. Springer (2018)
7. Horizon CLIC Project: CLIC: Circular models Leveraging Investments in Cultural heritage adaptive reuse, CLIC, Project, Fact Sheet, H2020. https://cordis.europa.eu/project/id/776758. Accessed 28 July 2025
8. Gravagnuolo, A., Fusco Girard, L., Kourtit, K., Nijkamp, P.: Adaptive re-use of urban cultural resources: contours of circular city planning. City Cult. Soc. **26**, 100416 (2021). https://doi.org/10.1016/j.ccs.2021.100416
9. Donders, Y.M.: Towards a right to cultural identity? Leiden J. Int. Law. **16**, 639–669 (2003)
10. He, Z., Chen, M., Gu, D.: How digital village construction affects to the effectiveness of rural governance?—research on the NCA and QCA methods. Cities. **156**, 105514 (2025)
11. Purvis, B., Else, T., Genovese, A., Jimenez, A., Venkataraman Guru, R.: Grassroots initiatives for a bottom-up transition to a circular economy: exploring community repair. Local Environ. **30**, 1240–1256 (2025). https://doi.org/10.1080/13549839.2025.2467385
12. Sauer, C.O.: The morphology of landscape. 2007 Found. Pap. Landsc. Ecol. 36–70 (1925)
13. Council of Europe: European Landscape Convention. (2000)
14. Brumann, C., Gfeller, A.É.: Cultural landscapes and the UNESCO world heritage list: perpetuating European dominance. Int. J. Herit. Stud. **28**, 147–162 (2022)
15. Magnaghi, A.: Politiche e progetti di territorio per il ripopolamento rurale (2012)
16. Foster, J.B.: The Return of Nature: Socialism and Ecology. NYU Press (2020)
17. Zhang, X., Zhang, Z., Zhang, X., Zhang, Z.: How do smart villages become a way to achieve sustainable development in rural areas? Smart Village planning and practices in China. Sustainability. **12** (2020). https://doi.org/10.3390/su122410510
18. Malfatti, F., Grisot, F.: Paesaggi rurali, resilienza e innovazione. Antropologi ed antropologhe a supporto delle Comunità. Antropol. Pubblica. **10**, 381–395 (2024)
19. Mantino, F., Lucatelli, S.: Le aree interne in Italia: un laboratorio per lo sviluppo locale. Agri. **45**, 1–3 (2016)

20. Secundo, G., Ndou, V., Vecchio, P.D., De Pascale, G.: Sustainable development, intellectual capital and technology policies: a structured literature review and future research agenda. Technol. Forecast. Soc. Change. **153**, 119917 (2020). https://doi.org/10.1016/j.techfore.2020.119917
21. Hennessy, T.: Can a rural renaissance be a reality? Eur. View. **23**, 64–71 (2024). https://doi.org/10.1177/17816858241242882
22. Ivanko, J.D., Kivirist, L.: Rural Renaissance: Renewing the Quest for the Good Life. New Society Publishers (2009)
23. Presidenza del Consiglio dei Ministri: PSNAI - Piano Strategico Nazionale per le Aree Interne (2025). https://politichecoesione.governo.it/it/documenti-ed-esiti-istituzionali/documenti-strategici-di-inquadramento/programmazione-2021-2027/piano-strategico-nazionale-delle-aree-interne-2021-2027-psnai-e-allegati/
24. Lega delle Autonomie Locali: Ricucire-i-territori-position-paper-Aree-interne-impaginato. https://aliautonomie.it/wp-content/uploads/2025/07/Ricucire-i-territori-position-paper-Aree-interne-impaginato-1.pdf.
25. Williams, J.: Circular cities. Urban Stud. **56**, 2746–2762 (2019). https://doi.org/10.1177/0042098018806133
26. Horizon Europe: Circular bioeconomy start up villages (HORIZON CL6 2024 CircBio 01 9). https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/calls-for-proposals?callIdentifier=HORIZON-CL6-2024-CIRCBIO-01&isExactMatch=true&status=31094501,31094502,31094503&order=DESC&pageNumber=1&pageSize=50&sortBy=startDate. Accessed 28 July 2025.
27. Williams, J.: Circular Cities: A Revolution in Urban Sustainability. Routledge, NY Routledge 2021, Abingdon, Oxon, New York (2021). https://doi.org/10.4324/9780429490613.
28. Parlato, M.C., Pezzuolo, A.: Sustainable reuse of traditional rural buildings: review, conceptual framework, and future research directions. In: International Mid-Term Conference of the Italian Association of Agricultural Engineering, pp. 1032–1038. Springer (2024)
29. Directive (EU) 2025/2360 of the European Parliament and of the Council of 12 November 2025 on soil monitoring and resilience (Soil Monitoring Law) (2025)
30. Munafò, M. ed: Consumo di suolo: Dinamiche territoriali e servizi ecosistemici (2021). https://iris.unibas.it/bitstream/11563/167965/1/ISPRA_2021_Saganeiti_et_alii.pdf
31. Garda, E.: Un nuovo ruolo per gli spazi urbani sottoutilizzati tra sottrazioni e permeabilità. Presented at the (2020)
32. Fusco Girard, L., Gravagnuolo, A. (eds.): Adaptive Reuse of Cultural Heritage: Circular Business, Financial and Governance Models. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4
33. Rovai, M., Andreoli, M., Gorelli, S., Jussila, H.: A DSS model for the governance of sustainable rural landscape: a first application to the cultural landscape of Orcia Valley (Tuscany, Italy). Land Use Policy. **56**, 217–237 (2016)
34. Li, Y., Westlund, H., Liu, Y.: Why some rural areas decline while some others not: an overview of rural evolution in the world. J. Rural. Stud. **68**, 135–143 (2019). https://doi.org/10.1016/j.jrurstud.2019.03.003
35. Grecchi, M., Colucci, A., Malighetti, L.E., Speciale, F.: Reinventing Mountain and Rural Villages: Strategies for Regeneration in the Perspective of a Circular Economy. Springer Nature (2025)
36. Buglione, F., Morte, P.D., Angrisano, M., Gravagnuolo, A., Girard, L.F.: Circular models for the sustainable regeneration of Italian rural villages: a critical analysis of good practices toward the definition of a circular Rural Village. Sustainability. (2026). https://doi.org/10.3390/su18031405
37. Thomas, D.R.: A general inductive approach for analyzing qualitative evaluation data. Am. J. Eval. **27**, 237–246 (2006). https://doi.org/10.1177/1098214005283748

38. Sheydayi, A., Dadashpoor, H.: Conducting qualitative content analysis in urban planning research and urban studies. Habitat Int. **139**, 102–878 (2023). https://doi.org/10.1016/j.habitatint.2023.102878
39. Magnaghi, A.: Il principio territoriale. Bollati Boringhieri, Torino (2020)
40. Barbera, F., Rossi, A.D.: Metromontagna: Un progetto per riabitare l'Italia. Donzelli Editore (2021)
41. Davoudi, S., et al.: Porter: resilience: a bridging concept or a dead end? Edited by Simin Davoudi and Libby porter. Plan. Theory Pract. **13**, 299–333 (2012). https://doi.org/10.1080/14649357.2012.677124
42. Stockdale, A.: Migration: pre-requisite for rural economic regeneration? J. Rural. Stud. **22**, 354–366 (2006). https://doi.org/10.1016/j.jrurstud.2005.11.001
43. Lewis, G.: Rural migration and demographic change. In: The Geography of Rural Change. Routledge (1998)
44. Grabher, G.: The weakness of strong ties; the lock-in of regional development in Ruhr area. Embed. Firm Socioecon. Ind. Netw., 255–277 (1993)
45. Cerasoli, M., Mattarocci, G.: Rigenerazione urbana e mercato immobiliare. Roma TrE-Press (2018)
46. Amaduzzi, S., Bressan, G., Guaran, A., Pascolini, M.: Gian Pietro Zaccomer: I paesaggi che cambiano: quali risposte contro il degrado? Presented at the (2019)
47. MacDonald, D., et al.: Agricultural abandonment in mountain areas of Europe: environmental consequences and policy response. J. Environ. Manag. **59**, 47–69 (2000). https://doi.org/10.1006/jema.1999.0335
48. Nora, P.: Between memory and history: les Lieux de Mémoire. Representations, 7–24 (1989). https://doi.org/10.2307/2928520
49. Landslide Hazard and Risk: Wiley Online Books. https://onlinelibrary.wiley.com/doi/book/10.1002/9780470012659#page=194. Accessed 30 Mar 2026
50. Thomson, H., Bouzarovski, S.: Addressing Energy Poverty in the European Union: State of Play and Action (2019). https://energy-poverty.ec.europa.eu/system/files/2024-05/paneureport2018_updated2019.pdf
51. Settis, S.: Settis, S: Paesaggio costituzione cemento: la battaglia per l'ambiente contro il degrado civile, pp. 3–313. Einaudi, Torino (2010)
52. Camagni, R.: Regional competitiveness: towards a concept of territorial capital. In: Capello, R. (ed.) Seminal Studies in Regional and Urban Economics: Contributions from an Impressive Mind, pp. 115–131. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-57807-1_6
53. Ostrom, E.: Governing the Commons: the Evolution of Institutions for Collective Action. Cambridge University Press (1990)
54. Martini, E.: La comunità tra nuove sfide e progettualità : capitale sociale, resilienza e innovazione. 1–131 (2025).
55. Niu, G., Segall, R.S., Zhao, Z., Wu, Z.: A survey of open source statistical software (OSSS) and their data processing functionalities: Int. J. Open Source Softw. Process. **12**, 1–20 (2021). https://doi.org/10.4018/IJOSSP.2021010101
56. Ikotun, A.M., Ezugwu, A.E., Abualigah, L., Abuhaija, B., Heming, J.: K-means clustering algorithms: a comprehensive review, variants analysis, and advances in the era of big data. Inf. Sci. **622**, 178–210 (2023). https://doi.org/10.1016/j.ins.2022.11.139
57. Marchettini, N., Brebbia, C.A., Pulselli, R., Bastianoni, S.: The Sustainable City IX: urban regeneration and sustainability (2 volume set). WIT Press. (2014)
58. Claude, S., et al.: Evaluating retrofit options in a historical city center: relevance of bio-based insulation and the need to consider complex urban form in decision-making. Energ. Buildings. **182**, 196–204 (2019). https://doi.org/10.1016/j.enbuild.2018.10.026

59. Silva, L.M.S. de: Built heritage-making and socioeconomic renewal in declining rural areas (2012)
60. Grodwohl, M.: Architecture vernaculaire et paysages. J. Anthropol. Assoc. Fr. Anthropol., 221–241 (2013). https://doi.org/10.4000/jda.4808
61. Koninckx, G.: Les parcours de transition d'Ungersheim : Être aspiré par le chaos des crises ou emprunter des trajectoires de résiliences. Acta Eur. Syst. **8**, 307–312 (2020). https://doi.org/10.14428/aes.v8i1.56483
62. Stefanovic, L.: SDG performance in local organic food systems and the role of sustainable public procurement. Sustainability. **14**, 11510 (2022). https://doi.org/10.3390/su141811510
63. Morena, M., Truppi, T., Del Gatto, M.L.: Sustainable tourism and development: the model of the Albergo Diffuso. J. Place Manag. Dev. **10**, 447–460 (2017). https://doi.org/10.1108/JPMD-08-2016-0057
64. Amer, M., Ginzarly, M., Renzi, M.-F.: Civita di Bagnoregio, Italy: towards a people-centred heritage branding approach. J. Herit. Tour. **18**, 483–503 (2023). https://doi.org/10.1080/1743873X.2023.2188450
65. Kim, J.Y.E.: The multi-faceted practice of architectural sustainability found in the regional architecture of Vorarlberg, Austria. Sustainability. **15**, 2430 (2023). https://doi.org/10.3390/su15032430
66. Lulovicova, A., Bouissou, S.: Environmental assessment of local food policies through a territorial life cycle approach. Sustainability. **15**, 4740 (2023). https://doi.org/10.3390/su15064740
67. Currás, R., Escriva, M.: Case study: Don't say street art, just say Fanzara. In: Entrepreneurship in Culture and Creative Industries, pp. 265–279. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-65506-2_15
68. Kassam, A.: How embracing graffiti stopped one Spanish village going to the wall (2015). https://www.theguardian.com/artanddesign/2015/apr/14/street-art-fanzara-spain-graffiti-artists
69. Petroni, M.: Civita di Bagnoregio, the dying town between depopulation and overtourism. Domus. (2021)

Heritage-Led Regeneration and Multidimensional Impact Assessment in a Historic Coastal City: The Case of the Community Cooperative *"Identity and Beauty"* in Sciacca, Sicily (Italy)

Antonia Gravagnuolo[1(✉)], Francesca Buglione[1], and Mariarosaria Angrisano[2]

[1] CNR ISPC, Institute of Heritage Science, Naples, Italy
antonia.gravagnuolo@cnr.it
[2] Engineering Department, Pegaso Telematic University, Naples, Italy

Abstract. Historic small and medium-sized cities are increasingly recognised as strategic laboratories for place-based regeneration, particularly when cultural heritage is activated through community-led initiatives. However, evaluation frameworks capable of assessing the multidimensional impacts of heritage-led tourism remain limited, especially in peripheral and fragile territories. This paper addresses this gap through the case of Sciacca, a historic coastal city in southern Sicily, where the Community Cooperative Identità e Bellezza ("Identity and Beauty") developed the Museo Diffuso dei 5 Sensi ("Widespread Museum of the 5 Senses") as a model of community-based, heritage-led regeneration.

The initiative mobilises local cultural, social, and environmental resources through experiential tourism, adaptive reuse, storytelling, and collaborative governance, coherently with the Heritage Community approach promoted by the Faro Convention. The paper adopts the Circular Cultural Tourism Dashboard developed within the Horizon 2020 Be.CULTOUR project as a multidimensional and multicriteria assessment framework structured around six dimensions: cultural, social and human, environmental, economic, man-made, and institutional.

The Sciacca pilot demonstrates how multidimensional indicators can capture not only tourism-related outcomes, but also processes of community activation, heritage reuse, local entrepreneurship, socio-cultural reinvestment, and governance innovation. The findings highlight the emergence of a regenerative territorial metabolism in which cultural heritage acts as enabling infrastructure for local development. Finally, the paper discusses methodological challenges related to data availability, mixed qualitative-quantitative evidence, and the need for longitudinal monitoring, arguing for evaluation as an adaptive and participatory governance tool supporting circular and inclusive regeneration trajectories.

Keywords: Heritage-Led Regeneration · Community-Based Tourism · Multidimensional Impact Assessment

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 506–522, 2027.
https://doi.org/10.1007/978-3-032-30518-3_32

1 Introduction

Historic small and medium-sized cities, as well as peripheral and less accessible territories, are increasingly recognised as strategic laboratories for place-based development because of their capacity to combine cultural identity, social embeddedness and environmental resources within locally grounded regeneration trajectories. At the same time, many of these territories are affected by demographic decline, marginal accessibility, fragile service systems, underused heritage assets and limited institutional capacity, conditions that often constrain their ability to transform heritage resources into long-term development opportunities [1–3]. In such contexts, the regeneration of cultural heritage cannot be reduced to conservation alone; rather, it must be linked to broader processes of community empowerment, economic diversification, social cohesion and environmental responsibility.

Within this debate, cultural tourism has increasingly been interpreted as a possible lever for territorial regeneration, particularly when it is rooted in local identity, intangible heritage, creative production and active community participation [4–7]. However, the relationship between cultural tourism and local development remains ambivalent. While heritage tourism may enhance visibility, diversify local economies and reactivate underused cultural assets, it may also generate commodification, seasonal dependence, selective valorisation and new forms of socio-spatial pressure if not supported by adequate governance and monitoring tools [8–10]. For this reason, heritage-led development requires evaluation frameworks capable of capturing both positive and negative effects across multiple dimensions and scales.

This paper contributes to this discussion by focusing on the case of Sciacca, in Sicily, where the Community Cooperative *Identità e Bellezza (Identity and Beauty)* has developed the *Museo Diffuso dei 5 Sensi (The Widespread Museum of the 5 Senses)*, an emblematic experience of community-based, heritage-led regeneration. The initiative is particularly relevant because it transforms local inhabitants from passive users into active co-producers of cultural value, tourist experiences, and care practices for abandoned or underused heritage places [11–15]. This approach resonates with the concept of Heritage Community introduced by the Faro Convention, according to which heritage value is generated and transmitted through the active involvement of people who identify with specific cultural resources and wish to sustain them for present and future generations.

The paper adopts a multicriteria and multidimensional impact assessment perspective, using the Circular Cultural Tourism Dashboard developed in the Horizon 2020 Be.CULTOUR project [16]. The main research question is the following: how can multidimensional indicators support the assessment of heritage-led, community-based regeneration processes in peripheral historic cities? A second question concerns the governance role of evaluation itself: can assessment function as a co-learning and co-decision tool for local actors, rather than merely as a technical reporting exercise?

Beyond documenting a local case, the paper contributes to the current debate in three ways. First, it operationalises the concept of Heritage Community through a structured set of multidimensional indicators, moving from a normative and interpretive notion towards an evaluative one [17]. Second, it shows how circular cultural tourism can be understood not merely as a sectoral tourism strategy, but as a regenerative territorial approach in which heritage, local entrepreneurship, social activation and care practices

are mutually reinforcing [18, 19]. Third, it tests the usefulness of a multidimensional and multicriteria impact assessment framework for interpreting complex place-based transformations in a small historic city where quantitative and qualitative evidence must be combined rather than artificially separated.

The paper argues that the Sciacca case shows the relevance of a multidimensional evaluation architecture capable of capturing cultural vibrancy, community commitment, entrepreneurship, adaptive reuse, and institutional innovation as interconnected outcomes of a circular and regenerative tourism strategy.

2 Heritage Communities, Circular Cultural Tourism and the Need for Multicriteria Assessment

The growing literature on sustainable, inclusive and regenerative tourism has highlighted the limits of conventional tourism models based predominantly on visitor growth, destination competitiveness and short-term economic returns. More recent approaches emphasise instead the importance of territorial identity, local value chains, social inclusion, environmental responsibility and transformative experiences, arguing that tourism should not simply consume places, but contribute to regenerating them [6, 7, 20, 21]. This shift is particularly relevant in peripheral and heritage-rich territories, where tourism may function not only as an economic sector, but also as a catalyst for local mobilisation, social innovation and new narratives of territorial futures [10, 22, 23].

At the same time, the heritage field has progressively moved from monument-centred conservation towards people-centred and community-centred paradigms. In this perspective, cultural heritage is increasingly understood as a living resource that enables collective memory, place attachment, creativity and local development, rather than as a static asset to be preserved in isolation. The concept of Heritage Community introduced by the Faro Convention is particularly relevant here, as it frames heritage value as something generated and transmitted through the active involvement of people who identify with specific cultural resources and wish to sustain them for present and future generations [17]. This perspective is especially useful for interpreting community-led regeneration processes in which local actors become co-producers of heritage value, tourism experiences and stewardship practices, rather than passive beneficiaries of top-down policies [24, 25].

Within this broader framework, the concept of circular cultural tourism is especially relevant because it extends circularity beyond resource efficiency alone. Circularity in tourism should be interpreted in an expanded sense, including not only the circular use of material resources but also the regeneration of social, cultural and relational assets. In this perspective, heritage reuse, local entrepreneurship, skill activation, digital mediation, community care and collaborative governance all become part of a circular territorial metabolism [18, 19]. Yet this broader understanding of circularity generates a methodological challenge: it cannot be measured through single-sector indicators alone. Existing tourism indicator systems provide important tools for sustainability monitoring, but they rarely capture the combined effects of heritage valorisation, community activation, adaptive reuse, local identity and regenerative circular development in less-known destinations [26, 27].

A multidimensional and place-based evaluation is crucial to address the complex relationship between tourism and historic centres, particularly in fragile and inner areas where standardized tourism-led enhancement processes may negatively affect local identities, community dynamics, and long-term territorial sustainability [28].

In this perspective, the evaluation of circular tourism impact should be conducted through multiple criteria and dimensions. including different meanings, weights and temporalities. Cultural vitality, social cohesion, entrepreneurship, governance capacity and environmental responsibility do not evolve in linear or fully commensurable ways; rather, they require a framework capable of recognising trade-offs, synergies and differentiated priorities [29, 30]. For heritage-led regeneration processes, evaluation should therefore support diagnosis, collective learning and governance, functioning as an embedded and adaptive component of the regeneration process itself.

3 The Case Study: Sciacca and the Community Cooperative *"Identity and Beauty"*

Sciacca is a historic coastal city located in the province of Agrigento, on the south-western coast of Sicily. Despite its rich stratification of tangible and intangible heritage, artisanal traditions, cultural landscapes, thermal resources, and maritime identity, Sciacca can still be regarded as a peripheral destination in infrastructural and tourism terms. Its heritage potential has long been only partially activated, with several places suffering from underuse, weak accessibility, and limited integration into broader local development strategies.

Within this context, the Community Cooperative *Identità e Bellezza* ("Identity and Beauty"), established in 2020, has developed an innovative model of local activation that can be regarded as one of the most significant community-based heritage regeneration experiences in Sicily. Its flagship initiative, the Museo Diffuso dei 5 Sensi ("Widespread Museum of the 5 Senses"), frames the city as an open and relational heritage system, in which local stories, crafts, food knowledge, neighbourhoods and cultural sites are reconnected through participatory experiences and distributed forms of stewardship. Residents are not external to the project; rather, they act as narrators, hosts, organisers and co-producers of value, while visitors are invited to engage with the city as temporary citizens rather than detached consumers. In this sense, Sciacca can be interpreted as a Heritage Community in action, where cultural heritage operates not only as a symbolic resource, but also as an enabling infrastructure for local development, community mobilisation and adaptive reuse [17, 22, 24, 25]. The cooperative does not merely promote attractions; it activates a social ecosystem of stewardship, in which heritage enhancement is linked to local entrepreneurship, collaborative storytelling, adaptive reuse, distributed visitor services, and community care. The project also contributes to reconnecting scattered cultural assets through a relational model instead of relying on large-scale infrastructural investment. In this sense, Sciacca offers a valuable case for examining how community-led heritage activation can produce multidimensional impacts.

4 Methodology: The Circular Cultural Tourism Dashboard as a Multidimensional Assessment Tool

The paper uses the Circular Cultural Tourism Dashboard developed within the Be.CULTOUR project as the main assessment framework. The dashboard was conceived as an operational self-assessment tool to support stakeholders and decision-makers in evaluating tourism strategies aimed at transforming less-known destinations into sustainable, circular and human-centred ones. The dashboard is structured around six dimensions: cultural, human and social, environmental, economic, man-made, and institutional.

The adoption of this framework is also justified in relation to the broader landscape of tourism and cultural indicator systems. International frameworks such as the UNWTO indicators for sustainable tourism, the UNEP-WTO guide for policy-making, the ETIS toolkit and the GSTC destination criteria have significantly advanced destination monitoring, but they tend to remain either too general, too management-oriented, or insufficiently focused on the specific interplay between heritage, circularity, community participation and regenerative local development [26, 27, 31]. At the same time, cultural monitoring systems such as the Cultural and Creative Cities Monitor provide valuable insights into cultural vibrancy and enabling conditions, yet they are not specifically designed to assess heritage-led community-based tourism transformations in less-known destinations [29]. The Be.CULTOUR dashboard is therefore used here as a more context-sensitive framework for evaluating circular cultural tourism in a peripheral heritage setting.

The choice of this framework is particularly appropriate for the Sciacca case because it allows impacts to be read not as isolated outputs but as interconnected effects distributed across multiple domains. It also accommodates mixed methods, combining surveys, checklists, mapping activities, statistical evidence, and qualitative notes. This is crucial in community-based heritage projects, where many relevant effects - such as sense of belonging, care practices, partnership density, and symbolic recognition - cannot be fully captured through quantitative metrics alone. Following this logic, the application of the framework to Sciacca can be described through four phases:

Literature-Based Framework Definition. The first phase, documented in previous research, consisted of reviewing the main tourism sustainability and cultural tourism indicator systems, identifying a gap in relation to circular, human-centred, and heritage-led tourism in peripheral areas. This phase justified the adoption of the Be.CULTOUR dashboard as a more suitable multidimensional architecture [32, 33].

Co-selection of Relevant Indicators. The co-selection phase consisted of a focus group involving three representatives of the Sciacca cooperative, who acted as key informants. Not all dashboard indicators are equally relevant to every local context. In Sciacca, the most pertinent indicators were selected through interaction with local actors and project collaborators, with the aim of focusing on those dimensions that best reflect the cooperative's actual strategy and impacts. This step is methodologically important because it transforms assessment from a top-down checklist into a context-sensitive multicriteria selection process. More specifically, indicator selection was guided by four criteria: relevance to the local regeneration strategy, data availability, sensitivity to short- and

medium-term change, and capacity to represent both tangible and intangible effects. This filtering process made it possible to prioritise indicators that were meaningful for the Sciacca pilot while preserving the multidimensional coherence of the dashboard. In this sense, the multicriteria nature of the assessment lies not only in the interpretation of results, but also in the prior selection and calibration of the indicators themselves, consistently with the broader literature on composite and multidimensional evaluation systems [29, 30]. The indicators initially proposed are those included in the Be.CULTOUR Dashboard [16]. A subset of indicators was selected and is presented in Table 1 (Sect. 5).

Co-collection of Evidence. The third phase involved collecting available data through different instruments. Some indicators were directly quantified, while others were reconstructed through descriptive evidence, proxy variables, and qualitative data. This reflects a realistic situation in local pilot contexts, where evaluation often progresses through partial but meaningful evidence rather than through fully standardised datasets. The collected dimensions refer to the engagement of multiple informants for the identification of data sources, while for some indicators more than one organization was involved in the data collection process when necessary. To collect information from multiple organizations and sources, shared datasheets were employed. Thus, the assessment should not be interpreted as a fully standardised impact measurement exercise, but as a robust exploratory framework for structured comparison, iterative monitoring and informed local decision-making, especially in contexts where meaningful change is only partially visible through conventional sectoral metrics [18, 26, 27].

Co-Evaluation and Co-reflection. Finally, the results were interpreted through an initial co-evaluation perspective that could be further improved and implemented over time. The initial phase involved members of the cooperative *Identità e Bellezza* ("Identity and Beauty") through a final focus group in which the relevant, reliability, effectiveness and replicability of the impacts assessment process was discussed. This phase, if further implemented, can successfully supports collective reflection, comparison over time, benchmarking with other places or best practices, and comparison across scales between organizations micro-level action and the broader urban contexts. In this sense, evaluation becomes a key component of governance process.

5 Results: Multidimensional Impacts of the Sciacca Pilot

The Sciacca pilot provides relevant evidence of a heritage-led regeneration trajectory whose impacts are distributed across several dimensions.

The Circular Cultural Tourism Dashboard was developed based on the indicator framework structured within the Horizon 2020 Be.CULTOUR project [16]. Figure 1 shows the dimensions and criteria of the Dashboard. In particular, the data collected in the Dashboard are categorized into six dimensions. Based on the selection criteria mentioned above, the representatives of the Sciacca cooperative identified the indicators considered most relevant for the activities carried out in the city of Sciacca. Beyond their role as monitoring and ex post impact assessment tools, the indicators can also support impact planning through the Theory of Change (ToC), as discussed in previous studies [16, 32].

In this case, since the analysis represents an ex-post evaluation exercise, it mainly assumes a monitoring and overall assessment function, aimed at supporting the revision of objectives and strategies towards a circular cultural tourism model in the city of Sciacca. More specifically, the assessment process helps identify positive results and effective actions, as well as less effective actions that may require revision during implementation. In this sense, evaluation triggers and enables a critical assessment of results and therefore supports strategic planning and decision-making processes towards a fully circular cultural tourism economy in the case study.

Moreover, a clear understanding of actual results and impacts by the involved communities enables a less subjective and biased self-assessment process, potentially supporting more informed decisions, particularly within participatory arenas. Clearly, impact assessment processes and tools alone are not sufficient to ensure successful participatory decision-making, which depends on multiple factors, including the formation and consolidation of the heritage community itself. However, the lack of evaluation capacity supported by coherent and reliable data may seriously undermine the effectiveness of participatory processes [34].

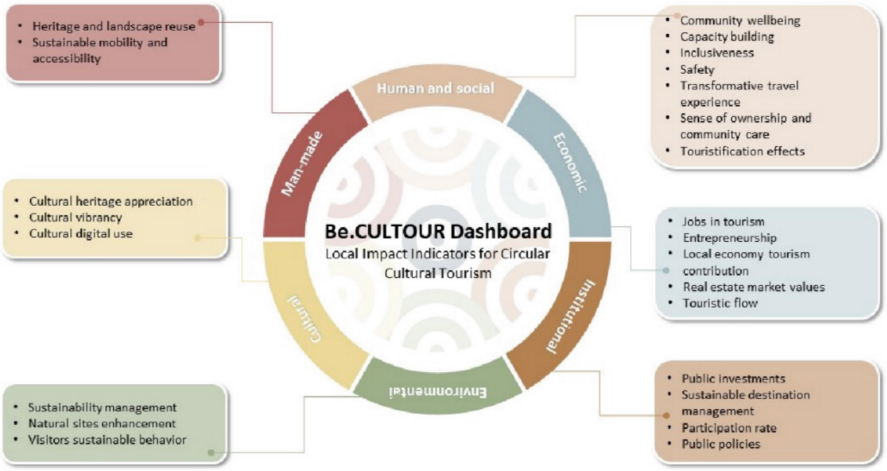

Fig. 1. Dimensions and criteria of the evaluation framework.

Based on the Dashboard, the relevant indicators were selected as described above. The list of selected indicators is reported in Table 1. Data are referred to the 2023.

According to indicators selected, a shared datasheet was established, enabling the members of cooperative to fill in the relevant data, highlighting data sources trough hyperlinks or other references. If needed, a specific section with general "notes" was provided, allowing further explanation or qualitative information. The resulting data set was used as the base for the assessment of impacts.

5.1 Cultural Impacts

The cultural dimension reveals a significant enhancement of local heritage attractiveness and cultural vibrancy. The dataset records 9 international awards and labels obtained over the last decade, indicating strong external recognition of the local heritage-based

Table 1. Selected multidimensional indicators of the evaluation framework.

Dimension	Criteria	Indicators	Units of Measure	Data collected
———Cultural dimension	Cultural heritage appreciation	Number of international awards/labels recognised to local tangible and intangible cultural heritage in the past 10 years	No. in the past 10 years	9
	——Cultural vibrancy	Number of cultural facilities open and accessible	No.	17
		Number of people visiting cultural sites per year	No./year	6.600
		Number of people participating in cultural experiences per year	No./year	922
		Number of cultural experiences for the enhancement of local tangible and intangible cultural heritage per year	No./year	45
		Number of cultural experiences for the enhancement of local tangible and intangible cultural heritage per year	No./year	4.628
Environmental dimension	Natural sites enhancement	Number of experiences to reconnect people and nature per year	No./year	4

(*continued*)

Table 1. (*continued*)

Dimension	Criteria	Indicators	Units of Measure	Data collected
——Human and Social dimension	—Sense of ownership and community care	Number of partnership activated for the care and valorizzation of cultural and natural heritage per year	No./year	5
		Volume of donations received for the management, enhancement, preservation and valorization on tangible and intangible cultural heritage per year	€/year	20.000
		Number of partnership activated for the care and valorizzation of cultural and natural heritage per year	No./year	65
	–Inclusiveness	Number of accommodations, culture and tourism structures offering supporting tools and services to breaking down language barriers	No.	9 info points with audio guides
		Number of women, youth, people with special needs and minority cultures included in enterprises	No./years	8

(*continued*)

Table 1. (*continued*)

Dimension	Criteria	Indicators	Units of Measure	Data collected
Economic dimension	Entrepreneurship	Number of new touristic enterprises in the last three year	No./3 years	5
Man-made dimension	Heritage and landscape reuse	Number of heritage/historic sites reused for cultural and tourist activities in the last five years	No. in 5 year	Via Perollo Grotte Vaporose
Institutional dimension	Public policies	Presence of policies to regulate the overcrowding	Yes/No	Yes

strategy. The city counts 17 cultural facilities, with the cooperative directly managing or activating a subset of heritage sites over time. The dataset also reports 6,600 visitors to cultural sites per year, 922 participants in cultural experiences in 2023 compared with 792 in 2022, and 45 cultural experiences dedicated to the enhancement of tangible and intangible heritage. These include food heritage workshops, craft experiences, wine and olive oil activities, carnival-related initiatives, coral traditions, and other experiential formats rooted in local identity.

These figures suggest that the cooperative's strategy is not centred on a single flagship monument, but on the creation of a distributed cultural ecosystem. This is consistent with the very idea of a diffuse museum and with the Heritage Community perspective, where value is generated through participation, interpretation, and local transmission rather than through passive consumption of heritage.

The cultural dimension also includes digital mediation. The dataset reports 4,628 downloads/views of digital audio-guide contents in 2023, after 8,788 in 2022. Rather than reading this only as a decline, it may be interpreted as a signal that digital tools need to be monitored longitudinally and contextualized in relation to changing visitor profiles, platform dynamics, and hybrid physical-digital experience design. Even so, the existence of 16 audio-guides confirms the relevance of digital infrastructure in expanding accessibility and narrative reach.

5.2 Human and Social Impacts

The human and social dimension is perhaps the most distinctive in the Sciacca case. The project is explicitly based on the activation of residents as narrators, hosts, volunteers, craftspeople, and partners. Although some social indicators still lack complete numerical standardisation, the available evidence clearly documents processes of community commitment, inclusion, and social learning.

A particularly emblematic case is the regeneration of Via Perollo, described in the dataset as a formerly marginal and weakly frequented street in the historic centre, transformed through ceramic installations, plants, and collective care practices. The notes report that the intervention involved 27 ceramic pots, 57 ceramic tiles, and 9 ceramic craftsman from the cooperative. Also, the additional qualitative notes highlight some little but relevant social outcomes: residents now water the plants, keep the street clean, and show a stronger sense of ownership, for example by contributing to the beaty of the street with additional ceramics provided by the community. This is a classic example of how a heritage-led micro-intervention can reactivate community stewardship and transform symbolic perception of place.

The dataset also reports 5 partnerships activated in 2023 within the relevant indicator, while the notes mention a much broader network including public-private agreements, educational programmes, school pathways, and internships. Likewise, donation-related indicators show a substantial capacity for collective mobilisation, with €20,000 reported under the volume of donations, alongside crowdfunding and sponsorship experiences. These elements confirm that social capital is not a background condition in Sciacca; it is one of the key outputs of the regeneration process.

5.3 Economic Impacts

The economic dimension indicates that the project is contributing to local entrepreneurship and heritage-based value creation. The dataset records 5 new tourism enterprises in the last three years, including language mediation services. Data referred to Via Perollo describes local business effects, with a reported +30% turnover for a tobacco shop and + 20% turnover for a ceramic workshop after the regeneration of the street.

While these figures are localized and not yet sufficient for a full economic impact model, they are highly relevant from a multicriteria perspective because they show how heritage enhancement can generate micro-economic spillovers through improved attractiveness, public realm quality, and storytelling value. The notes also mention increased real estate value and the emergence of new holiday homes, which suggests the need for future monitoring of both positive revitalisation effects and possible early touristification dynamics.

5.4 Built Environment and Spatial Impacts

The man-made dimension is captured above all through adaptive reuse and public space reactivation. In the dataset, Via Perollo again emerges as a paradigmatic example of heritage reuse for cultural and tourist purposes. Its transformation has produced a new urban scene, one of the most photographed places in town, and a venue for cultural and theatrical itineraries. The importance of this result lies not only in the physical upgrade of space, but in the fact that adaptive reuse here operates as a spatial mediator of social and economic regeneration.

5.5 Environmental Impacts

Although environmental evidence is less systematically available than cultural, social and economic data, the Sciacca case still suggests some relevant environmental implications of heritage-led regeneration. First, the distributed and proximity-based nature of the Museo Diffuso dei 5 Sensi ("Widespread Museum of the 5 Senses") promotes forms of slow, walkable and experience-based tourism that are potentially less resource-intensive than conventional destination consumption models, especially in a historic coastal city where cultural and everyday spaces are closely intertwined [35]. Second, the adaptive reuse of existing cultural assets and public spaces reflects a regenerative logic based on the valorisation of what already exists, thereby limiting the need for new material consumption and contributing indirectly to more sustainable place management [19]. Third, community-led actions of care and beautification, such as those observed in regenerated public spaces, contribute to a more liveable and environmentally attentive urban environment. At the same time, this dimension remains underdeveloped in the current assessment and points to the need for future monitoring through more specific indicators, such as sustainable mobility patterns, waste reduction practices, energy performance of reused spaces and environmental awareness generated through tourism experiences.

Among the six dimensions, the environmental one currently remains the least empirically consolidated and should therefore be interpreted as a preliminary field for future indicator refinement rather than as a fully measured impact domain.

5.6 Institutional Impacts

The institutional dimension highlights the role of governance arrangements and management tools. The data collected indicates the presence of mechanisms for regulating tourist flows, including audio-guides, reservation systems and ad hoc schedules, suggesting a proactive attempt to avoid congestion and manage visitor experience. This is particularly relevant because it demonstrates that even in a small historic city, heritage-led tourism can be accompanied by elementary but meaningful governance tools rather than left to spontaneous growth alone.

6 Discussion

The Sciacca case offers at least three broader insights for the debate on multicriteria impact assessment of heritage-led regeneration. These findings contribute to the literature by showing that, in small and medium-sized historic cities, heritage-led regeneration should not be evaluated through isolated sectoral indicators alone. Rather, it should be understood as a place-based transition process in which heritage operates as a connective infrastructure linking cultural production, social mobilisation, economic experimentation and institutional learning.

First, the case highlights the interdependence among impact dimensions. Cultural vitality supports entrepreneurship; adaptive reuse strengthens social ownership; digital tools expand access; partnerships reinforce governance capacity. A single-dimension indicators approach would fail to capture these relationships. The multidimensional and

multicriteria structure of the dashboard is therefore not only methodologically useful but conceptually necessary. In this respect, the Sciacca case confirms several insights emerging from recent studies on inclusive tourism, collective tourism entrepreneurship and circular cultural tourism, while also extending them by demonstrating how these dynamics can be operationalised through a multidimensional evaluation framework in a peripheral coastal city [18, 22, 25].

Second, the case confirms that impacts evaluation should be participatory and adaptive. In community-led projects, many relevant impacts are emergent, relational, and qualitative. Co-selection and co-interpretation of indicators increase the relevance of the framework and make assessment itself part of the regeneration process. In this sense, evaluation becomes a form of co-governance, fully coherent with the Heritage Community approach.

Third, the evidence from Sciacca suggests that circularity in heritage-led tourism should be interpreted in expanded territorial terms. The most significant effects do not concern only waste reduction or resource efficiency, but also the circulation of stories, trust, skills, care practices, and local economic benefits. This broader understanding of circularity is particularly suitable for historic villages, coastal towns and rural territories where intangible assets and community relations are central development resources.

The Sciacca case also suggests that experiential and digitally mediated heritage interpretation, for example through digital audio guides and itineraries, may strengthen the quality and memorability of visitor engagement while broadening access to distributed cultural assets. In line with the literature on visitor experience and human-centred tourism, the combination of storytelling, local participation and digital mediation can enhance not only satisfaction, but also emotional connection and sense-making, thereby reinforcing the regenerative potential of heritage-led tourism [20, 21, 36].

The multicriteria value of the framework lies in its capacity to interpret uneven evidence across dimensions, not by reducing them to a single aggregate score, but by supporting the comparative reading of dimensions with different levels of visibility, relevance and maturity.

More specifically, the Sciacca case both confirms and advances current scholarships. It confirms the importance of community-based tourism and collective entrepreneurship as mechanisms for mobilising local resources, enhancing social participation and generating more inclusive forms of territorial development [22, 25]. At the same time, it advances this literature by showing that, in heritage-rich small and medium-sized cities, these dynamics can be embedded within a more explicitly heritage-led regeneration logic, in which cultural heritage acts as a connective infrastructure linking identity, stewardship, adaptive reuse and local economic experimentation. Finally, the case contributes methodologically by showing that evaluation can play a role that goes beyond impact measurement alone. When organised through a multidimensional and multicriteria framework, assessment can support co-governance, reveal uneven and relational effects across domains, and help orient future decisions in complex place-based regeneration processes [18].

Overall, the results confirm that Sciacca's regeneration process is not reducible to a tourism promotion initiative. It is better understood as a multi-actor territorial strategy in which heritage functions as a platform for community activation, distributed entrepreneurship, place care, and governance innovation.

Taken together, these findings show that the strongest impacts of the Sciacca pilot emerge in the cultural and human-social dimensions, followed by significant effects in the economic, spatial and institutional spheres. The environmental dimension is currently less measurable, yet it is conceptually embedded in the project's regenerative logic through adaptive reuse, proximity-based tourism and collective care practices. This uneven distribution of evidence further confirms the usefulness of a multicriteria reading, since different dimensions display different levels of maturity, visibility and measurability.

At the same time, the case also reveals some limitations. Not all indicators are complete or directly measurable; several depend on surveys or local collaboration; some results rely on qualitative data rather than standardised time series. In the case of Sciacca, specific surveys were designed but not yet implemented thus this limited the selection of some indicators, especially those referring to citizens and tourists' perception and satisfaction. Moreover, the available evidence is stronger on activation outputs and proximate effects than on long-term outcomes such as demographic stabilization, housing affordability, or structural employment transformation. For this reason, the framework should be understood not as a closed measurement system, but as an adaptive evaluative architecture open to refinement, comparison and progressive data consolidation. Moreover, the interpretation of qualitative but also quantitative data, can be influenced by subjective bias, which was not possible to assess in the case of Sciacca as relevant benchmarks were not available for the specified indicators. The continue use of the proposed indicators framework could however address this limitation through future research and application in other case studies, collecting more information over time and possibly comparing it through available datasets and official sources where possible.

Finally, the transferability of the Sciacca experience to other contexts should be considered with caution. The effectiveness of community-based heritage regeneration depends on local leadership, social capital, heritage density, place identity and institutional openness, conditions that cannot be mechanically replicated elsewhere [10]. Nevertheless, the evaluative architecture proposed here may be adapted to other historic villages, coastal heritage cities and rural territories as a flexible framework for context-sensitive comparison, especially where conventional tourism metrics fail to capture the relational and regenerative nature of local development processes.

7 Conclusions

This paper has examined the case of Sciacca as an example of heritage-led regeneration driven by a local Heritage Community and assessed through a multidimensional multicriteria framework. The experience of the Community Cooperative *Identità e Bellezza* ("Identity and Beauty") and the *Museo Diffuso dei 5 Sensi* ("Widespread Museum of the 5 Senses") demonstrates how cultural heritage can become the nucleus of a wider

regenerative process, capable of activating local actors, reusing underused spaces, diversifying tourism experiences, strengthening territorial identity, and generating economic and social value.

The adoption of the Circular Cultural Tourism Dashboard proved especially useful because it enabled the interpretation of impacts across six dimensions and supported a mixed-method approach combining quantitative indicators and qualitative evidence. The results show that the Sciacca pilot has generated visible advances in cultural vibrancy, community engagement, partnerships, local entrepreneurship, adaptive reuse, and heritage recognition. At the same time, the case confirms that assessment must remain open, participatory and longitudinal, especially when dealing with complex community-based transformations.

From a policy and research perspective, the main implication is that evaluation should be embedded into regeneration processes from the outset. In peripheral historic cities, where resources are limited and change is often incremental, multidimensional indicators can help actors identify where value is being created, where risks are emerging, and how strategies may be adjusted over time. Rather than serving only as ex-post verification, multicriteria impact assessment can function as a strategic device for guiding circular, inclusive and heritage-led local development.

In this respect, the Sciacca case suggests that multidimensional and multicriteria assessment can play a strategic role not only in measuring change, but also in enabling it. When embedded into local decision-making, evaluation becomes a governance device capable of making visible the relational, place-specific and often intangible effects of heritage-led regeneration. This is precisely the kind of evaluative approach needed for the circular regeneration of historic villages, landscapes and rural territories, where development outcomes are distributed across cultural, social, economic, spatial, environmental and institutional domains and cannot be adequately understood through conventional sectoral metrics alone [18, 19, 27].

Acknowledgments. A.G. developed the overall research concept and assessment criteria, introduction and research questions, identification of the case study, coordination of the research, discussion and conclusions, as well as funding acquisition. F.B., M.A., A.S., specifically developed the methodology, data collection and analysis, manuscript drafting and review. All co-authors co-developed the theoretical study, literature review, results section and conclusions. Authors acknowledge the precious collaboration of the Community Cooperative *Identità e Bellezza* ("Identity and Beauty") based in Sciacca, Italy, for data collection.

Funding. This research was co-funded under the framework of Horizon 2020 research project Be.CULTOUR "Beyond cultural tourism". This project has received funding from the Euro pean Union's Horizon 2020 research and innovation programme under grant agreement Number 101004627. This research was co-funded by Italian Ministry of University and Research (MUR), in the context of the Projects of Relevant National Interest—National Recovery and Resilience Plan (PRIN PNRR). Project Title: Italian historic regeneration through circular ecological heritage communities. Grant Number B53D23029490001.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Casavola, P.: A Strategy for Inner Areas in Italy: Definition, Objectives, Tools and Governance, Rome (2014)
2. Bertolini, P., Pagliacci, F.: Quality of life and territorial imbalances. A focus on italian inner and rural areas. Bio-based and applied. Economics. **6**, 183–208 (2017). https://doi.org/10.13128/BAE-18518
3. Hospers, G.-J., Reverda, N.: Managing Population Decline in Europe's Urban and Rural Areas (2015)
4. Almeida García, F., Balbuena, A., Macias, R.: Resident's attitudes towards the impacts of tourism. Tour. Manag. Perspect. **13** (2015). https://doi.org/10.1016/j.tmp.2014.11.002
5. Timothy, D.: Cultural Heritage and Tourism: an Introduction (2011). https://doi.org/10.21832/9781845411787
6. Richards, G.: Rethinking Cultural Tourism (2021). https://doi.org/10.4337/9781789905441
7. Richards, G.: Cultural tourism: a review of recent research and trends. J. Hosp. Tour. Manag. **36** (2018). https://doi.org/10.1016/j.jhtm.2018.03.005
8. Glasson Professor, J.: Oxford: a Heritage City under Pressure Visitors, Impacts and Management Responses (1994)
9. García-Hernández, M., la Calle-Vaquero, M., Yubero, C.: Cultural heritage and urban tourism: Historic City Centres under pressure. Sustainability. **9** (2017). https://doi.org/10.3390/su9081346
10. Fonseca, F., Ramos, R.: Heritage tourism in peripheral areas: development strategies and constraints. Tourism Geographies-Tour GEOGR. **14**, 1–27 (2011). https://doi.org/10.1080/14616688.2011.610147
11. Francesca, A., Lami, I.M., M.B.: Performance indicators framework to analyse factors influencing the success of six urban cultural regeneration cases. In: Carmelina, B., Calabrò, F., D.S.L. (eds.) New Metropolitan Perspectives, pp. 886–897. Springer International Publishing, Cham (2021)
12. Bullen, P.: Factors influencing the adaptive re-use of buildings. J. Eng., Des. Technol. **9**, 32–46 (2011). https://doi.org/10.1108/17260531111121459
13. Wong, L.: Adaptive Reuse: Extending the Lives of Buildings (2016). https://doi.org/10.1515/9783038213130
14. Stefano, M., De Franco, A., Vacant Buildings, B.B.M.: Distinguishing heterogeneous cases: public items versus private items; empty properties versus abandoned properties. In: Lami, I.M. (ed.) Abandoned Buildings in Contemporary Cities: Smart Conditions for Actions, pp. 9–18. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-35550-0_2
15. Ingaramo, R., Lami, I.M., Robiglio, M.: How to activate the value in existing stocks through adaptive reuse: an incremental architecture strategy. Sustainability. **14** (2022). https://doi.org/10.3390/su14095514
16. Be.CULTOUR Project: Design of systemic policy-oriented Be.CULTOUR policy dashboard
17. Council of Europe: Framework Convention on the Value of Cultural Heritage for Society (Faro Convention)
18. Nocca, F., Bosone, M., De Toro, P., Fusco Girard, L.: Towards the human circular tourism: recommendations, actions, and multidimensional indicators for the tourist category. Sustainability. **15**, 1845 (2023). https://doi.org/10.3390/su15031845
19. Fusco Girard, L., Nocca, F.: From linear to circular tourism 1. AESTIMUM. **70**, 51–74 (2017). https://doi.org/10.13128/Aestimum-21081

20. Prayag, G., Hosany, S., Muskat, B., Del Chiappa, G.: Understanding the relationships between tourists' emotional experiences, perceived overall image, satisfaction, and intention to recommend. J. Travel Res. **56**, 41–54 (2017). https://doi.org/10.1177/0047287515620567
21. Packer, J., Ballantyne, R.: Conceptualizing the visitor experience: a review of literature and development of a multifaceted model. Visitor Stud. **19**, 128–143 (2016). https://doi.org/10.1080/10645578.2016.1144023
22. Scheyvens, R., Biddulph, R.: Inclusive tourism development. Tour. Geogr. **20**, 1–21 (2017). https://doi.org/10.1080/14616688.2017.1381985
23. Garau, C.: Perspectives on cultural and sustainable rural tourism in a smart region: the case study of Marmilla in Sardinia (Italy). Sustainability. **7**, 6412–6434 (2015). https://doi.org/10.3390/su7066412
24. Boussaa, D., Madandola, M.: Cultural heritage tourism and urban regeneration: the case of fez Medina in Morocco. Front. Architect. Res. **13**, 1228–1248 (2024). https://doi.org/10.1016/j.foar.2024.04.008
25. Jørgensen, M., Hansen, A., Sørensen, F., Fuglsang, L., Sundbo, J., Jensen, J.: Collective tourism social entrepreneurship: a means for community mobilization and social transformation. Ann. Tour. Res. **88** (2021). https://doi.org/10.1016/j.annals.2021.103171
26. Indicators of sustainable development for tourism destinations a guidebook (English version). World Tourism Organ. (UNWTO) (2004). https://doi.org/10.18111/9789284407262
27. Office, P.: The European Tourism Indicator System ETIS toolkit for sustainable destination management (2016). https://doi.org/10.2873/982144
28. Torrieri, F., Rossitti, M., Oteri, A.: The Role of Architectural Heritage in the National Strategy for Inner Areas: Evidence from the Project Areas (2024). https://doi.org/10.13138/2039-2362/3398
29. Valentina, M.: The Cultural and Creative Cities Monitor 2019th (eds.) Publications Office of the European Union (2019)
30. Gómez Vega, M., Picazo-Tadeo, A.: Ranking world tourist destinations with a composite indicator of competitiveness: to weigh or not to weigh? Tour. Manag. **72**, 281–291 (2018). https://doi.org/10.1016/j.tourman.2018.11.006
31. GSTC: www.gstcouncil.org GSTC Destination Criteria Performance indicators and SDGs (2019)
32. La Rocca, L., et al.: Towards a Circular Cultural Tourism Impact Assessment Framework for Decision Support in Less-Known and Remote Destinations. Presented at the March (2024). https://doi.org/10.1007/978-3-031-65537-1_8.
33. Gravagnuolo, A., Marasco, A., Fasciglione, M., Xydia, S., Matei, A., Garzillo, C.: Be.CULTOUR: Beyond CULtural TOURism: human-centred innovations for sustainable and circular cultural tourism Deliverable 3.1 Protocol/Methodology for human-centred innovation in sustainable and circular cultural tourism-Type of action: IA (Innovation action)
34. Carrick, J., Bell, D., Fitzsimmons, C., Gray, T., Stewart, G.: Principles and practical criteria for effective participatory environmental planning and decision-making. J. Environ. Plan. Manag. **66**, 2854–2877 (2023). https://doi.org/10.1080/09640568.2022.2086857
35. Gössling, S., Scott, D., Hall, C.: Global trends in length of stay: implications for destination management and climate change. J. Sustain. Tour. **26**, 1–15 (2018). https://doi.org/10.1080/09669582.2018.1529771
36. Stankov, U., Gretzel, U.: Tourism 4.0 technologies and tourist experiences: a human-centered design perspective. Inf. Technol. Tour. **22**, 477–488 (2020). https://doi.org/10.1007/s40558-020-00186-y

Strategic Design for Energy Self-sufficiency in Historic Villages: Renewable Energy Communities and Circular Regeneration in Grecìa Salentina

Grazia Neglia(✉), Ferdinando Verardi, and Ippolita Mecca

Engineering Department, Pegaso Telematic University, Naples, Italy
grazia.neglia@unipegaso.it

Abstract. The energy transition in rural European territories is often hindered by the structural conflict between heritage conservation and technological innovation. This study develops a systemic de-risking roadmap for establishing Renewable Energy Communities (RECs) in historic villages, utilizing a Design Science Research (DSR) approach integrated with Strategic Design and Foresight. The methodology formalizes the decision-making process by synthesizing institutional policy analysis with a comparative evaluation of radical and incremental social innovation models. The research proposes an operational roadmap structured into three strategic actions: the establishment of multidisciplinary "Transition Teams" (Cognitive Interface), the design of circular financing models to mitigate energy poverty, and the implementation of multidimensional evaluation protocols supported by Digital Twins for monitoring territorial capital. Results demonstrate that this strategic engine facilitates the transition of historic villages from passive assets toward adaptive regenerative ecosystems, effectively mediating between decarbonization mandates and the preservation of local identity. By defining the parameters for the Grecìa Salentina (Italy), the study provides a replicable roadmap that establishes the mandatory pre-implementation conditions for grounding technical infrastructures within systemic territorial coherence.

Keywords: Circular Rural Regeneration · Renewable Energy Communities · Strategic Design

1 Introduction

Rural resilience in Europe is challenged by metabolic and demographic contractions [1–3]. The RED III Directive mandates an accelerated energy transition, positioning Renewable Energy Communities (RECs) as a mechanism to decouple the built environment from carbon dependence [4]. This transition must align with New European Bauhaus (NEB) principles, reconciling sustainability with aesthetic and inclusive values to prevent the erosion of territorial identity [5–7], while leverage structured accounting systems for material cultural heritage to map territorial values [8]. This regenerative perspective is particularly relevant for terraced landscapes and rural territories, where circular economy models act as catalysts for eco-systemic recovery [9, 10].

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 523–540, 2027.
https://doi.org/10.1007/978-3-032-30518-3_33

The transition requires engaged communities as determinant factors for system adaptability [1, 11, 12]. Within the Italian context, data reveal a persistent systemic gap in environmental performance in the Southern "Mezzogiorno", where infrastructure scarcity hampers sustainability transitions [2, 3, 13, 14]. While 55% of Italian RECs are in small municipalities, development in historic contexts is hindered by administrative complexity [15–17], requiring circular business, financial, and governance models tailored for adaptive reuse [18–20].

The limited efficacy of many energy policies in rural contexts reflects a misalignment between centralized objectives and local "daily practice". Applying "human-centered engineering", this discrepancy is identified as a "gulf of execution" where the absence of integrated planning tools prevents a coherent mapping between regulatory intent and operational reality [21, 22]. From a design-science perspective, this represents a lack of "procedural rationality", necessitating a transition toward participatory and co-evolutionary evaluation processes [23–25]. Rural resilience is defined as the capacity of territorial systems to absorb external shocks [26]. According to recent assessments, this involves recognizing the "social well-being" derived from heritage as a primary indicator of territorial health [27], necessitating foresight-driven strategies to anticipate future crises [28, 29]. In vulnerable southern contexts, identifying the socio-economic drivers of resilience is vital [2, 30]. This necessitates a human-centered approach to energy justice to ensure the transition becomes a pathway toward inclusive growth [22]. Within this framework, the circular economy functions as a restorative metabolic model [31, 32], where simultaneously the intrinsic value of cultural heritage drives human-centered adaptive reuse [18, 19]. This involves moving from abandoned assets toward active nodes in the urban metabolism. Consistent with the Faro Convention, the "heritage community" emerges as the central actor [33]. Effective implementation of these models relies on understanding best practices for adaptive reuse [34], reinterpreting the village as a circular ecosystem where "circular governance" acts as the bridge between metabolic circularity and cultural identity [1, 2, 24, 35]. Technical robustness and physical safety during adaptive transformations are further ensured by integrating optimization-based structural strengthening frameworks for historic masonry buildings [36].

The Grecìa Salentina (Apulia, Italy) is a paradigmatic case. Beyond the impact of *Xylella fastidiosa*, which devastated the olive-growing landscape, the region suffers from an erosion of identity [3, 37]. The response must integrate multi-criteria decision-making to manage trade-offs between agricultural recovery and energy transition [37], positioning the REC as a hub for energy and bio-resource circularity [38, 39]. The circular metabolism can be reinforced by valorizing local agricultural residues for renewable energy [40, 41]. In line with the Historic Urban Landscape (HUL) approach, this territory represents a "vulnerable node" where the loss of traditional agricultural metabolism necessitates a design-driven recovery of territorial capital [2, 7, 42, 43]. This study evaluates the compatibility between technological requirements and the intrinsic value of built heritage, proposing systemic nodes within the territorial metabolism [1, 39, 44, 45].

1.1 Paper Objective and Structure

The study proposes a decision-support model grounded in the convergence of Design Science Research (DSR), Strategic Design, and Foresight methodologies [23, 28, 46]. The output is a socio-technical artifact designed to mitigate the "cognitive friction" and "administrative barriers" characterizing the implementation of energy infrastructures in protected contexts [17, 21]. Strategic Design functions here as an integrative field where experts facilitate "social conversations" to ground European initiatives within local circular programs [6, 47]. Within this architecture, Strategic Design provides the "kernel theory" for integrating organizational strategy with design-led innovation [46], while Foresight acts as the protocol to ensure resilience against future shocks by integrating inbound and outbound change [28]. RECs serve as instruments for the shared management of cultural heritage [39], and for valorizing public heritage "awaiting reuse" [44]. The analysis investigates the operational translation of circular metabolic principles (RQ1), the role of energy citizenship (RQ2), the capacity of RECs to mitigate systemic risks (RQ3), and the optimal governance/financing models (RQ4). The latter is crucial as financial viability remains a primary obstacle, requiring innovative capital structures and the involvement of "heritage communities" [48–50]. This pre-implementation roadmap acknowledges the reality of social innovation in marginalized areas, where expectations often clash with structural constraints [51, 52]. By prioritizing the governance engine, the research seeks to bridge the gap between European policy goals (RED III [4], NEB [5]) and local identity [23, 53].

The document is organized as follows: Sect. 2 explores the theoretical and regulatory framework concerning rural regeneration and RECs; Sect. 3 presents the research methodology; Sect. 4 presents the results and discussion; finally, Sect. 5 outlines conclusions and future research perspectives.

2 Literature Review

2.1 Circular Metabolism and Adaptive Territorial Networks

The construction sector is a critical pressure point for environmental stability, accounting for 40% of energy consumption and 36% of greenhouse gas emissions [31]. This necessitates a transition toward circular regenerative loops that prioritize the maintenance of existing value. In the European landscape, where over 80% of the building stock predates 1990, circular regeneration is a structural necessity for territorial resilience [1]. In historic contexts, this transition requires a multi-scalar understanding of territorial capital where cultural heritage acts as a non-renewable resource [2, 18, 38]. This implies a shift toward the Circular City model, where the reuse of the built environment is a core component of territorial governance [19, 35], evolving toward a "human-centered" village model capable of regenerating both material and non-material values [24].

The historic village is re-evaluated as a vulnerable node within territorial networks, where environmental resilience can be interpreted through the lens of small-scale aggregates' ability to self-organize resources [54]. This systemic integrity is maintained through dynamic co-evolutionary processes [1, 2, 18], requiring structured transformation management [55] and rigorous methodological frameworks for heritage accounting [8]. However, reconciling modern performance with heritage constraints remains a

"metabolic gap" [1, 38]. Closing this gap requires evaluating best practices in adaptive reuse to identify scalable circular governance models [34]. In the Italian context, energy efficiency must be balanced with the preservation of cultural-identity features [3, 7]. To ground these interventions without compromising structural safety, optimization-based frameworks for the strengthening of masonry buildings must be cross-referenced with energy retrofit scenarios [36]. To measure this transition, Human-Centered Indicators (HCI) are adopted to evaluate the regeneration of vulnerable heritage in terms of social and human well-being [56], ensuring that circularity integrates "beauty" and "sustainability" in line with the NEB [5]. The adaptive reuse paradigm reconfigures the historic village as a living organism [18]. By extending the life cycle of buildings, historic structures evolve from passive assets into active energy hubs [31, 44]. While reaching Nearly Zero-Energy Building (nZEB) standards in heritage contexts presents significant constraints, innovative technologies and bio-based materials allow for substantial CO_2 reductions while respecting historical integrity [31, 45, 57]. The circular metabolism of rural villages is further enhanced by the valorization of local agricultural residues for energy production [40, 41]. Empirical precedents in specific Italian contexts demonstrate that such circular interventions can transform abandoned assets, such as the former military ammunition depot in Sangano, into energy-active hubs for the community [44], provided that the overarching technological deployments are systematically integrated with nZEB adaptive reuse frameworks [31], and governance models are tailored for circular adaptive reuse [20].

2.2 Socio-technical Interfaces: Designing for "Heritage Citizenship"

The transition to decentralized energy via RECs is essentially a "social innovation" process [16]. However, social innovation often faces structural constraints in marginalized areas [51]. To mitigate this, Strategic Design must focus on the "affordance" of governance – the actual properties of the REC that determine the ease of interaction between the community and the system [21]. Citizen engagement is a dynamic process influenced by transparency and the visibility of benefits [58]. This engagement is a fundamental driver for implementing Circular City models, ensuring that regeneration strategies align with local needs [11]. Co-design acts as the operational tool, converting citizens into active "prosumers" [17, 57].

The "heritage community" [33] serves as the social engine, yet its activation requires a "neo-endogenous" development approach where local actors are supported by external networks without losing territorial sovereignty [52]. Activation is conditioned by the "governance of proximity" [59], necessitating participatory and co-evolutionary evaluation models to foster deep-rooted commitment [25]. In line with the New Leipzig Charter, social inclusion must be a structural requirement of the energy metabolism [60]. While Strategic Design and Foresight provide the operational protocols for long-term planning [24, 28], recent psychosocial research highlights that the communicative processes and behavioral drivers are fundamental for driving actual citizen involvement in RECs [61, 62]. Within this framework, the heritage community becomes the primary stakeholder in a "circular human-centered" governance, where cultural identity acts as the glue for collective action [18, 24].

2.3 Regulatory Gaps and "Systemic de-Risking" in Marginal Areas

The implementation of the RED III Directive [4] and the Italian regulatory framework (D.Lgs 199/2021 [63]; D.M. 414/2023 [64]) occurs amid regional disparities [15]. A "place-based" deficiency – applying uniform solutions to heterogeneous territories – often results in social discontent in the "Mezzogiorno" [65]. This is fueled by the "digital and energy divide" affecting inner areas [13, 66]. Addressing these disparities requires a holistic view of the circular economy supported by a clear taxonomy and mapping of existing energy community models across the European landscape [67]. Integrating the UNDRR GAR 2025 framework shifts the focus to "systemic de-risking", where energy self-sufficiency acts as insurance against cascading climate risks [26]. This is inherently linked to "energy justice" [22].

The regeneration of specific landscape features, such as terraced systems, demonstrates how circularity mitigates systemic risks while providing economic value [9, 10]. Without multi-level governance that integrates NEB and Faro principles, RECs risk technical alienation [51, 58]. The National Strategy for Inner Areas Plan (PSNAI) acknowledges this vulnerability but requires tools for "sovereign resilience" [68]. Sustainability is ultimately linked to circular well-being, where proceeds are reinvested into the maintenance of the Historic Urban Landscape [42] and monitored through multidimensional indicators [56]. Evaluating the attractiveness of these nodes remains crucial for guiding investment and identifying development potential [30]. Financial viability depends on leveraging regional funding and operational support from the Gestore dei Servizi Energetici (GSE) while managing the constraints of social innovation [48, 51, 69, 70].

3 Methodology

The research adopts a Design Science Research (DSR) approach, a paradigm engineered to extend organizational capabilities by creating innovative socio-technical artifacts [23]. The study operates at the "front-end of innovation," focusing on the architectural configuration of the decision-making process rather than the empirical monitoring of existing infrastructures [28, 46]. The robustness of this approach is ensured by the recursive interaction between three cycles – relevance, rigor, and design – which establish the epistemic conditions for a replicable roadmap in fragile contexts (Fig. 1).

3.1 The Relevance Cycle: Contextualizing the "Wicked Problem"

The Relevance cycle bridges the research with the specific environment of the Grecìa Salentina. This phase serves as a diagnostic protocol to identify the requirements of the artifact.

Problem Framing: The study identifies the "metabolic friction" caused by post-ecological shocks and regulatory complexity. This justifies a circular governance model [24, 47, 56] that reintegrates abandoned biomass and marginal lands into the local energy cycle [37, 38]. The empirical feasibility of this approach is backed by recent technological assessments on the bioenergetic potential of local lignocellulosic residues and agricultural biomass [40, 41].

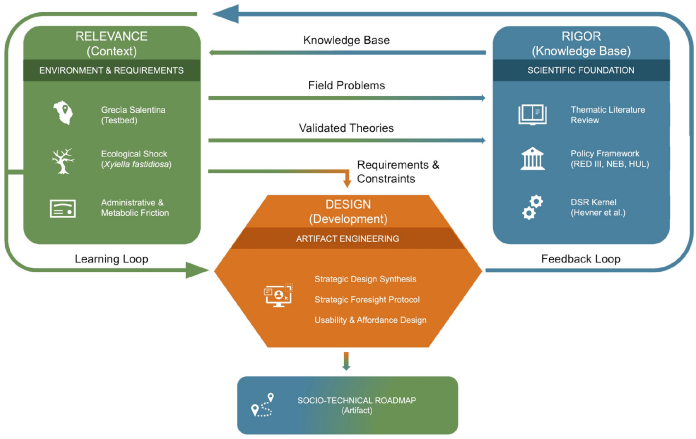

Fig. 1. Methodological framework based on DSR cycles [23] (Authors' elaboration)

Territorial Requirements: Utilizing institutional grey literature and regional policies, the research defines the socio-technical constraints [65]. The assessment of urban attractiveness in Southern contexts serves as a benchmark for identifying these requirements [30], ensuring the roadmap addresses the specificities of the "Mezzogiorno".

3.2 The Rigor Cycle: Literature as Foundation Building

The Rigor cycle provides scientific grounding through a purposive thematic review (2013–2026), structured as a "foundation building" process to justify the pre-implementation stage.

Data Collection Logic: A hybrid documentary base was synthesized by cross-referencing peer-reviewed literature (Scopus/Google Scholar) with multi-level policy frameworks, including the RED III Directive, New European Bauhaus, and UNESCO Historic Urban Landscape recommendations to ensure "convergent validity" [58]. This involves extracting best practices for the adaptive reuse of cultural heritage to identify successful circular business and governance models [20, 34]. This ensures that the roadmap's strategic actions (Sect. 4.3) are consistent with international standards of conservation and participation [58].

Strategic Scoping: The review utilized a combination of Boolean operators and keywords (e.g., "Energy Self-sufficiency", "Rural areas", "Historic Villages", "Renewable Energy Communities", "Circular Economy", "Rural Regeneration", "Place-based", "Community transition", "Energy transition", "Grecìa Salentina", "Apulia", "Regional Policy"), to capture the rapid evolution of the REC regulatory landscape and the socio-economic dynamics of heritage intrinsic value [24]. To ensure comprehensive boundary conditions, the artifact acknowledges the physical preservation of the built fabric, recognizing the role of structural strengthening frameworks for masonry buildings in ensuring long-term safety [36].

3.3 The Design Cycle: Strategic Design and Usability Interface

The core of the methodology lies in the Design cycle, where the roadmap is engineered not as a passive document, but as an "active usability interface" for territorial governance [46].

Strategic Foresight Protocol: Foresight is employed to "rehearse the future," integrating global trends with local identity [46]. This structure is grounded in the multidimensional benefits of landscape regeneration as a driver for ecosystemic and economic recovery [9, 10], transforming the roadmap into a systemic de-risking tool against the cascading risks identified in the UNDRR GAR 2025 framework [26].

Systemic Usability Engineering: The roadmap integrates principles of human-centered engineering to mitigate the "design failure" common in top-down policies through [21].

Functional Affordance: Action A (Transition Team) suggests and facilitates actionable paths for local administrators.

Signifiers and Cognitive Mapping: The roadmap guides the "heritage community", aiming to reduce the "cognitive friction" between regulatory complexity and local agency.

Strategic Feedback Loops: Action C (Digital Twin) is hypothesized as a dynamic repository of metabolic flows, leveraging digital tools and spatial data frameworks for historic building management [31, 57], while social well-being indicators are derived from human-centered monitoring models [56]. This cycle ensures that community engagement remains the primary stakeholder in governance [11, 19].

3.4 Case Study Selection and Comparative Logic

A multi-scale comparative analysis (Sect. 4.2) extracts key variables (governance, financing, participation) for roadmap replicability. The Grecìa Salentina district was selected as a critical case study (Sect. 4.1) due to the convergence of post-ecological shock conditions, geographical marginality, and institutional activism, exemplified by the European project GRowing participative models for EnErgy commuNities (GREEN) [71, 72]. This Union of Municipalities allows for testing "circular city and region models" at a sub-regional scale, where linguistic identity acts as a social glue for the "heritage community" [3, 52, 72]. The case functions as a testbed to transition from abandoned assets toward active nodes within a regenerative metabolism [32, 35], bridging theoretical premises with operational actions [2, 43, 69].

3.5 Pre-Implementation Rigor and Evaluative Boundaries

As a pre-implementation artifact, the roadmap's utility is evaluated through its internal consistency and potential to solve administrative paralysis in protected contexts [23, 50]. The absence of immediate post-operational data is a deliberate methodological boundary. The research aims to establish "epistemic prerequisites" as mandatory precursors to physical deployment, prioritizing de-risking and social consensus [28, 51]. This evaluative framework follows participatory, dynamic, and co-evolutionary models designed

to monitor the long-term impacts of circular adaptive reuse [25]. By defining structural governance and circular financial loops, the methodology provides a framework to stabilize conditions for future pilot testing [23, 26, 31].

4 Results and Discussion

4.1 Case Study: the Grecìa Salentina Context

The Union of Municipalities of Grecìa Salentina [72] represents a rural system where vernacular forms define a long-term co-evolutionary landscape [3]. The ecological shock of *Xylella fastidiosa* has converted this productive fabric into "residual capital" requiring systemic reconversion [2, 37]. This crisis necessitates a "human-centered" recovery where community engagement is the determinant factor for adaptability [58], leveraging the cultural "Grìko" identity as a catalyst for social capital regeneration [52, 72]. In this context, the Circular City model ensures that the recovery of identity coincides with economic resilience [35].

Current efforts via the GREEN Project [71], in which the Union of Municipalities of Grecìa Salentina is a strategic partner, aim to establish RECs as instruments for cohesion. However, the transition faces a "metabolic friction": integrating technology within the HUL without compromising morphological integrity [42]. The REC is designed to overcome the "gulf of execution" between regional targets and heritage constraints by adopting a "circular governance" model that reconciles ecological needs with humanistic regeneration [24].

4.2 Cross-Case Analysis and Comparative Lessons

The comparative analysis (Table 1) highlights that REC success depends on the ability to integrate heritage protection with mutualistic management [16, 49, 50]. Adaptive reuse requires specific circular business and governance models [20, 34]. Evidence suggests that social innovation in rural areas requires a "neo-endogenous" approach where local actors maintain sovereignty over resources [52].

Table 1 serves as the empirical base for deriving the key strategic variables – governance, financing, and participation – identified as essential for the replicability of the circular regeneration model.

Radical trajectories, such as Birse or San Giovanni a Teduccio, demonstrate that innovation lies in the collective ownership of resources [50, 59]. Specifically, the empirical dynamics of the Birse case, as investigated through the radical social innovation framework [59], illustrate how a community-led enterprise can actively build a new development trajectory by reclaiming ancient rights over natural heritage, despite deep-seated landownership concentrations and subsequent internal ideological conflicts regarding community empowerment. These models collectively demonstrate how the "intrinsic value" of heritage drives circular human-centered adaptive reuse [18]. Conversely, the incremental model prevalent in Italy reflects a municipal-led leadership that acts as a guarantor of institutional trust [53]. Cases like Melpignano or Ferla demonstrate that participation is often mediated by local authorities to facilitate access to PNRR funds

Table 1. Multi-scale comparative analysis (Authors' elaboration)

Case Study	Technological Mix & Heritage Constraint	Governance & Social Innovation Trajectory	Financing Model	Strategic Value & Heritage Integration	Ref.
Ferla (Italy)	Distributed PV (311 kW) & IoT monitoring. Constraint: Minor historic village integrity	Incremental: Municipal-led trajectory with prosumer involvement	TREPESL Project & regional funding	Education in "heritage citizenship" through collective management	[14]
Villanovaforru (Italy)	PV (44.3 kW), storage & smart meters. Constraint: Rural & historic landscape protection	Incremental: Consolidation of local resilience through institutional trust	PNRR (Mission 2, C2, I.1.2) & municipal budget	Common energy financing for proximity services	[16]
S. Giovanni a Teduccio (Italy)	55 kWp PV & 10 kWh storage. Constraint: Peripheral urban integrity	Radical: Solidarity model for combating energy poverty	Fondazione con il Sud and national public incentives	Social algorithm for inclusive well-being; regeneration of the urban-social fabric	[50]
Melpignano (Italy)	PV on historic roofs & water management. Constraint: Historic center integration	Incremental: Strengthening cohesion through a community cooperative	Local resources and PNRR (Mission 1)	Valorization of heritage as a living system (HUL)	[53]
Birse (Scotland)	Forest biomass & soil management. Constraint: Integrity of natural systems	Radical: Trajectory-creating process based on collective land ownership	Self-financing from community assets & local resources	Sovereignty over natural resources as a form of care for the common good	[59]

(*continued*)

Table 1. (*continued*)

Case Study	Technological Mix & Heritage Constraint	Governance & Social Innovation Trajectory	Financing Model	Strategic Value & Heritage Integration	Ref.
South Tyrol (Italy)	Hydroelectric & district heating. Constraint: Alpine landscape conservation	Incremental / Historical: Evolution of centuries-old mutualistic patterns	Social capital of cooperatives & energy sales	Energy self-sufficiency as a pillar of cultural identity	[67]
San Severo (Italy)	Onshore wind farm. Constraint: Landscape integration	Incremental: Active civic participation model	Civic crowdfunding: investment open to residents	Democratization of energy; shared productive infrastructure	[70]

[14, 53]. When expanding this perspective to the macro-regional level, the historical evolution of mutualistic patterns – such as those historically observed in alpine cooperative networks [70] – further confirms that regional long-term energy self-sufficiency can act as a structural pillar of local identity, reinforcing the broader taxonomies and mapping criteria established for European energy communities [67]. This confirms that the transition must be framed as a pathway toward energy justice, ensuring benefit distribution among vulnerable populations [22].

4.3 The Strategic Design Roadmap: Operationalizing the "Human-Centered" Interface

The proposed roadmap (Fig. 2) translates the RED III Directive [4] and New European Bauhaus [5] into an operational framework.

Action A – The Transition Team (Cognitive Interface): This unit is intended to function as a technical bridge for participatory co-design, potentially reducing the "cognitive friction" deriving from administrative complexity and regulatory burden [21, 42]. By aligning with NEB initiatives and HUL criteria, it identifies specific investment sectors where heritage and energy synergy is highest [6, 19].

Action B – Civic Crowdfunding (Circular Financial Loop): This action integrates public grants with digital crowdfunding to support "procedural justice" [22, 48]. Retaining added value locally is essential to fund the restoration of heritage assets like dry-stone walls within compromised agricultural and terraced contexts [3, 9]. Aligning with UNDRR GAR 2025 [26], this action proposes to treat community investment as an experimental asset to mitigate energy poverty while preserving cultural identity [2, 39].

Action C – Digital Twin Monitoring (Strategic Evaluation Feedback): This tool evaluates "territorial capital quality" through participatory, dynamic, and co-evolutionary

protocols [25, 28]. It transforms data into "strategic intelligence" [57], monitoring the impact on the intrinsic value of heritage and the well-being of the population [27, 56]. This ensures the REC is treated as an "active adaptive reuse project" capable of responding to post-ecological shocks [31, 37, 44].

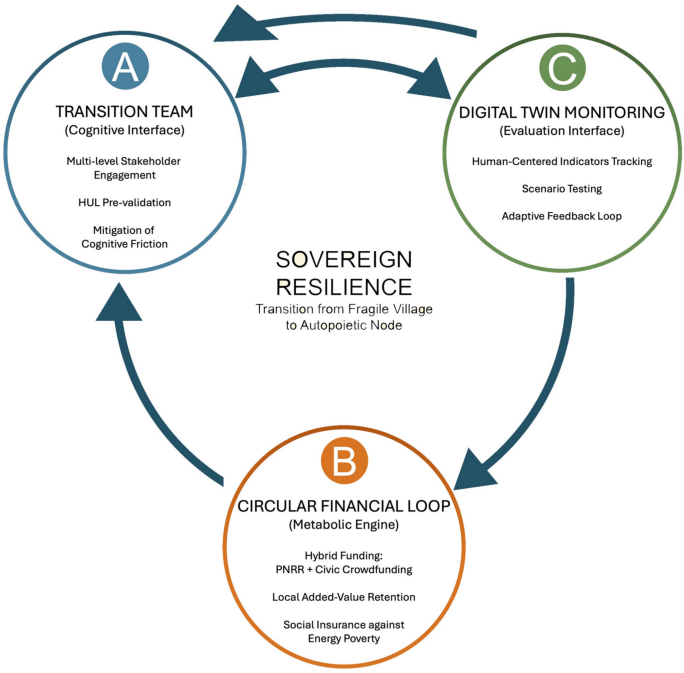

Fig. 2. Strategic Design framework for circular rural regeneration an integrated de-risking engine (Authors' elaboration)

4.4 Design Architecture for Territorial Resilience: A Systemic Synthesis

The innovation of the proposed roadmap lies in its architectural configuration as a potential multi-layer de-risking engine for fragile heritage contexts. By leveraging community engagement as the primary driver, the roadmap transforms urban regeneration into a shared management of common goods [11].

Administrative De-risking (Action A): It aims to address the conflict between conservation and innovation [45, 55]. By establishing a multidisciplinary Transition Team, the roadmap seeks to neutralize the risk of "veto-players" (e.g., *Soprintendenze*) through early-stage co-design, it hypothesizes a transition from a bureaucratic "wicked problem" into a process of human-centered engineering [12, 21], where the Faro Convention [33] and HUL criteria [42] are not constraints but design parameters. It seeks to ensure that circular governance reduces the risk of failure in top-down initiatives by fostering local consensus [24]. This pre-validation phase is intended to reduce the "time-to-market" of

energy communities in areas under landscape protection, where delays often lead to the loss of funding [17, 64, 69].

Financial and Metabolic De-risking (Action B): Community investment serves as "social insurance" against energy poverty [26, 49]. By integrating grants (e.g., PNRR) [64, 69] with crowdfunding, the roadmap attempts to solve the "place-based deficiency", proposing a financial metabolism that could prevent the "draining of value", aiming for "sovereign resilience" [52, 54, 65], where the economic surplus is recycled into the maintenance of the territorial capital [2, 7].

Adaptive and Foresight De-risking (Action C): It is intended to prevent "identity erosion" via feedback loops and Digital Twin technology [18, 28, 29]. This enables the REC to function as an adaptive socio-technical system capable of absorbing ecological shocks through simulated diagnostic scenarios [1, 50], incorporating active citizen engagement into its long-term evolution [58], monitored through Human-Centered Indicators [56]. The use of Digital Twins allows for "scenario-testing", helping the heritage community to visualize the impact of energy infrastructures before they are built, thus increasing social acceptance [57, 58].

4.5 Critical Evaluation of the Pre-implementation Roadmap

As a pre-implementation artifact, the proposed roadmap requires a balanced assessment of its potential utility and inherent limitations.

Potential Strengths (Pros): The primary hypothesized advantage is the reduction of "cognitive friction" through a structured governance interface [21]. By providing a clear roadmap before physical deployment, local administrations can coordinate the NEB pillars (inclusion, aesthetics, sustainability) [5] with the New Leipzig Charter's focus on the common good [60], justifying the human-centered approach as the core requirement for resilient system design [58, 65].

Challenges and Limitations (Cons): Effectiveness remains linked to the "social activation capacity" of the community, which may be limited in marginalized areas [51, 61, 62]. Metrics currently serve as predictive hypotheses requiring validation through longitudinal pilot testing [23].

Ultimately, while the roadmap does not claim to solve all the systemic issues of inner areas, it aims to provide a replicable operational method to potentially ensure that the energy transition is a regenerative process tailored to the specific metabolic needs of heritage-rich territories [2, 3, 7, 27, 37, 56].

5 Conclusions

The persistent marginalization of inner areas necessitates a paradigm shift from passive preservation toward proactive, circular regeneration. This research establishes a methodological bridge between European energy ambitions and the granular constraints of historic rural contexts. Adopting a Design Science Research approach, the work confirms that the viability of RECs in fragile territories depends on the architectural configuration of governance models rather than technological endowment alone [23, 24].

Addressing RQ1 and RQ2, the research demonstrates that circular metabolism functions as a socio-technical transition where community energy practices serve as primary procedural assets [2, 11, 16]. The proposed roadmap narrows the gap between policy targets and the operational reality of marginalized landscapes, identifying administrative complexity and regulatory barriers as the primary obstructions to local agency [2, 24]. To address RQ3 and RQ4, the research proposes a synthesis of three integrated de-risking mechanisms: a multidisciplinary Transition Team to navigate regulatory hurdles [1, 6]; a circular financial loop to retain economic value within the heritage community [48, 49]; and a Digital Twin-based monitoring system to safeguard historical integrity and track human-centered well-being [25, 31, 56, 57]. This tripartite structure provides a strategic response to the structural and financial constraints of social innovation in marginal territories [51, 52]. Furthermore, the Strategic Foresight protocol allows local administrations to anticipate future trajectories and absorb external shocks through an adaptive socio-technical artifact [28, 29]. This approach shifts de-risking from a purely financial dimension to a socio-epistemic one, where the usability of governance is a prerequisite for success.

While the pre-implementation scope constitutes a methodological boundary, the architecture provides a replicable protocol for both national (PSNAI [68]) and international (RED III, NEB [4, 5]) frameworks. The Grecìa Salentina case study serves as a scalable model for Mediterranean "vulnerable nodes", where RECs mitigate environmental shocks by transforming residual capital into multidimensional value [19, 37, 42, 56]. Ultimately, if validated through future longitudinal testing, this roadmap offers a pathway for fragile villages to evolve toward "sovereign resilience", where resource autonomy and cultural persistence converge into a unified trajectory of territorial regeneration.

Author Contributions. Conceptualization, G.N., F.V., and I.M.; methodology, G.N. and F.V.; formal analysis, G.N. and F.V.; investigation G.N.; data curation G.N., and F.V.; writing – original draft preparation, G.N., F.V.; writing, review and editing, G.N., F.V., and I.M.; supervision, G.N., F.V. All authors have read and agreed to the published version of the manuscript.

Funding. This research was funded by Italian Ministry of University and Research (MUR), in the context of the Projects of Relevant National Interest – National Recovery and Resilience Plan (PRIN PNRR). Project Title: Italian historic regeneration through circular ecological heritage communities. Grant Number B53D23029490001.

G.N. was supported by a Ph.D. fellowship programme within the XXXIX Cycle of Doctoral Course in "Digital Transformation", funded by the Ministero dell'Università e della Ricerca (MUR) – D.M. 117/2023, Grant No. 39–033-E1-DOT238M9ZB-9959, CUP B63C23001310003.

Data Availability Statement. Data are contained within the article.

Conflicts of Interest. The authors have no competing interests to declare that are relevant to the content of this article.

Disclosure. The authors declare that no Generative AI was used in the creation of the scientific content of this manuscript. AI Assisted Technology was used solely for grammar refinement.

References

1. Ciampa, F., Marchiano, G., Girard, L.F., Angrisano, M.: The Rural Village regeneration for the European built environment: from good practices towards a conceptual model. Sustainability. **17**, 2787 (2025). https://doi.org/10.3390/su17072787
2. Verardi, F., Angrisano, M., Fusco Girard, L.: New development policies for the internal areas of southern Italy. General principles for the valorization of rural areas in Calabria region: Nuove politiche di sviluppo per le aree interne del Mezzogiorno. Principi generali per la valorizzazione Dei borghi rurali della Regione Calabria. Valori Valutazioni. **33**, 105–116 (2023). https://doi.org/10.48264/VVSIEV-20233308
3. Barbanente, A., Grassini, L.: Landscape regeneration and place-based development in marginal areas: learning from an integrated project in southern Salento. City Territ. Archit. **11**, 26 (2024). https://doi.org/10.1186/s40410-024-00247-3
4. Directive (EU) 2023/2413 of the European Parliament and of the Council of 18 October 2023 amending Directive (EU) 2018/2001, Regulation (EU) 2018/1999 and Directive 98/70/EC as regards the promotion of energy from renewable sources, and repealing Council Directive (EU) 2015/652. Official Journal of the European Union. L series. 31.10.2023
5. European Commission: New European Bauhaus. Beautiful, Sustainable, Together. Brussels, 15.9.2021. COM (2021) 573 final (2021)
6. Angrisano, M., Gravagnuolo, A., Bottero, M., Fusco Girard, L.: Towards the implementation of new European Bauhaus initiatives in circular cities programmes: analysis of best practices to identify investment sectors. Front. Built Environ. **11**, 1601770 (2025). https://doi.org/10.3389/fbuil.2025.1601770
7. Buglione, F., Della Morte, P., Angrisano, M., Gravagnuolo, A., Fusco Girard, L.: Circular models for the sustainable regeneration of Italian rural villages: a critical analysis of good practices toward the definition of a circular Rural Village. Sustainability. **18**, 1405 (2026). https://doi.org/10.3390/su18031405
8. ESPON: HERMES–Material cultural heritage satellite account–methodological framework, Final report. February 2025. ISBN: 978–2–919839-07-0 (2025)
9. Gravagnuolo, A., Varotto, M.: Terraced landscapes regeneration in the perspective of the circular economy. Sustainability. **13**, 4347 (2021). https://doi.org/10.3390/su13084347
10. Fusco Girard, L., Gravagnuolo, A., De Rosa, F.: The multidimensional benefits of terraced landscape regeneration: an economic perspective and beyond. In: Varotto, M., Bonardi, L., Tarolli, P. (eds.) World Terraced Landscapes: History, Environment, Quality of Life, pp. 273–293. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-319-96815-5_17
11. Angrisano, M., Nocca, F.: Urban regeneration strategies for implementing the Circular City model: the key role of the community engagement. In: Gervasi, O., Murgante, B., Rocha, A.M.A.C., Garau, C., Scorza, F., Karaca, Y., Torre, C.M. (eds.) Computational Science and its Applications–ICCSA 2023 Workshops, pp. 359–376. Springer Nature, Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-37117-2_25
12. Gravagnuolo, A., Micheletti, S., Bosone, M.: A participatory approach for "circular" adaptive reuse of cultural heritage. Building a heritage Community in Salerno, Italy. Sustainability. **13**, 4812 (2021). https://doi.org/10.3390/su13094812
13. Legambiente: Ecosistema Urbano rapporto sulle performance ambientali delle città 2025. A cura di Laurenti, M., Trentin, M. In collaborazione con Ambienteitalia e il Sole 24 Ore (2025)
14. Legambiente: Comunità Energetiche Rinnovabili. Il punto della situazione in Italia (2025)
15. Zhu, Y., Salvalai, G., Zangheri, P.: Italian renewable energy communities: status and prospect development analysis. Energ. Buildings. **348**, 116404 (2025). https://doi.org/10.1016/j.enbuild.2025.116404

16. Musolino, M., Farinella, D.: Renewable energy communities as examples of civic and citizen-led practices: a comparative analysis from Italy. Land. **14**, 603 (2025). https://doi.org/10.3390/land14030603
17. Barbaro, S., Napoli, G.: Towards a participatory energy transition. Critical issues and potentials of regulatory and financial instruments for renewable energy communities (RECs) in Italy. Valori Valutazioni. **35**, 69–95 (2024). https://doi.org/10.48264/VVSIEV-20243506
18. Fusco Girard, L., Vecco, M.: The "intrinsic value" of cultural heritage as driver for circular human-centered adaptive reuse. Sustainability. **13**, 3231 (2021). https://doi.org/10.3390/su13063231
19. Angrisano, M., Bottero, M., Cavana, G., Fabbrocino, F., Gravagnuolo, A., Fusco Girard, L.: Adaptive reuse of cultural built heritage: towards the implementation of the circular city model. Front. Built Environ. **11**, 1561982 (2025). https://doi.org/10.3389/fbuil.2025.1561982
20. Fusco Girard, L., Gravagnuolo, A. (eds.): Adaptive Reuse of Cultural Heritage: Circular Business, Financial and Governance Models. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4
21. Norman, D.A.: The design of everyday things: Revised and expanded edition. Basic Books. New York (2013)
22. Tobari, T.: Human-centered approaches to energy justice: pathways toward inclusive and sustainable economic growth. IJBLE. **7**, 83–91 (2026). https://doi.org/10.56442/ijble.v7i1.1355
23. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research 1. MIS Q. **28**, 75–105 (2004)
24. Fusco Girard, L.: The evolutionary circular and human centered city: towards an ecological and humanistic "re-generation" of the current city governance. HSM. **40**, 753–775 (2021). https://doi.org/10.3233/HSM-211218
25. Gravagnuolo, A., Bosone, M., Micheletti, S., Angrisano, M., Fusco Girard, L.: Towards participatory, dynamic, co-evolutionary evaluation for circular adaptive reuse of cultural heritage: the experimentation of Salerno Circular City of health. In: Fusco Girard, L., Gravagnuolo, A. (eds.) Adaptive Reuse of Cultural Heritage, pp. 349–376. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4_13
26. UNDRR: Global Assessment Report on Risk Reduction. Resilience Pays, Financing and Investing for our Future (2025)
27. ESPON: HERIWELL–Cultural Heritage as a Source of Societal Well-being in European Regions. Final report. June 2022. ISBN: 978–2–919816-64-4 (2022)
28. Buehring, J., Bishop, P.C.: Foresight and design: new support for strategic decision making. She Ji: The J. Des. Econ. Innov. **6**, 408–432 (2020). https://doi.org/10.1016/j.sheji.2020.07.002
29. Heritage, A., Iwasaki, A., Wollentz, G.: Anticipating Future for Heritage. ICCROM Foresight Initiative. Horizon Scan Study 2021. Rome, Italy. ICCROM. ISBN: 978–92–9077-339-9 (2023)
30. Gravagnuolo, A., Angrisano, M.: Assessment of urban attractiveness of port cities in southern Italy—a case study of Torre Annunziata. Sustainability. **5**, 3906–3925 (2013). https://doi.org/10.3390/su5093906
31. Angrisano, M., Gravagnuolo, A., Cavalaglio, G., Neglia, G., Fabbrocino, F.: The circular adaptive reuse project of cultural heritage buildings in Salerno (Italy): from abandoned assets to nZEB buildings. In: Gervasi, O., Murgante, B., Garau, C., Karaca, Y., Faginas Lago, M.N., Scorza, F., Braga, A.C. (eds.) Computational Science and its Applications–ICCSA 2025 Workshops, pp. 102–119. Springer Nature, Switzerland, Cham (2026). https://doi.org/10.1007/978-3-031-97645-2_7

32. Gravagnuolo, A., Angrisano, M., Girard, L.F.: Dismissed historic buildings and new reuse strategies: the case of Edifici Mondo in the City of Salerno (Italy). In: Calabrò, F., Madureira, L., Morabito, F.C., Piñeira Mantiñán, M.J. (eds.) Networks, Markets & People, pp. 214–224. Springer Nature, Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-74501-0_23
33. Council of Europe: Framework Convention on the Value of Cultural Heritage for Society. Faro (2005)
34. Buglione, F. et al.: Understanding best practices of cultural heritage adaptive reuse in the perspective of the circular economy: in-depth assessment of case studies. In: Fusco Girard, L., Gravagnuolo, A. (eds.) Adaptive Reuse of Cultural Heritage, pp. 193–223. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4_7
35. Gravagnuolo, A.: Il Modello di Economia Circolare nella Città e nel Territorio. Documenti geografici, 451–457 (2022). https://doi.org/10.19246/DOCUGEO2281-7549/202102_28
36. Ramaglia, G., Mecca, I., Angrisano, M., Santagata, R., Olivieri, C.: An optimization-based framework for the structural strengthening of masonry buildings. Int. J. Space Struct. **40**, 170–179 (2025). https://doi.org/10.1177/09560599251374517
37. Coluccia, B., Tunno, V., Agnusdei, G.P.: The agricultural regeneration of Salento (Apulia, Italy) after the Xylella fastidiosa crisis: managing the shocks through multi-criteria decision-making methods. Sustainability. **17**, 8812 (2025). https://doi.org/10.3390/su17198812
38. Leone, A., Ripa, M.N., Vomero, M., Verardi, F.: Comunità energetiche e territorio, binomio indissolubile. BDC. Bollettino Del Centro Calza Bini. 221–239 (2023). https://doi.org/10.6093/2284-4732/10097.
39. Sessa, V.M.: Le Comunità Energetiche Rinnovabili quali strumento e modello per una gestione condivisa del patrimonio culturale/Renewable Energy Communities as an instrument and model for shared management of cultural heritage. Il Capitale Culturale. Stud. Value Cult. Heritage. **29**, 67–81 (2024). https://doi.org/10.13138/2039-2362/3389
40. Russo, G., Gelosia, M., Fabbrizi, G., Angrisano, M., Policastro, G., Cavalaglio, G.: Valorization of xylose-rich medium from Cynara cardunculus stalks for lactic acid production via microbial fermentation. Polymers. **16**, 3577 (2024). https://doi.org/10.3390/polym16243577
41. Cavalaglio, G., Fabbrizi, G., Cardelli, F., Lorenzi, L., Angrisano, M., Nicolini, A.: Lignocellulosic residues from fruit trees: availability, characterization, and energetic potential valorization. Energies. **17**, 2611 (2024). https://doi.org/10.3390/en17112611
42. UNESCO: the UNESCO Recommendation on the Historic Urban Landscape, Paris, France (2019)
43. Puglia, R.: Strategia Aree Interne: Strategia Area Interna, Sud Salento (2019)
44. Malavasi, G.: Enhancing public 'awaiting reuse' architectural heritage through RECs: the former ammunition depot in Sangano: La valorizzazione del patrimonio architettonico pubblico "in attesa di riuso" attraverso le CER: l'ex Polveriera Militare di Sangano. Valori Valutazioni. **39**, 45–67 (2026). https://doi.org/10.48264/VVSIEV-20263904
45. Nicolini, E.: Energy self-sufficiency of smaller rural centers: experimental approaches. Buildings. **14**, 1862 (2024). https://doi.org/10.3390/buildings14061862
46. Auernhammer, M.J.: Strategic design: the integration of the two fields of strategy and design. In: Gray, C., Ciliotta Chehade, E., Hekkert, P., Forlano, L., Ciuccarelli, P., Lloyd, P. (eds.) DRS 2024: Boston, 23–28 June. USA, Boston (2024). https://doi.org/10.21606/drs.2024.1399
47. Manzini, E.: Design em um mundo conectado e em transformação. SDRJ. **7**, 95–99 (2014). https://doi.org/10.4013/sdrj.2014.72.06
48. Khorrami, S., Falvo, M.C., Pompili, M.: Financial opportunities and challenges in energy communities: revenue, costs, and capital structures. Energies. **19**, 937 (2026). https://doi.org/10.3390/en19040937
49. Romero-Castro, N., Ángeles López-Cabarcos, M., Miramontes-Viña, V., Ribeiro-Soriano, D.: Sustainable energy transition and circular economy: the heterogeneity of potential investors in

rural community renewable energy projects. Environ. Dev. Sustain. **27**, 18051–18076 (2023). https://doi.org/10.1007/s10668-022-02898-z
50. Taromboli, G., et al.: Renewable energy communities: frameworks and implementation of regulatory, technical, and social aspects across EU member states. Sustainability. **17**, 4195 (2025). https://doi.org/10.3390/su17094195
51. Slee, B., Lukesch, R., Ravazzoli, E.: Social innovation: the promise and the reality in marginalised rural areas in Europe. World. **3**, 237–259 (2022). https://doi.org/10.3390/world3020013
52. Navarro-Valverde, F., Labianca, M., Cejudo-García, E., De Rubertis, S.: Social innovation in rural areas of the European Union learnings from neo-endogenous development projects in Italy and Spain. Sustainability. **14**, 6439 (2022). https://doi.org/10.3390/su14116439
53. Roversi, R., Boeri, A., Pagliula, S., Turci, G.: Energy Community in Action—Energy Citizenship Contract as tool for climate neutrality. Smart Cities. **5**, 294–317 (2022). https://doi.org/10.3390/smartcities5010018
54. Verardi, F.: Urban networks and environmental resilience. Fractal urban aggregates' role. Bulletin of the Calza Bini Center. 285–306 (2024). https://doi.org/10.6093/2284-4732/11391.
55. Ripp, M., Jigyasu, R., Gustafsson, C.: Transformation Management of Heritage Systems. Heritage. **9**, 28 (2026). https://doi.org/10.3390/heritage9010028
56. Bosone, M., Ciampa, F.: Human-Centred indicators (HCI) to regenerate vulnerable cultural heritage and landscape towards a Circular City: from the Bronx (NY) to Ercolano (IT). Sustainability. **13**, 5505 (2021). https://doi.org/10.3390/su13105505
57. Angrisano, M., Neglia, G., Mecca, I.: Building the future while preserving the past: digital tools and community engagement to "re-generate" historic buildings in post-disaster scenarios. In: Gervasi, O., Murgante, B., Garau, C., Karaca, Y., Faginas Lago, M.N., Scorza, F., Braga, A.C. (eds.) Computational Science and its Applica-Tions – ICCSA 2025 Workshops, pp. 299–315. Springer Nature, Switzerland, Cham (2026). https://doi.org/10.1007/978-3-031-97654-4_19
58. Landa-Oregi, I., Gonzalez-Ochoantesana, I., Anaya-Rodriguez, M.: Understanding the engagement of citizens in the regeneration of urban spaces through human-centered design: a systematic literature review. Cities. **168**, 106414 (2026). https://doi.org/10.1016/j.cities.2025.106414
59. Vercher, N., Bosworth, G., Esparcia, J.: Developing a framework for radical and incremental social innovation in rural areas. J. Rural. Stud. **99**, 233–242 (2023). https://doi.org/10.1016/j.jrurstud.2022.01.007
60. European Commission: The New Leipzig Charter: The transformative power of cities for the common good. Adopted at the Informal Ministerial Meeting on 30 November 2020 (2020)
61. Scardigno, R., Sportelli, C., De Simone, E., Gatti, F., D'Errico, F.: The communicative processes in the renewable energy communities (REC): a systematic review. J. Environ. Stud. Sci. (2025). https://doi.org/10.1007/s13412-025-01070-z
62. De Simone, E., Rochira, A., Procentese, F., Sportelli, C., Mannarini, T.: Psychological and social factors driving citizen involvement in renewable energy communities: a systematic review. Energy Res. Soc. Sci. **124**, 104067 (2025). https://doi.org/10.1016/j.erss.2025.104067
63. Decreto Legislativo 8 novembre 2021, n. 199. (Attuazione della direttiva (UE) 2018/2001 del Parlamento europeo e del Consiglio, dell'11 dicembre 2018, sulla promozione dell'uso dell'energia da fonti rinnovabili, come modificata dalla direttiva (UE) 2023/2413 del Parlamento europeo e del Consiglio, del 18 ottobre 2023, che modifica la direttiva (UE) 2018/2001, il regolamento (UE) 2018/1999 e la direttiva n. 98/70/CE per quanto riguarda la promozione dell'energia da fonti rinnovabili e che abroga la direttiva (UE) 2015/652 del Consiglio). GU n. 285 del 30-11-2021- Suppl. Ordinario n. 42
64. Decreto Ministeriale n. 414/2023. Decreto CER. Ministero dell'Ambiente e della Sicurezza Energetica. 07 Dec 2023

65. McCann, P.: How Have Place-Based Policies Evolved to Date and What Are They For Now?. In: Background paper for the OECD-EC High-Level Expert Workshop Series on "Place-Based Policies for the Future", Workshop 1, 14 April 2023 (2023)
66. Circular Economy Network: 7° Rapporto sull'Economia Circolare in Italia 2025. Italy, Rome (2025)
67. Koltunov, M., et al.: Mapping of energy communities in Europe: status quo and review of existing classifications. Sustainability. **15**, 8201 (2023). https://doi.org/10.3390/su15108201
68. Presidenza del Consiglio dei Ministri: Piano Strategico Nazionale delle Aree Interne (2025)
69. Regione Puglia: Programma Regionale Puglia 2021–2027. Avviso Pubblico per la selezione di proposte progettuali finalizzate alla costituzione di Comunità Energetiche Rinnovabili (CER). https://politiche-energetiche.regione.puglia.it/-/avvisocer?redirect=%2F-%2Favvisocerproroga16aprile. Access 5 May 2026
70. Legambiente, G.S.E.: Comuni Rinnovabili. Come le rinnovabili hanno cambiato il Paese e portato Vantaggi, Sfide e obiettivi ancora da raggiungere (2025)
71. Interreg IPA South Adriatic Program: GRowing participative models for EnErgy commuNities – GREEN https://etmi-al.org/project/green-growing-participative-models-for-energy-communities/. Access 5 May 2026)
72. Unione dei Comuni della Grecìa Salentina: https://www.unionegreciasalentina.le.it. Access 5 May 2026

Co-designing Alternatives of Adaptive Reuse: The Case of the Church of the Holy Name of Jesus in Moncalieri

Maria Tartari[1], Giulio Cavana[2(✉)], Mariarosaria Angrisano[3], Gravagnuolo Antonia[3], Luigi Fusco Girard[4,5], and Marta Bottero[2]

[1] Libera Università di Lingue e Comunicazione (IULM), Milan 20143, Italy
[2] Interuniversity Department of Regional and Urban Studies and Planning (DIST), Politecnico di Torino, 39, Viale Mattioli, 10125 Turin, TO, Italy
giulio.cavana@polito.it
[3] Engineering Department, Pegaso Telematic University, Naples 80132, Italy
[4] Institute of Heritage Science, National Research Council (CNR-ISPC), Naples 80134, Italy
[5] Architecture Department, University of Naples Federico II, Naples 80134, Italy

Abstract. This paper presents a participatory co-design process focused on the adaptive reuse of the Chiesa del SS. Nome di Gesù, a deconsecrated church in the historic centre of Moncalieri, Northern Italy, developed within the PRIN TREE 2022 research project on circular economy models for architectural heritage reuse. Through a full-day participatory workshop with 14 participants and semi-structured validation interviews with seven key informants, the study employed a Circular Business Model Canvas to bridge participatory outputs and governance planning. Findings consistently identify the Church as a potential generative infrastructure for a new intergenerational community, a civic common where different generations, cultural backgrounds, and forms of knowledge enter into durable relationship. Three co-generative value dimensions are mapped (cultural, social, and economic) alongside a polycentric governance model. Interpreted through the concept of heritage community as defined in the Council of Europe's Faro Convention, the results suggest that the most sustainable form of religious heritage reuse is one that simultaneously activates a new community of meaning around the heritage site. Implications for heritage governance, participatory design methodology, and Faro Convention operationalisation are discussed.

Keywords: Reuse of cultural heritage · circular economy · co-programming · co-participation · Circular business canvas model

1 Introduction

The construction sector is a major consumer of raw materials and energy, and a significant source of carbon emissions [1, 2]. On top of this, the sector is also responsible of the production of a substantial amount of waste deriving from demolition processes [3]. Within this context, adaptive reuse emerges as a potential way to update the building stock, avoiding the resource-intensive cycle of demolition and reconstruction [4].

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 541–559, 2027.
https://doi.org/10.1007/978-3-032-30518-3_34

Adaptive reuse is broadly defined as any intervention aimed at modifying a building's capacity, function, or performance in order to meet new requirements [5]. Through both functional and physical transformations, adaptive reuse enables existing buildings to be given a second life.

Beyond its environmental benefits, this practice contributes to the social, economic and cultural dimensions of sustainability, representing a comprehensive approach that addresses multiple challenges simultaneously [6, 7]. By proposing new functions that are compatible with the heritage value of buildings, adaptive reuse promotes the long-term sustainable use of resources, in line with the principles of the circular economy. Furthermore, by preserving the cultural significance of buildings, it could provide an opportunity to enhance local community engagement while fostering economic well-being [8]. Reusing built heritage could be a driving force in the revitalisation of urban and peri-urban areas, or those affected by long-term demographic trends of depopulation, which leave a large proportion of the building stock underutilised, as is the case in marginal and inner areas [9].

With over 500,000 structures, religious buildings constitute a substantial proportion of Europe's built heritage and could present opportunities for adaptive reuse [10]. In peripheral areas in particular, this proportion of the building stock suffers from secularisation, demographic decline and urbanisation, leading to abandonment and deterioration, with the need to repurpose these buildings for new, non-liturgical functions [11–14]. Indeed, the dissolution of religious institutions and the transition to alternative forms of worship has had a profound impact on the religious heritage of many Western countries, effectively transforming it into an economic burden. In addressing this issue, tourism-related uses emerge as a prevalent adaptive reuse strategy. However, this approach is not without its limitations. Firstly, it is susceptible to the commodification of cultural values within the so-called "economy of experience" [15]. Secondly, it often fails to deliver sustainable or regenerative outcomes. Furthermore, it does not consistently represent a viable solution in peripheral or marginal areas. It is evident that economic viability continues to be a primary factor in determining success. The financial sustainability of uses is a pivotal aspect in ensuring the long-term preservation of heritage assets, which in turn contributes to broader sustainability goals [4].

Nonetheless, a number of instances of successful adaptive reuse illustrate the capacity for regeneration, particularly when participatory processes such as co-design and co-planning are initiated to establish novel uses for religious buildings [16, 17]. Engaging local communities is pivotal in ensuring that projects reflect shared values and needs, thereby fostering a sense of ownership and long-term commitment. The selection of appropriate functions constitutes a critical stage in the process [18], from both a technical standpoint (with regard to spatial configurations, structural characteristics, and material constraints) and a social standpoint (taking into account possible impacts and outcomes for the community). In this regard, several methodologies have been widely explored in the literature both ex-ante and ex-post [8].

Despite the recognised importance of community participation, its practical implementation remains challenging. While stakeholders generally recognise its value, the involvement of the public at every stage of the process is often deemed impractical due to the complexity of decision-making and the diversity of interests involved. As has

been noted by a number of practitioners, this results in community input being expressed informally through local groups and associations [4]. The feasibility of uses is ultimately dictated by market dynamics, especially in cases where there is competition between public and private interests [4].

In this context, the PRIN-TReE research project (Grant PRIN 2022 No. 2022TN5M7F) has been conceived with the objective of identifying evaluation processes that can assist decision-makers in defining adaptive reuse strategies for cultural built heritage. The project involves engaging local communities in the co-construction of adaptive reuse processes, with a view to ensuring long-term feasibility. The present contribution delineates the present stage of the process initiated for the adaptive reuse of the abandoned Church of the Holy Name of Jesus (Chiesa del Santissimo Nome di Gesù) in Moncalieri, in Turin, Italy. In particular, it presents the participatory co-design process activated within a municipal governance context that had already taken a significant normative step toward formalising the relationship between citizens and the administration in the management of shared urban resources.

2 Case Study and Problem Situation

The Church of the Holy Name of Jesus (floor plan in Fig. 1), located in Moncalieri, was constructed in 1710 and subsequently abandoned during the 1950s. Following the consent of the Diocese of Turin to reopen the disused church, with the intention of repurposing the structure for the benefit of the local community, a preliminary workshop involving a panel of 12 experts was convened to collaboratively establish the knowledge framework surrounding the building. The in-depth description of the case study and its context, as well as the results of this preliminary workshop, was presented in a previous publication [19]. Here, the main results are summarized.

During the first workshop held in February 2025, a SWOT analysis was employed as a problem structuring method (PSM) to assist the analyst and the relevant actors in constructing the problem situation in which the decision must be made. Among the available tools, SWOT analysis is considered to be a simple and convenient method of structuring the knowledge base, engaging different experts and stakeholders, and improving communication and information exchange [20–22]. The problem-structuring exercise revealed a set of features of the case study. These were related to constraints and opportunities concerning the physical consistency, maintenance status and level of degradation of the church, with potential uses being limited by the depth of requalification intervention that could be proposed. In terms of the social context, the opportunity to ‘promote a sense of citizenship and belonging (artists of citizenship) and enhance the preservation of community memory’ also emerged. These outputs were considered to potentially benefit from establishing a relational network with other systems, such as worship and cultural/educational circuits, at local and territorial scales. Furthermore, links with institutions and local stakeholders represent potential opportunities for capitalising on synergies with local associations and possible partnership arrangements, fostering the creation of a real heritage community network.

The city of Moncalieri activated the “Regulations on Citizen-Administration Collaboration for the Care, Regeneration and Shared Management of Common Goods” in

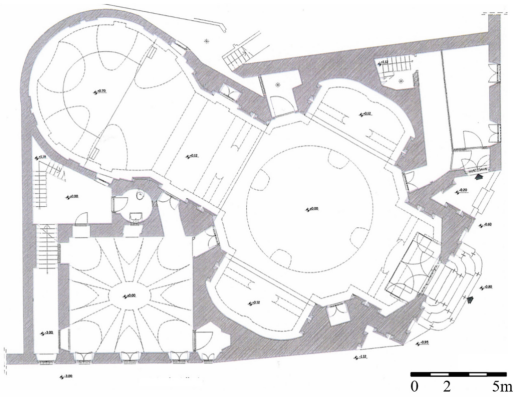

Fig. 1. Floorplan of the Church of the Holy Name of Jesus (source: [17])

April 2025. This legal instrument, modelled on the Italian "Regulation on Urban Commons" framework pioneered by the city of Bologna in 2014 and subsequently adopted by hundreds of Italian municipalities, allows citizens to cooperate with the public administration to actively and directly take care of public goods. The adoption of such tool took place independently from the research, involving the entire set of common goods in its territory, establishing a fertile ground for further initiatives.

The participatory co-design process presented in this paper was thus activated within a municipal governance context that had already taken a significant normative step towards institutionalising the relationship between citizens and the administration in the management of shared urban resources. The Moncalieri regulation establishes the legal and procedural conditions under which citizens and the administration may enter into collaboration pacts for the care, regeneration, and shared governance of material and immaterial assets recognised as functional to community wellbeing, to the exercise of fundamental rights, and to the interests of future generations. Its scope is explicitly wide-ranging, encompassing cultural and architectural heritage, and its underlying philosophy, namely subsidiarity, co-responsibility, and civic activation, is coherent with the principles of participatory heritage governance promoted by the Faro Convention. This regulatory backdrop constitutes one of the structural conditions that make the co-design process described in this paper both meaningful and actionable.

The second workshop convened around the Church on January 14, 2026, can be read, in this light, not only as a research activity but as an operationalisation of the governance model the regulation envisions: a process in which citizens, associations, cultural practitioners, educators, and institutional actors collectively imagine and begin to co-produce the future of a shared heritage asset. The collaboration pact thus functions as the institutional hinge between the abstract principle of the commons and the situated, relational practice of community building around a specific heritage site. In this sense, the Moncalieri case offers a particularly instructive example of how a local governance innovation, the commons regulation, can create the enabling conditions for a Faro-style heritage community to emerge: not as a formal declaration, but as a living process of civic co-design grounded in the concrete life of a place.

3 Methodology

The methodological approach combines participatory design practices, qualitative stakeholder engagement, and business model analysis. This hybrid strategy was chosen to address the complexity of adaptive reuse processes, which require not only technical feasibility but also social legitimacy, cultural relevance, and long-term sustainability.

The research process was structured into two sequential and complementary parts:

1. a participatory co-design workshop, aimed at generating shared visions and identifying possible functions for the space;
2. a series of semi-structured interviews with key informants, aimed at validating, refining, and contextualizing the preliminary outcomes.

Together, these phases enabled the co-production of knowledge between researchers, local stakeholders, and institutional actors (in Table 1 participants and informants are listed).

Table 1. Participants and informants.

Participants to workshop	Informants interviewed
• n° 2 members of the research group	• Local artist
• n° 2 Local architects	• Art Manager
• Staff of local public library	• Municipality bureaucrat
• Local NGO representative	• Public library manager
• Municipality representative	• National Energy
• Representative of the Diocese	• Local headteacher
• Representative of the parish	• Local deputy headteacher
• Diocese consultant	• National Agency for New Technologies, Energy, and Sustainable Economic Development researcher
• n° 2 University students	
• Urban economic expert	
• Renewable energy community representative	

Part 1 consisted of a one-day participatory design workshop held on the 14th of January, 2026. Fourteen participants took part in the activity, representing a heterogeneous group including cultural practitioners, local stakeholders, researchers, and representatives of institutions. The workshop was designed to collect situated knowledge about local needs, expectations, and opportunities related to the potential reuse of the church. The facilitation process followed principles of human-centred design and collaborative planning, ensuring that participants could contribute experiential knowledge and visions for the future of the space. The workshop was organized into a sequence of four structured phases aimed at progressively translating imaginaries into potential organizational and economic configurations.

The first activity "Imagination Phase" aimed to stimulate future-oriented thinking and creative reflection. Participants were invited to imagine the church as a fully active space in the year 2035. To facilitate this process, a visual storytelling tool based on illustrated cards was used. Participants selected images that resonated with their perceptions, aspirations, and concerns regarding the future of the church. Through associative storytelling, the images triggered reflections on possible roles for the space, including creativity, collaboration, intergenerational exchange, and cultural identity. This phase allowed participants to articulate both desires and fears, such as the risk that the building might remain unused or closed, while also envisioning the church as a vibrant cultural and civic space connected to the life of the city.

The second phase "Value Mapping" translated the imaginaries generated during the previous activity into explicit value categories. Participants worked in small groups using a large visual board structured around three dimensions: cultural value, social value and economic value. Using colored post-it notes, participants identified the types of value that the reuse of the church could generate for the community. This activity encouraged participants to move beyond the description of specific activities and focus instead on the broader impact and meaning of the space. The outcome of this phase was a collective value map, highlighting the potential role of the church as a place for cultural production, community building, learning, and social wellbeing.

In the third phase, "Stakeholder Mapping", participants identified the actors necessary to activate and sustain the envisioned functions. A stakeholder mapping exercise was conducted using a diagram structured around four categories of possible contributions: time, competences, relationships and resources. Participants mapped local institutions, associations, cultural organizations, and informal networks that could contribute to the project. This exercise revealed the presence of a local ecosystem of potential partners, ranging from schools and cultural organizations to environmental associations and local businesses. The output of this phase was a preliminary stakeholder map and governance hypothesis.

The fourth and final phase of the workshop focused on the economic and organizational sustainability of the project. Participants were invited to reflect on how the identified activities and stakeholders could generate and circulate value. Guided by the facilitators, the discussion explored questions such as who produces value within the system, who benefits from it and how people participate (as users, contributors, or supporters). The insights collected during the workshop were synthesized into a preliminary Circular Economy Business Model Canvas (CEBMC).

CEBMC [23] responds to the necessity of extending the traditional BMC to capture the complexity Circular Economy principles [24–26]. While the analysis of circular adaptive reuse strategies [27, 28] and circular business models for cultural heritage reuse have become a studied field [17, 29], studies applying a formalised tool are still not common [30].

In this research, the participatory CEBMC for cultural heritage adaptive reuse is tested as an analytical tool to organize emerging insights into categories such as vision, activities, stakeholders, governance, and sustainability mechanisms.

The second Part of the research process aimed to validate and refine the preliminary results of the workshop through targeted interviews with key informants. This

step allowed the researchers to deepen specific themes and test the feasibility of the ideas that emerged during the participatory process. Seven key informants were selected based on their strategic position within the local ecosystem, through a snowball activation. The interviewees represented different sectors, including cultural and artistic production, municipal administration, public cultural services (public library), energy and environmental research, secondary education institutions. This diversity ensured a multiperspective understanding of the project, combining technical, institutional, cultural, and educational viewpoints.

The interviews were conducted in a semi-structured format, allowing flexibility while maintaining a consistent thematic framework. Each interview lasted approximately 30–45 min and focused on three main topics:

1. the potential future functions of the church
2. the identification of relevant local stakeholders and partnerships
3. conditions required for long-term sustainability

Interviewees were also invited to comment on the preliminary CEBMC, providing feedback on its feasibility and identifying possible risks or opportunities.

Data collected were analyzed using a qualitative thematic approach. Notes and observations were reviewed to identify recurring themes, patterns, and points of convergence across participants. The findings were then integrated into an updated version of the CEBMC, which synthesizes the main components of the proposed reuse model. This iterative process allowed to move from collective imagination to structured project scenarios, ensuring that the proposed reuse strategies remained grounded in local needs while aligned with broader principles of circular heritage regeneration.

4 Results

The findings from the participatory co-design process, involving the two complementary phases (a participatory workshop and validation interviews with local stakeholders and key informants), converge on a coherent vision of adaptive reuse that integrates cultural, social, and economic dimensions within a circular framework. The findings also identify the structural conditions necessary for long-term sustainability.

4.1 Workshop Results: Emergent Value Patterns

Participants used illustrated cards as a non-directive creative device to help them project themselves into the year 2035 and imagine the church as a fully operational space. Two dominant narratives emerged. On the one hand, participants expressed anxiety about renewed closure or loss of vitality. On the other hand, the image of the church as a living hub for creative, performative and relational activity came to the fore. Recurrent visual metaphors depicted the church as a unifying and influential force in the historic centre, as a multigenerational space where different life cycles, artistic traditions and contemporary aspirations coexist, and as a site oriented simultaneously towards memory and transformation. These visions were consistent across both working groups, suggesting a shared vision that transcended individual disciplinary perspectives.

In the second phase ("Value Mapping"), participants translated their ideas into explicit value statements, organising contributions along three interdependent cultural, social and economic dimensions. These dimensions were not considered as distinct categories, but rather as co-generative processes. Cultural value was primarily associated with the construction of identity, both individual and collective; the provision of slow, non-digital physical spaces conducive to creativity, reflection and deep attention; and the promotion of encounters between different generations and cultures. The Church was envisaged as a place where memory and innovation coexist to produce shared intercultural narratives. Social value was linked to the idea of community as a process rather than a given entity. The Church was imagined as a place that could transform weak ties into stronger forms of solidarity through everyday practices such as encounter, play, care, and collaboration. Prominent themes included intergenerational integration, particularly between the elderly and young people, facilitated through traditional crafts, arts, and knowledge transmission. Another theme was the activation of self-organisation and shared responsibility, consistent with a commons-based governance model. Economic value, conceived in non-reductive terms, was considered a prerequisite rather than an end in itself: the Church's long-term survival would depend on a hybrid resource model combining monetary and non-monetary contributions, mutualistic exchanges, cross-subsidy logics and reduced social costs by producing collective well-being. Participants identified various revenue-generating activities, such as paid courses, artistic residencies, performances, temporary local product markets, event space rental, crowdfunding and subscriptions to cultural programming, as well as activities that generate high social value but require economic coverage, including free youth programmes, intergenerational play laboratories, study spaces, non-monetary exchange initiatives and time banks.

In phase 3 ("Stakeholder Mapping"), a multi-layered map of potential contributors to the adaptive reuse project was created and organised according to the four types of contribution they could make: time, competencies, relationships and resources. The map revealed a strong cluster of actors involved in education, culture and lifelong learning, including schools, universities of the third age, cultural associations and music organisations. These actors were identified as central nodes for programme development and competency activation. A second cluster encompassed local associative and institutional networks, including municipal councillors, civic committees, and environmental organisations. Their contribution was seen as essential for territorial legitimisation and connectivity to supra-local cultural circuits. A third cluster comprised proximity economic actors, including merchants' associations, local farmers, and nearby businesses. These actors were not positioned as traditional sponsors, but rather as partners in a circular territorial economy. Overall, the stakeholder map outlined a polycentric governance model, where sustainability depends on the ability to activate and coordinate diverse contributions, rather than relying on a single responsible entity.

Phase 4 involved translating the accumulated insights into a preliminary Circular Business Model Canvas as shown in Figure NN. The vision formulated for 2035 described the church as a living civic, cultural and community space that holds together memory and innovation. It is recognised as an urban common and is inhabited daily by multiple generations. Identified key functions included cultural workshops, creative laboratories, courses and training activities, intergenerational exchange programmes,

performances and exhibitions, study spaces and reading groups, and collective events tied to the local calendar. The Church would also be used daily as a neighbourhood meeting place. The target communities were residents of Moncalieri of all ages, with a particular focus on the elderly, adolescents, families with children, and individuals experiencing social fragility. These communities also included artists, researchers, educators, and cultural associations. The proposed governance model was participatory and shared, involving a coordinating entity, a community manager, and a variety of institutional, associative, and informal local actors who would share responsibilities and decision-making in a transparent and inclusive manner. Expected impacts included strengthening local cultural identity, increasing civic participation and social cohesion, improving individual and collective well-being, sustainably valorising religious heritage, reducing social isolation, activating local and circular economies, and creating a replicable adaptive reuse model (Table 2).

Table 2. Parameters used in the different Scenarios.

1. VISION FOR THE SITE. What had the Church become by 2035?		
By 2035, the Church will have become a vibrant and welcoming civic and cultural hub, where tradition meets innovation. It will be a shared, non-digital urban space for connection, creativity and growth, bringing together spiritual, cultural and social life across generations.		
2. VALUE CREATED. What value does this place bring to people and the local area?		
Cultural value. • Creation of meaning, identity and shared memory. • Space for artistic, intellectual and spiritual creativity. • Place for continuous learning, both formal and informal • Meeting point for different generations, cultures and faiths. • Creation of new intercultural memories.	Social value. • Building and strengthening the local community and relationships • Intergenerational inclusion (older people, young people and children). • Creating a space for listening, play, care and connection • Fostering self-organisation and shared responsibility.	Economic value. • Mobilisation of local monetary and non-monetary resources. • Support for cultural, educational and social initiatives. • Reduction of social costs promoting well-being • A local, mutualist and circular economy • Long-term sustainability of the area as community infrastructure.

(*continued*)

Table 2. (*continued*)

3. KEY FUNCTIONS AND ACTIVITIES.	**4. TARGET AUDIENCE / TARGET COMMUNITY.**	**5. STAKEHOLDER.**
1. Cultural, creative and artistic workshops. 2. Courses, classes, training activities (crafts, practical skills, culture and adult education). 3. Spaces for study, reading, film screenings, discussion and intergenerational exchange. 4. Performances, concerts and immersive exhibitions 5. Bartering, reusing and recycling books, games, objects and clothes. 6. Community centre and neighbourhood café hosting community events.	1. Citizens and residents of Moncalieri. 2. Older people, young people, and teenagers. 3. Families with children 4. Artists, researchers, educators and trainers 5. Cultural and social organisations 6. Local traders 7. People in vulnerable situations 8. Public authorities, private individuals and companies organising events. 9. People who do not usually visit cultural venues 10. Future generations.	1. Active local communities 2. Education, training and lifelong learning. 3. Culture, arts and creativity 4. The third sector and social care 5. Institutions and local government 6. The local economy and local supply chains 7. Environment, sustainability and regional networks 8. Supra-local cultural networks.
6. GOVERNANCE AND ROLES.	**7. SUSTAINABILITY AND RESOURCES.**	**8. EXPECTED IMPACTS.**
• The Church is managed through a participatory and shared governance model. • It is viewed as an urban common good, and a coordinating body, community manager and range of local stakeholders — including institutional, voluntary and informal groups — share transparent and inclusive responsibilities and decision-making processes.	• Donations, crowdfunding and fundraising, Sponsorship from local businesses and traders. • Educational and cultural activities on a contribution basis and Season tickets for programmed activities • Hire of spaces for compatible events • Sale of books and local products (temporary shop). • Exchange of goods, skills and time (time bank). • Skilled volunteering • Inclusion in a renewable energy community	• Strengthening local cultural identity: • Increased civic participation and social cohesion. • Improved individual and collective well-being. • Sustainable enhancement of religious heritage • Reduced social isolation and vulnerability. • Boosting local and circular economies • Creation of a replicable model for adaptive reuse.

4.2 Stakeholder Interviews: Validation and CEBMC Refinement

Phase 2 involved conducting semi-structured qualitative interviews with seven key informants, who were selected to provide insights into various aspects of territorial knowledge, including cultural and artistic production, public administration, associative networks, library and cultural services, energy, technical feasibility and secondary education. The interviews, which took place online in February 2026, aimed to validate, deepen and integrate the preliminary CEBMC produced during the workshop. For brevity, the interviews are not quoted, but are instead critically analysed to extract the main points raised by the informants.

Interviews with an artist and an arts manager emphasised Moncalieri's distinct environmental and historical identity, which is deeply connected to the Po River, the rolling hills and local agricultural traditions. This identity should not be overshadowed by its proximity to Turin. The proposal involved a strategic axis linking environmental sustainability, agricultural heritage and local food traditions, with the church acting as a cultural hub connected to biodiversity networks, UNESCO Biosphere Reserve initiatives and local supply chains. Synergies with the corporate social responsibility programmes of regional firms were also identified.

A municipal officer responsible for environmental matters confirmed the absence of a structured social space in the historic centre, outlining the potential of the church as a venue for events, exhibitions and performances, capable of hosting established, city-wide festivals and seasonal events. The key strategic objective identified in the interview was the importance of "bringing the city into the church", thereby strengthening the link between cultural programming and urban vitality. The informant also emphasised the strength of existing associative networks as a resource that could be utilised.

The director of the Arduino public library identified the lack of study rooms and conference facilities in the city centre as a structural gap, and proposed leveraging the library's physical proximity to the church to create a functional extension of cultural services. The priority target users were identified as those aged 10–18. Proposals included the creation of a dedicated, themed library section with youth-oriented collections, the provision of study spaces, coordination with local schools and the engagement of civic service workers as permanent staff. The most plausible governance scenario identified was a municipal management model with dedicated personnel.

A researcher from the Italian National Agency for New Technologies, Energy, and Sustainable Economic Development (ENEA) emphasised the importance of conservation constraints affecting the building, highlighting the need for lightweight, modular and reversible interventions that are compatible with heritage protection requirements. Such types of intervention would respond to the double necessity of respecting the physical consistency of the building, while also reducing the generation of waste from renovation interventions, ultimately linking the preservation of the historical memory, with the necessity of reducing construction related increase of the asset embodied carbon. Opportunities were identified for passive cooling strategies that exploit the building's chimney-effect morphology, natural ventilation and humidity management. The church was also identified as a potential urban thermal refuge during summer heatwaves, contributing to a low-cost local climate adaptation policy. The importance of preliminary technical measurements and administrative support for accessing funding was emphasised.

Interviews with the headteacher and deputy headteacher of a local high school revealed a clear gap in youths' lives: a lack of non-domestic, non-educational spaces where they can exercise autonomy in a safe and stimulating environment. Suggested programme activities included visual arts, music and theatre workshops, peer tutoring, university and professional orientation, and weekend activities with public showcases of student work. The interviewees emphasised the importance of involving students as co-creators of the activities, and warned that the Church could be perceived as merely an extension of the school, which would undermine its transformative potential.

5 Discussion

An integrated analysis of the workshop outputs and interview data reveals a set of findings that are both convergent and structurally coherent. Four strategic axes stand out.

Firstly, there is a consistent demand for a third space that is neither purely institutional nor commercial. This demand is particularly evident among adolescents who require non-scholastic environments in which to study, be creative, and develop peer relationships. It is also evident among adults and elderly residents who require a daily social anchor in the historic centre. Secondly, the educational and intergenerational function of the church as a study space, conference venue, orientation laboratory and creative workshop appears not only desirable, but also structurally necessary, particularly in light of its proximity to the civic library. Thirdly, Moncalieri's environmental and identity-based characteristics emerge as a distinctive asset, and the church's programme should reflect the local relationship with the Po, the hillside, biodiversity and agricultural heritage, thus positioning the site within wider environmental and cultural networks. Fourthly, governance and programming continuity are identified as the primary organisational risk factors. Episodic activation would be insufficient to generate community belonging and sustainability. This would require a stable coordinating figure, a structured programme schedule and broad partnerships with schools, associations, cultural networks and local economic actors.

Collectively, these findings suggest that the Church of the Holy Name of Jesus could function as a multifunctional civic platform and intergenerational cultural and educational hub. This infrastructure would be rooted in local identity and aligned with the principles of the ecological economy, generating diffuse and durable shared value for the city of Moncalieri.

A transversal reading of the workshop outputs and interview data reveals a theme that surfaced repeatedly and with notable consistency across both phases of the research: the idea that the renewed church should not merely host activities or services, but become the centre of a new intergenerational social network — a place around which a community is actively and continuously produced, not just gathered. Participants described this aspiration from different disciplinary vantage points and in different registers, yet the underlying intuition remained remarkably stable: the church as a site where different generations, backgrounds and forms of knowledge enter into meaningful and durable relationships with one another and where caring for shared heritage constitutes community. This finding invites a theoretical reading that goes beyond urban regeneration frameworks or cultural economics alone. It directly aligns with the concept of a heritage community, as outlined in the Council of Europe's Framework Convention on the

Value of Cultural Heritage for Society (the Faro Convention, 2005), which prioritises the relationship between individuals, heritage, and community formation within cultural heritage policy and practice. The following discussion engages with this convergence as its primary theoretical axis.

5.1 Intergenerationality as a Structural Principle

One of the most striking features of the workshop findings was how frequently and consistently intergenerational integration emerged, not as a secondary theme, but as a structural principle of the imagined community. Participants did not merely acknowledge that both the elderly and young people might use the space; they articulated a vision in which encounters between the generations would constitute the very fabric of the community. Knowledge transmission, creative collaboration, peer tutoring, sharing traditional crafts and contemporary artistic practices, and mentorship relationships between older residents and adolescents all appeared not as incidental programme elements, but as the relational core of what the Church should facilitate.

This emphasis is significant for several reasons. Firstly, it reflects a diagnosis of social fragmentation that is not unique to Moncalieri, but rather is characteristic of contemporary European urban contexts. Here, age segregation in areas such as housing, education, leisure and care tends to weaken the informal solidarity and mutual recognition that are essential for community cohesion. The Church was collectively envisaged as a structural antidote to this fragmentation: a place where encounters between different life stages are designed in, rather than left to chance.

Secondly, this intergenerational approach is directly connected to the temporal dimension of heritage communities in the Faro Convention. The Convention explicitly situates heritage communities in relation to future generations, emphasising the obligation to maintain and transmit cultural heritage across time. The workshop participants operationalised this temporal horizon in a strikingly concrete way, with several contributions explicitly referencing not only current residents, but also 'those who have not yet been born', thus framing the church's reactivation as an investment in Moncalieri's long-term social and cultural continuity. This is not merely rhetorical, but has direct implications for governance, programming and the long-term stewardship of the space.

Thirdly, the intergenerational emphasis complicates any straightforward mapping of user groups onto predetermined categories of beneficiaries, productively so. The vision that emerged from the co-design process is not of a space serving different populations sequentially or in parallel, but of a space in which relationships between age groups are the service offered. This has implications for programming logic, the design of the physical space and governance structures. It requires organisational forms that pay attention to the dynamics of encounters and translations between different generations rather than simply focusing on the efficient delivery of segmented services.

5.2 Heritage Community and Commons-Based Governance

The Faro Convention's concept of the heritage community carries implicit governance implications that are made explicit in the co-design findings. The stakeholder mapping and the preliminary Circular Business Model Canvas both point toward a polycentric,

commons-based model of governance, one in which no single actor holds exclusive authority over the space, and in which sustainability depends on the capacity to coordinate heterogeneous contributions of time, competencies, relationships, and resources. This model is not simply a pragmatic response to resource constraints; it is a normative choice, grounded in the conviction that the Church's identity as a civic and community infrastructure can only be maintained if governance arrangements reflect the distributed character of the community it serves.

The Faro Convention supports this orientation explicitly. Its provisions call for the active participation of communities in the definition, management, and transmission of cultural heritage, and emphasise the importance of shared responsibility between public institutions, civil society, and citizens. The governance model outlined in the CEBMC, a coordinating entity, a community manager figure, and a broad coalition of institutional, associative, and informal actors sharing decision-making transparently, is consistent with this framework, and with a growing body of European practice in the governance of urban commons and cultural spaces.

The interview findings add a further layer of complexity to this picture. While the commons-based model was broadly endorsed, interviewees also highlighted the risk of governance fragmentation: a space governed by too many actors, without clear coordination and stable leadership, may become a site of conflict rather than collaboration. Several key informants stressed the need for a recognisable, professionally capable coordinating figure, the community manager, whose role would be not to direct but to facilitate and sustain the relational processes on which the community model depends. This tension between distributed governance and operational coherence is a known challenge in commons-based management and will require careful attention in the design of the Church's governance architecture.

5.3 Adaptive Reuse of Religious Heritage and the Production of Civic Identity

The findings also contribute to a broader scholarly conversation on the adaptive reuse of religious heritage in European post-secular urban contexts. The progressive reduction in liturgical use of church buildings across Western Europe has generated both a patrimony management challenge, how to preserve architecturally, historically, and symbolically significant structures in the absence of their original communities of use, and a set of opportunities for civic and cultural repurposing that remains underexplored in policy and practice. Most existing frameworks for the adaptive reuse of religious buildings tend to prioritise architectural conservation, economic viability, or functional fit, without adequately addressing the symbolic and relational dimensions of heritage transformation.

The Moncalieri case offers a more integrated perspective. The co-design process generated a vision in which the physical transformation of the Church is inseparable from a social transformation: the building changes function because the community changes, and the community changes because the building is reimagined. This circularity, between space, programme, and community, is precisely what the Faro Convention's concept of heritage community is designed to capture. It suggests that the most sustainable form of adaptive reuse is not one that finds a new use for an old building, but one that activates a new community of meaning around a heritage site, enabling that community to define

the terms of its relationship with the past in ways that are genuinely generative for the present and the future.

The interview with the cultural practitioners further enriched this dimension by drawing attention to the specificity of Moncalieri's territorial identity, its relationship with the Po River, the hillside, the agricultural landscape, and the local heritage of production and craft. This territorial rootedness is not incidental to the heritage community model; on the contrary, it is one of its constitutive dimensions. The Faro Convention emphasises that heritage communities are always situated: they are defined not only by what they value but by where they are, and by the particular landscape, history, and social ecology that shapes their sense of place. In Moncalieri, the reactivation of the Church thus has the potential to function as a territorial anchor, a site at which the community's relationship with its broader environmental and cultural heritage is continuously renegotiated and renewed.

5.4 The CBM as a Tool for Heritage Community Governance

A distinctive methodological contribution of this research lies in the application of the Circular Business Model Canvas as an instrument for structuring and operationalising the co-design process. The CEBMC proved effective not primarily as a financial planning tool, its conventional function, but as a mediating device between the imaginative and relational outputs of participatory processes and the analytical categories required for governance and sustainability planning. By translating emergent values, stakeholder configurations, and activity proposals into a shared schema of functions, resources, relationships, and impacts, the CEBMC enabled participants and researchers to maintain a systemic view of the Church's potential without prematurely foreclosing the openness of the co-design process.

From the perspective of the Faro Convention, this methodological choice has a further significance. The Canvas, as used here, is not a blueprint for a heritage product but a map of a heritage community in formation: its cells are populated not with predetermined solutions but with collectively produced insights, tensions, and questions that remain alive and generative as the process continues. The validation interviews in Phase 2 demonstrated the CEBMC's capacity to absorb and integrate contributions from different sectoral perspectives, cultural, educational, environmental, administrative, technical, while maintaining the coherence of the overall vision. This suggests that the CEBMC, combined with a participatory co-design methodology, may constitute a replicable instrument for heritage community governance in comparable adaptive reuse contexts across Europe.

6 Conclusions

Adopted in 2005 and progressively ratified across Europe over the following two decades, the Council of Europe's Faro Convention introduces a conceptual and normative shift in the way cultural heritage is understood and governed. Rather than focusing on monuments, buildings or collections as objects of protection, the Convention highlights the

relationship between people and heritage, and the meanings, practices and social processes that heritage activates and sustains. The Convention's central concept is that of the heritage community, defined as 'people who value specific aspects of cultural heritage which they wish, within the framework of public action, to maintain and transmit to future generations'. This definition is deliberately open; a heritage community is not a pre-existing social group and is not necessarily bounded by geography, institutional affiliation or shared religious identity. Rather, it is a community in the making, constituted through the shared act of recognising, caring for, and passing on something held to be of collective significance.

The findings of this research speak directly to this theoretical framework with a specificity that deserves careful analytical attention. Throughout both phases of the co-design process, the participatory workshop and the stakeholder interviews, a consistent and recurring aspiration emerged: that the Chiesa del SS. Nome di Gesù should not merely become a venue for cultural activities, but the generative infrastructure of a new intergenerational community. Participants did not present this as nostalgia for a pre-existing parish community or as a mere extension of existing associative life in Moncalieri. Rather, they articulated something more demanding and interesting: the idea that the process of collectively deciding what the church should become and then inhabiting and co-managing it would produce the social fabric they were imagining. In Faro terms, the heritage community and the heritage project are simultaneous and mutually constitutive, rather than sequential (first the community, then the project). Community is not a precondition of adaptive reuse, but one of its primary outputs.

This interpretation fundamentally changes the way we view the adaptive reuse of religious heritage. It suggests that the most significant value generated by the reactivation of the Chiesa del SS. Nome di Gesù would not be architectural, economic or narrowly cultural, although it may be all of these, but would be fundamentally relational: producing new forms of intergenerational solidarity, shared responsibility and civic identity. This aligns closely with the Faro Convention's assertion that heritage has a 'human rights dimension' and that democratic participation in heritage processes is a social good in itself, not just a procedural requirement.

6.1 Limitations and Open Questions

Several limitations of the present study should be acknowledged. The research presented here is necessarily preliminary: Phase 2 interviews were completed at the time of writing, but the broader participatory process, including the formalisation of governance arrangements, the definition of an operational programme, and the mobilisation of resources, remains in progress. The findings thus reflect a specific moment in an evolving process, and should be read as inputs to ongoing governance deliberation rather than as definitive prescriptions.

A further limitation concerns the composition of the workshop participant group. Although the 14 participants were selected to represent a broad range of disciplinary perspectives and sectoral affiliations, the sample inevitably reflects the existing networks of the research team and the local organisational context. Future stages of the co-design process will need to extend participation to groups currently underrepresented, in particular, adolescents and young people, whose needs were extensively discussed

but whose voices were largely mediated through institutional interlocutors rather than directly expressed.

Finally, the theoretical articulation between adaptive reuse, circular economy principles, and the Faro Convention's heritage community framework remains at an early stage of elaboration. The convergences identified in this paper suggest that this triangulation is analytically productive and policy-relevant, but a more systematic theoretical development, grounded in comparative case studies and a deeper engagement with the heritage studies and urban commons literatures, will be required to consolidate its contributions.

Funding. Funding Research Grant PRIN 2022 No. 2022TN5M7F on "TReE - Supporting the Transition to Ecological Economy in Italian cities Regeneration: circular model tools for reusing architecture and infrastructures". Italian Ministry of University and Re-search (MUR).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Disclosure. The authors declare that no Generative AI was used in the creation of the scientific content of this manuscript. AI Assisted Technology was used solely for translation and grammar refinement.

References

1. Lederer, J., Gassner, A., Fellner, J., Mollay, U., Schremmer, C.: Raw materials consumption and demolition waste generation of the urban building sector 2016–2050: a scenario-based material flow analysis of Vienna. J. Clean. Prod. **288** (2021). https://doi.org/10.1016/j.jclepro.2020.125566
2. Cavana, G., Becchio, C., Bottero, M.: Feasibility and evolution studies on renewable energy communities in cities. Renew. Sust. Energ. Rev. **213**, 115477 (2025). https://doi.org/10.1016/j.rser.2025.115477
3. Yang, X., Hu, M., Zhang, C., Steubing, B.: Urban mining potential to reduce primary material use and carbon emissions in the Dutch residential building sector. Resour. Conserv. Recycl. **180** (2022). https://doi.org/10.1016/j.resconrec.2022.106215
4. Yung, E.H.K., Chan, E.H.W.: Implementation challenges to the adaptive reuse of heritage buildings: towards the goals of sustainable, low carbon cities. Habitat Int. **36** (2012). https://doi.org/10.1016/j.habitatint.2011.11.001
5. Douglas, J.: Building Adaptation. Routledge (2006). https://doi.org/10.4324/9780080458519.
6. Della Spina, L.: Cultural heritage: a hybrid framework for ranking adaptive reuse strategies. Buildings. **11** (2021). https://doi.org/10.3390/buildings11030132
7. Angrisano, M., Bottero, M., Cavana, G., Fabbrocino, F., Gravagnuolo, A., Fusco Girard, L.: Adaptive reuse of cultural built heritage: towards the implementation of the circular city model. Front. Built Environ.
8. Andreucci, M.B., Karagözler, S.: Adaptive Reuse of Existing Buildings. Presented at the (2025). https://doi.org/10.1007/978-3-031-73490-8_11
9. Barbaro, S., Napoli, G., Trovato, M.R.: Circular Economy and Social Circularity. Diffuse Social Housing and Adaptive Reuse of Real Estate in Internal Areas. Presented at the (2022). https://doi.org/10.1007/978-3-031-12814-1_14
10. European Commission: The Future for religious heritage. European cultural heriatge strategy for 21th century (2021). www. coe.int/strategy21

11. Longhi, A.: Storie di architettura ecclesiale e processi di patrimonializzazione: valori, resilienza, adattività, riuso. BDC. Bollettino Del Centro Calza Bini. **19**, 9–26 (2019)
12. Barolomei, L., Longhi, A.: Italian debates, studies and experiences concerning reuse projects of dismissed religious heritage. In: Wandel und wertschatzung. Synergien fur die Zunkunft von Kirchenraumen, pp. 107–135 (2017)
13. Franco, G.: Adaptive reuse of religious and sacred heritage: preserving material traces and Spirit of place. Heritage. **7**, 4725–4754 (2024)
14. Sedova, A.: Impact analysis on adaptive reuse of obsolete ecclesiastical cultural heritage. Euro. J. Cult. Manag. Policy. **12** (2022). https://doi.org/10.3389/ejcmp.2022.11083
15. Pine, J.B., Gilmore, J.H.: The Experience Economy. Work is Theatre–Every Business a Stage. , Milan (2000)
16. Sadowy, K., Szemző, H.: Aesthetics, gentrification and new identities: the comparison of adaptive reuse practices in contemporary Budapest and Warsaw. J. Cult. Heritage Manag. Sustain. Dev. **14** (2024). https://doi.org/10.1108/JCHMSD-09-2022-0172
17. Buglione, F., et al.: Understanding best practices of cultural heritage adaptive reuse in the perspective of the circular economy: in-depth assessment of case studies. In: Adaptive Reuse of Cultural Heritage, pp. 193–223. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4_7
18. Abastante, F., Corrente, S., Greco, S., Lami, I.M., Mecca, B.: The introduction of the SRF-II method to compare hypothesis of adaptive reuse for an iconic historical building. Oper. Res. **22** (2022). https://doi.org/10.1007/s12351-020-00611-4
19. Cavana, G., Angrisano, M., Gravagnuolo, A., Girard, L.F., Bottero, M.: Adaptive reuse of religious cultural heritage: the case study of the church "Santissimo Nome di Gesù" in Turin (Italy). In: Lecture Notes in Computer Science (2026). https://doi.org/10.1007/978-3-031-97654-4_17
20. Kurttila, M., Pesonen, M., Kangas, J., Kajanus, M.: Utilizing the analytic hierarchy process (AHP) in SWOT analysis - a hybrid method and its application to a forest-certification case. Forest Policy Econ. **1** (2000). https://doi.org/10.1016/s1389-9341(99)00004-0
21. Margles, S.W., Masozera, M., Rugyerinyange, L., Kaplin, B.A.: Participatory planning: using SWOT-AHP analysis in buffer zone management planning. J. Sustain. For. **29**, 613–637 (2010). https://doi.org/10.1080/10549811003769483
22. Kajanus, M., Kangas, J., Kurttila, M.: The use of value focused thinking and the A'WOT hybrid method in tourism management. Tour. Manag. **25**, 499–506 (2004). https://doi.org/10.1016/S0261-5177(03)00120-1
23. Jayakodi, S., Senaratne, S., Perera, S.: Circular economy business model in the construction industry: a systematic review. Buildings. **14**, 379 (2024). https://doi.org/10.3390/buildings14020379
24. Lewandowski, M.: Designing the business models for circular economy-towards the conceptual framework (2016). https://doi.org/10.3390/su8010043
25. Hina, M., Chauhan, C., Kaur, P., Kraus, S., Dhir, A.: Drivers and barriers of circular economy business models: where we are now, and where we are heading. J. Clean. Prod. **333** (2022). https://doi.org/10.1016/j.jclepro.2021.130049
26. Braun, A.T., Schöllhammer, O., Rosenkranz, B.: Adaptation of the business model canvas template to develop business models for the circular economy. In: Procedia CIRP (2021). https://doi.org/10.1016/j.procir.2021.03.093
27. Gravagnuolo, A., Angrisano, M., Fusco Girard, L.: Circular economy strategies in eight historic port cities: criteria and indicators towards a Circular City assessment framework. Sustainability. **11**, 3512 (2019). https://doi.org/10.3390/su11133512
28. Foster, G., Saleh, R.: The adaptive reuse of cultural heritage in European Circular City plans: a systematic review. Sustainability. **13**, 2889 (2021). https://doi.org/10.3390/su13052889

29. Vellecco, I., Martone, A.: Business models for cultural heritage adaptive reuse. In: Adaptive Reuse of Cultural Heritage, pp. 463–481. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4_17
30. Saleh, R., Ost, C.: Circular business model for cultural heritage adaptive reuse: an iterative journey. In: Adaptive Reuse of Cultural Heritage, pp. 433–461. Springer International Publishing, Cham (2025). https://doi.org/10.1007/978-3-031-67628-4_16

A Multiscale Methodology for the Circular Regeneration of Historic Villages, Cultural Landscapes and Rural Territories

Albina Scioti(✉) and Mariarosaria Angrisano

Telematic University Pegaso, Naples, Italy
{albina.scioti,mariarosaria.angrisano}@unipegaso.it

Abstract. Historic villages, rural territories and cultural landscapes are increasingly recognized as strategic contexts for testing regeneration models that integrate heritage conservation, ecological transition and local development. However, interventions in these areas are still often fragmented and building-centred, with limited integration between environmental performance, conservation needs, local economies and community governance. This paper proposes a multiscale and multicriteria methodology for the circular regeneration of historic villages and rural territories. This article proposes a multiscale and multicriteria methodology for the circular regeneration of historic villages and rural territories, conceived not as a tool limited to the technical rehabilitation of individual buildings, but as a framework oriented toward territorial regeneration. From this perspective, historic building stock, adaptive re-use, the selection of materials, and the use of appropriate technologies are interpreted as integrated components of a broader process, which also involves governance arrangements, local economies, community participation, eco-logical infrastructures, and service systems. Within this approach, evaluation plays a central role: it is not understood merely as an ex-post tool for verifying implemented interventions, but as a continuous decision-support process capable of guiding design choices, strengthening governance, and aligning local transformations with a shared community vision.

Keywords: Circular Regeneration · Historic Villages · Rural Territories · Cultural Heritage · Multicriteria Evaluation · Adaptive Reuse · Circular Region · Community-Based Regeneration

1 Introduction

Recent debate on the circular economy has progressively extended from industrial systems and waste reduction to the built environment, where circularity is increasingly interpreted as a strategy for extending building life cycles, reducing re-source consumption, and strengthening local resilience (Foster, 2020; Finamore et al., 2024). In particular, recent studies have shown that, in the heritage and existing-building sectors, restoration, continuous use, rehabilitation, and adaptive reuse are central to improving circularity and reducing environmental impacts (Foster, 2020; Dişli, 2023). While this transition

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 560–571, 2027.
https://doi.org/10.1007/978-3-032-30518-3_35

has been widely explored in metropolitan and urban contexts, rural municipalities and historic village systems remain comparatively under-investigated, despite their strategic relevance for ecological transition and territorial cohesion.

This gap is particularly significant in historic villages and rural territories, where built heritage, landscape, local resources, and community practices are deeply interconnected. In such contexts, heritage should not be understood as an isolated stock of buildings, but as a territorial infrastructure capable of activating regeneration processes, local economies, and forms of ecological and social resilience. Recent contributions on heritage led rural regeneration and community-based planning suggest that these areas can operate as laboratories for sustainable development, especially when regeneration is linked to local participation, co-creation, and heritage-based innovation processes (De Luca et al., 2021; Egusquiza et al., 2021).

This challenge can be interpreted through the concepts of the circular city and circular region, with particular attention to the role of multidimensional and multicriteria evaluation, indicators, and participatory approaches in guiding circular transformations across villages, landscapes, and rural systems. Recent research has stressed both the need for urban and territorial scale circularity indicators and the value of participatory evaluation frameworks for adaptive reuse projects, particularly where heritage, community interests, and long-term sustainability goals must be balanced (Paoli et al., 2022; Gravagnuolo et al., 2024).

In this perspective, the paper proposes a methodology that links heritage conservation, circular regeneration, and territorial evaluation within a single decision-support frame-work. Unlike existing approaches, which tend to focus either on adaptive reuse at the building scale or on rural regeneration as a policy process, the proposed methodology integrates territorial profiling, transformability assessment, participatory scenario-building, and multidimensional evaluation within a single multiscale framework. This integration is consistent with recent work emphasizing the environmental benefits of circular adaptive reuse, the role of heritage communities in decision-making, and the relevance of participatory multicriteria assessment for selecting regeneration scenarios in complex contexts (Foster, 2020; Gravagnuolo et al., 2021; Gravagnuolo et al., 2024).

Accordingly, the term regeneration is used here to indicate an integrated territorial process that combines the recovery and adaptive reuse of historic buildings with broader strategies for landscape enhancement, ecological transition, service innovation, local economic activation and community-based governance. The framework therefore does not focus exclusively on building-stock rehabilitation; rather, it uses the building stock as one entry point within a wider village and rural-territorial system.

2 Theoretical Background

2.1 Circular Regeneration and Heritage-led Territorial Development

In the heritage field, circularity has progressively moved beyond simple material recycling to encompass adaptive reuse, reversibility, value retention, life-cycle thinking and the re-generation of existing assets. Recent studies highlight that the conservation and reuse of existing buildings represent one of the most effective strategies for reducing environmental im-pacts in the construction sector, as they preserve embodied energy,

materials and cultural values while avoiding demolition and new resource consumption (Foster, 2020; Dişli, 2023). In this perspective, adaptive reuse is not only a technical intervention, but a systemic approach that integrates evironmental, cultural, social and economic dimensions within a circular framework (Gravagnuolo et al., 2018) argue that adaptive reuse of cultural heritage is a key strategy for implementing circular economy models in cities and territories, because it preserves embodied cultural and material value while enabling the activation of new social, cultural and economic functions (Angrisano et al., 2025). Similarly, recent contributions emphasize the role of participatory processes and heritage com-munities in enhancing circularity outcomes, particularly in con-texts where decision-making must balance conservation needs with innovation and local development goals (Gravagnuolo et al., 2021; Egusquiza et al., 2021).

When applied to rural systems, however, circularity must be interpreted at a broader territorial scale, extending beyond in-dividual buildings to include settlement structures, landscape systems and local resource networks. Rural territories are characterized by low-density patterns, strong interdependencies between natural and built environments, and the presence of underused heritage assets, which require integrated and place-based approaches to regeneration. Recent European poli-cy frameworks highlight that circular economy strategies still face significant challenges in achieving integration across local and regional scales, particularly in non-urban and peripheral areas (European Commission, 2020; European Commission, 2024). In addition, European territorial policies emphasize the need to strengthen cohesion, resilience and place-based development strategies in rural and inner areas, recognizing their strategic role in the green transition (European Commission, 2021).

This policy dimension is particularly relevant in the Italian context, where historic villages and rural territories often overlap with inner areas affected by demographic decline, reduced accessibility to essential services and fragile local economies. For this reason, the proposed framework should be read in relation to the Italian National Strategy for Inner Areas (SNAI), which promotes place-based development, service accessibility and territorial cohesion in marginal areas. At the European level, it is also consistent with the Long-Term Vision for the EU's Rural Areas, the Circular Economy Action Plan and the Circular Cities and Regions Initiative, all of which stress the need to connect ecological transition, local development and multilevel governance in rural and peripheral con-texts.

Recent studies on heritage-led rural regeneration further suggest that villages and rural landscapes can function as laboratories for circular transition, especially when regeneration processes are linked to local resources, community participation and multilevel governance frameworks (De Luca et al., 2021; Egusquiza et al., 2021; Suárez et al., 2025; Scioti et al., 2022).

This perspective calls for the development of place-based and multiscale methodologies capable of translating circular economy principles into low-density territorial systems, where village networks, ecological infrastructures and local communities play a decisive role (Verardi et al., 2023; Angrisano, M.; 2023). In such contexts, circularity should be understood not only as resource efficiency, but as a broader process of territorial regeneration aimed at strengthening resilience, reactivating local economies and

enhancing the long-term value of cultural landscapes (Foster, 2020; Paoli et al., 2022; Gravagnuolo et al., 2024).

The literature reviewed in this paper can therefore be organized into three complementary strands. The first concerns circularity and adaptive reuse in the built heritage field, with attention to embodied value, reversibility, durability and com-patible technologies. The second addresses rural and inner-area regeneration, emphasizing territorial cohesion, local re-sources, community participation and multilevel governance. The third focuses on evaluation frameworks, indicators, KPIs and multicriteria decision-support tools, which are needed to compare alternative scenarios and make trade-offs explicit. The originality of the present contribution lies in bringing these strands together within a single methodological sequence for heritage-sensitive rural territories.

2.2 Evaluation as a Strategic Process

A second key issue concerns the role of evaluation. In heritage-led regeneration, evaluation is no longer understood only as an ex-post verification of impacts, but increasingly as a continuous and strategic process supporting the definition of scenarios, priorities and governance arrangements. This perspective is particularly strong in the literature on circular adaptive reuse, where evaluation criteria are used to balance cultural, environmental, social and economic dimensions (Gravagnuolo et al., 2018).

Similarly, the RURITAGE methodology shows that community-based regeneration in rural areas benefits from participatory and collaborative processes, organized through Rural Heritage Hubs and co-creation practices, in which evaluation supports decision-making and aligns actions with locally shared visions (Buglione et al., 2026; De Luca et al., 2021).

This perspective suggests that evaluation should be embedded within the regeneration process itself, not limited to the final measurement of impacts.

2.3 Research Gap and Aim of the Paper

Current scientific literature review tends to address circular economy, heritage conservation, adaptive reuse and rural re-generation through partially separated perspectives. Heritage studies frequently privilege conservation and building-scale adaptive reuse; circular economy studies often concentrate on material flows, resource efficiency and technical performance; while rural regeneration research tends to emphasize governance, local development and territorial cohesion. Although each of these strands is relevant, existing approaches rarely provide an operational methodology capable of connecting building stock, cultural landscape, local resource systems, community actors and evaluation criteria across multiple scales.

Historic built heritage represents the result of a long and complex evolutionary process, intrinsically linked to the functional role of buildings and to the territorial and environmental conditions in which they are embedded. The original con-figuration of these architectural structures, shaped by the geo-graphical and climatic features of the site, the availability of local materials, traditional construction techniques, and the socio-cultural practices of resident communities, has progressively changed over time under the

influence of evolving planning policies, management strategies, technological progress, and regulatory developments.

In numerous cases, the lack of comprehensive guidance frameworks and decision-support methodologies for interventions in complex historic environments—frequently subject to heritage protection constraints—has created substantial operational challenges both for technical professionals and for public administrations responsible for supervising transformation processes. As a consequence, many improvement initiatives have produced fragmented and often ineffective outcomes, contributing to widespread degradation across entire settlement systems, with only a few isolated building units benefiting from targeted interventions.

At the same time, where intervention processes have been undertaken, rehabilitation activities have commonly been implemented through a succession of isolated actions focused on individual buildings and developed according to a predominantly "case-by-case" design approach. Although these interventions have generally adopted technically reliable and technologically appropriate solutions, such a fragmented method-ology has significantly prolonged the overall recovery process while progressively weakening the perception of the settlement as a coherent and unified organism. This has ultimately generated functional and managerial fragmentation within the historic built environment.

From this perspective, addressing the issues related to sustainable and circular reuse requires the adoption of an integrated systemic approach. Such an approach should consider both the territorial context and the individual buildings as interconnected elements of a single complex system. Moreover, it enables the compensation or integration of operational limitations affecting specific buildings through opportunities and intervention potentials identifiable at broader territorial and settlement scales.

Within this interpretative framework, the comprehensive assessment and qualification of the environmental, contextual, architectural, constructional, and performance-related characteristics of the built heritage becomes a key objective. The aim is to define a balanced operational framework capable of ensuring the conservation of original material and formal values, the enhancement of both existing and potential relationships between the built environment and its surrounding context, and the technically compatible improvement of the actual performance of historical construction systems.

The specific gap addressed by this paper is therefore methodological: the lack of a clear multiscale framework for sup-porting circular regeneration in historic villages and rural territories, where technical choices on materials and technologies must be coordinated with wider territorial scenarios, governance processes and community priorities. Existing studies provide useful concepts and assessment tools, but they do not sufficiently clarify how these components can be combined into a decision-support sequence applicable to low-density and heritage-sensitive contexts.

The aim of the paper is to define a simplified but structured theoretical-methodological framework for the circular regeneration of historic villages, cultural landscapes and rural territories. The expected output is not an empirical validation, but the formalization of a transferable decision-support frame-work that identifies phases, actors, data categories, evaluation domains and possible outputs. The framework operates at multiple scales, from the building and open-space scale to the village, landscape

and territorial network scale, and is intended to support the co-definition and comparison of regeneration scenarios.

3 The Methodology for the Circular Regeneration of Historic Villages, Cultural Landscapes and Rural Territories

The framework was constructed through a theoretical synthesis of three bodies of literature: circular adaptive reuse and heritage conservation; rural and inner-area regeneration; and evaluation, indicators and multicriteria decision-support approaches. These strands were compared in order to identify recurrent analytical dimensions, actor groups, decision points and expected outputs. The four methodological phases there-fore derive from a structured conceptual integration rather than from a single case-study application. This makes the pa-per primarily methodological and theoretical, while future empirical testing will be required to validate and refine the framework (Fig. 1).

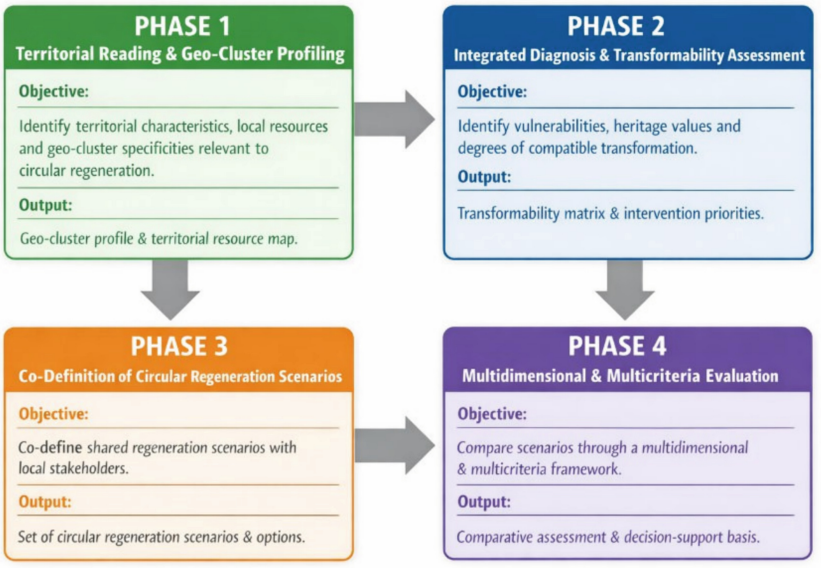

Fig. 1. Methodology for the Circular Regeneration of Historic Villages, Cultural Landscapes and Rural Territories

Phase 1 - Territorial reading and geo-cluster profiling focuses on the construction of a territorial profile of the village or network of villages.

The analysis includes environmental and climatic conditions, landscape structure, settlement morphology, recurrent building types, local materials, available energy resources, underused heritage assets, accessibility, services and socio-economic net-works. The concept of geo-cluster is used to identify relatively homogeneous areas sharing environmental, constructive and socio-cultural characteristics (Sesana M.M. et al., 2015; Lombardo G. et al., 2015; De Fino M., et al. 2017).

This phase is essential because circular regeneration depends on the recognition of territorial invariants and local potentials: reusable building stock, local supply chains, renewable energy sources, ecological assets, and community capacities. In line with the literature on rural circularity, the purpose is to move from a purely building-based reading toward a territorial under-standing of resource flows and development opportunities (Alonso et al., 2024). The output of this phase is to generate a geo-cluster profile and territorial resource map (Fig. 1).

The Phase 2 — Integrated diagnosis and transformability assessment focuses on the diagnosis of the historic building stock and the village system. It includes the assessment of conservation conditions, construction pathologies, energy weaknesses, environmental vulnerabilities, functional obsolescence and reuse potential. At the same time, it identifies the cultural and landscape values that must be preserved and the degrees of compatible transformability.

A key aspect of this phase is that performance is not limited to operational energy efficiency. Instead, it includes resilience, adaptability, maintenance capacity, landscape compatibility and contribution to territorial cohesion. This broader notion is consistent with circular adaptive reuse approaches, which stress that heritage interventions should be assessed through multiple dimensions rather than through energy criteria alone (Gravagnuolo et al., 2018). The Output is to develop a trans-formability matrix, critical issues and intervention priorities (Fig. 1).

Phase 3 - Co-definition of circular regeneration scenarios is based on an explicitly participatory approach. It involves local administrations, residents, associations, technical experts, economic actors and third-sector organizations in the co-definition of regeneration scenarios. In this sense, the method-ology aligns with the emphasis on co-design, co-programming and co-evaluation, as well as with the RURITAGE experience of com-munity-based heritage management (De Luca et al., 2021).

The scenarios may include adaptive reuse of historic buildings, regeneration of public spaces, strengthening of local circular production chains, renewable energy communities, eco-logical restoration, soil regeneration, mobility improvements and service innovation. Within each scenario, the choice of materials and technologies is framed not only by technical performance, but also by compatibility with historic fabric, reversibility, durability, maintainability, use of local or recycled resources and contribution to territorial circularity. The Output of this phase is to develop a set of shared circular regeneration scenarios and related material and technological options (Fig. 1).

Phase 4 - Multidimensional and multicriteria evaluation represents the core of the methodology. Evaluation is conceived as a continuous process supporting decision-making

before, during and after implementation. The assessment framework is built around five domains:

- heritage and landscape quality;
- environmental and energy performance;
- circularity of resources and processes;
- socio-economic and community impacts;
- territorial and governance effects.

This structure is consistent with the literature on circular adaptive reuse and with cur-rent research on indicator systems for cities and regions. In the proposed framework, indicators and KPIs are not intended as purely technical metrics, but as tools for comparing scenarios, clarifying trade-offs and supporting shared decisions (Gravagnuolo et al., 2018). The output of this phase is to develop a comparative evaluation of scenarios and decision-support basis for the selection of priority interventions (Fig. 1).

Taken together, the four phases define an operational sequence in which territorial knowledge, diagnosis, participatory scenario-building and evaluation interact progressively to support regeneration choices across scales.

In operational terms, each phase is associated with a specific decision-support function: Phase 1 builds the knowledge base through territorial data and geo-cluster profiles; Phase 2 identifies constraints, values and degrees of transformability; Phase 3 translates diagnosis into shared regeneration scenarios; and Phase 4 compares these scenarios through multidimensional indicators and KPIs. This sequence is intended to make the framework usable both for research-oriented assessment and for preliminary planning processes involving public ad-ministrations, technical experts and local communities.

For these reasons, the framework should not be interpreted as a fixed protocol, but as a flexible decision-support structure that can be adapted to different territorial conditions. Future applications should test the methodology in diverse heritage-sensitive rural contexts, compare alternative indicator systems and integrate formal multicriteria and life-cycle-based tools. Such empirical testing will be necessary to verify the robustness of the framework and to assess its capacity to support implementation-oriented regeneration strategies.

At the same time, the framework has some limitations. First, it remains a theoretical and methodological proposal and has not yet been validated through a real case study. Second, the indicators and KPIs are defined at the level of evaluation do-mains and categories, but they require further specification according to local data availability, policy priorities and stakeholder needs. Third, the transferability of the framework depends on institutional capacity, quality of territorial data, and the effective involvement of local actors in participatory processes.

The proposed framework has potential applicability in historic villages, rural municipalities, cultural landscapes and networks of small settlements where regeneration requires the coordination of conservation, circular resource management, service provision and local governance. Its main added value lies in treating evaluation as a strategic and iterative process, capable of supporting the formulation, comparison and adjustment of regeneration scenarios rather than merely measuring final impacts.

4 Discussion

In this sense, the outcome of the research is not a single de-sign solution, but a structured methodological device that clarifies how circular regeneration can be framed, discussed and evaluated across scales. The framework makes explicit the connection between building-level actions and wider territorial regeneration processes, thereby addressing the ambiguity be-tween adaptive reuse and territorial development by treating them as complementary components of the same circular re-generation process.

The framework also identifies the main data categories re-quired for implementation: spatial and environmental data, building and conservation data, socio-economic data, governance and stakeholder data, and performance indicators. These data categories support a decision logic based on progressive refinement: from territorial profiling, to diagnosis, to scenario construction, and finally to the evaluation and prioritization of interventions.

The result can be summarized through five core dimensions. The first is the territorial dimension, which includes settlement morphology, cultural landscape structure, ecological infra-structures, accessibility, services and socio-economic net-works. The second is the heritage and building-stock dimension, which includes conservation conditions, construction systems, material compatibility, adaptive reuse potential and degrees of transformability. The third is the circularity dimension, which concerns local resources, material flows, energy opportunities, reversibility, durability, maintenance and the use of local or recycled materials. The fourth is the social and governance dimension, which includes residents, local administrations, associations, economic actors, technical experts and third-sector organizations. The fifth is the evaluation dimension, which translates these elements into indicators, KPIs and multicriteria comparison criteria.

As a theoretical and methodological paper, the main research result is the formalization of a multiscale decision-support framework for the circular regeneration of historic villages, cultural landscapes and rural territories. The framework is organized as a sequence of four interconnected phases that trans-form territorial knowledge into shared regeneration scenarios and comparative evaluation outputs.

5 Conclusion

This paper has proposed a multiscale methodology for the circular regeneration of historic villages, cultural landscapes and rural territories. The framework clarifies regeneration as an integrated territorial process in which adaptive reuse of historic buildings, conservation of cultural landscapes, local circular economies, ecological transition and community governance are considered interdependent components.

The contribution of the paper is mainly theoretical and methodological. Its main result is the formalization of a decision-support framework that connects territorial reading, integrated diagnosis, participatory scenario-building and multidimensional evaluation. In doing so, the paper addresses the gap between building-centred adaptive reuse approaches and broader territorial regeneration strategies.

The proposed framework may support researchers, planners and local administrations in structuring early-stage regeneration processes, identifying relevant data categories, involving actor groups and comparing alternative scenarios. However, its generalization requires caution. The framework must be adapted to local contexts, data availability and governance conditions, and should be tested through empirical applications in different rural and heritage-sensitive settings.

Future developments should therefore include the application of the methodology to case studies, the construction of con-text-sensitive indicator systems and the integration of formal multicriteria methods and life-cycle-based tools. These steps will be necessary to verify the operational effectiveness, transferability and limitations of the framework.

This paper has proposed a multiscale methodology for the circular regeneration of historic villages, cultural landscapes and rural territories. The framework treats evaluation as an on-going decision-support process for regeneration strategies rather than as a final analytical exercise.

The proposed methodology contributes to current research in three main ways. First, it translates circular city and circular region principles into heritage-sensitive rural contexts. Second, it connects building rehabilitation with territorial development, community participation and local circular economies. Third, it structures evaluation as a multidimensional and multicriteria process capable of supporting decision-making across scales.

Its relevance is therefore threefold: for research, because it addresses the gap between heritage conservation and circular territorial development; for design, because it supports the se-lection of compatible materials, technologies and scenarios; and for policy, because it offers a framework that can strengthen governance and align local regeneration with shared eco-logical and community goals.

Future developments should include the application of the methodology to case studies, the construction of context-sensitive indicator systems and the integration with formal multicriteria methods and life-cycle-based tools. At the same time, further testing will be necessary to verify the transferability of the framework across different heritage and rural con-texts.

In this perspective, the regeneration of historic villages can be understood not as sum of isolated building interventions, but as a coordinated territorial process in which heritage, land-scape, community and circular resource management are mutually reinforcing.

At the same time, future developments of the research will include the application of the proposed framework to a real case study within an ongoing national research project. In particular, the methodology will be tested in the context of the project "Italian historic regeneration through circular ecological heritage communities", funded by the Italian Ministry of University and Research (MUR) under the Projects of Relevant National Interest—National Recovery and Resilience Plan (PRIN PNRR). This application will allow for the verification of the operational effectiveness of the framework and its adaptability to different heritage-sensitive rural contexts.

Author Contributions. Conceptualization, A.S., M.A.; methodology, A.S.; validation, M.A.; formal analysis, A.S., M.A.; investigation, S.A.; re-sources, M.A.; data curation, A.S.; writing—original draft preparation, A.S., M.A.; writing—review and editing, A.S., M.A.; funding acquisition M.A. All authors have read and agreed to the published version of the manuscript.

Funding. This research was funded by Italian Ministry of University and Research (MUR), in the context of the Projects of Relevant National Interest—National Recovery and Resilience Plan (PRIN PNRR). Project Title: Italian historic regeneration through circular ecological heritage communities. Grant Number B53D23029490001.

References

Alonso, I.B., González, A.R., Fernández, V.M., et al.: Following the circular economy in European rural areas: missing links at municipal and macro scales. J. Environ. Manag. **370**, 122640 (2024)

Angrisano, M., Bottero, M., Cavana, G., Fabbrocino, F., Gravagnuolo, A., Fusco Girard, L.: Adaptive reuse of cultural built heritage: towards the implementation of the circular city model. Front. Built Environ. **11**, 1561982 (2025). https://doi.org/10.3389/fbuil.2025.1561982

Buglione, F., Della Morte, P., Angrisano, M., Gravagnuolo, A., Fusco Girard, L.: Circular models for the sustainable regeneration of Italian rural villages: a critical analysis of good practices toward the definition of a circular rural village. Sustainability. **18**(3), 1405 (2026)

Cavana, G., Angrisano, M., Gravagnuolo, A., Fusco Girard, L., Bottero, M.: Adaptive reuse of religious cultural heritage: the case study of the church "Santissimo Nome di Gesù" in Turin, Italy. In: Lecture Notes in Computer Science LNCS, vol. 15896, pp. 266–282. Springer, Cham (2026)

Cucco, P., Dell'Ovo, M., Oppio, A.: An evaluation model for adaptive reuse of cultural heritage in a circular economy perspective (2023)

De Fino, M., Scioti, A., Cantatore, E., Fatiguso, F.: Methodological framework for assessment of energy behavior of historic towns in Mediterranean climate. Energ. Buildings. **144**, 87–103 (2017). https://doi.org/10.1016/j.enbuild.2017.03.029

De Luca, C., Kouklios, I., Carbone, F., et al.: Participatory process for regenerating rural areas through heritage-led plans: the RURITAGE community-based methodology. Sustainability. **13**(9), 5212 (2021)

Dişli, G.: Circular economy in the heritage conservation sector: An analysis of circularity degree in existing buildings. Developments in the Built Environment (2023).

Egusquiza, A., et al.: Systemic innovation areas for heritage-led rural regeneration: a multilevel repository of best practices. Sustainability. **13**(9), 5069 (2021)

European Circular Economy Stakeholder Platform: Reading Guide on Circularity in the Built Environment. European Circular Economy Stakeholder Platform (2025)

European Commission / Circular Cities and Regions Initiative: Indicators on a Circular Economy for Cities and Regions. European Commission, Brussels (2024–2025)

European Commission: A New Circular Economy Action Plan: for a Cleaner and more Competitive Europe. European Commission, Brussels (2020)

European Commission: The Long-Term Vision for the EU's Rural Areas: Towards Stronger, Connected, Resilient and Prosperous Rural Areas by 2040. European Commission, Brussels (2021)

European Commission: Circular Cities and Regions Initiative (CCRI): Indicators and Frameworks for Circular Economy Implementation. European Commission, Brussels (2024)

Finamore, M., et al.: Circular economy in construction: Findings from a literature review. Heliyon [volume/article number needed] (2024)

Foster, G.: Circular economy strategies for adaptive reuse of cultural heritage buildings to reduce environmental impacts. Resour. Conserv. Recycl. **152**, 104507 (2020)

Gravagnuolo, A., Fusco Girard, L., Ost, C., Saleh, R.: Evaluation criteria for a circular adaptive reuse of cultural heritage. BDC. Bollettino del Centro Calza Bini (2017–2018)

Gravagnuolo, A., et al.: A participatory approach for "circular" adaptive reuse of cultural heritage: building a heritage community in Salerno, Italy. Sustainability. **13**(9), 4812 (2021)

Gravagnuolo, A., Varotto, M.: Terraced landscapes regeneration in the perspective of the circular economy. Sustainability. **13**(8), 4347 (2021). https://doi.org/10.3390/su13084347

Gravagnuolo, A., Angrisano, M., Bosone, M., Buglione, F., De Toro, P., Fusco Girard, L.: Participatory evaluation of cultural heritage adaptive reuse interventions in the circular economy perspective: a case study of historic buildings in Salerno, Italy. J. Urban Manage. **13**, 107 (2024)

Lombardo, G., Cicero, C.: Rehabilitation strategies for energy-efficient buildings in GeoCluster. Adv. Build. Energy Res. **9**(1), 34–54 (2015). https://doi.org/10.1080/17512549.2014.923326

Paoli, F., et al.: Indicators for the circular city: a review and a proposal. Sustainability. **14**(19), 11848 (2022)

RURITAGE Project: Community-Based Heritage Management and Planning Methodology. RURITAGE Project (2020–2021)

Scioti, A., De Fino, M., Fatiguso, F.: Sustainable reuse of vernacular architecture for rural tourism development. In: Proceedings of the ReUso Documentation, Restoration and Reuse of Heritage, Porto, Portugal (2022)

Sesana, M.M., Cuca, B., Iannaccone, G., Brumana, R., Caccavelli, D., Gay, C.: Geomapping methodology for the GeoCluster mapping tool to assess deployment potential of technologies for energy efficiency in buildings. Sustain. Cities Soc. **17**, 22–34 (2015). https://doi.org/10.1016/j.scs.2015.02.006

Suárez, M.A., et al.: Towards a methodology of good practices for heritage-led rural regeneration (2025)

Verardi, F., Angrisano, M., Fusco Girard, L.: New development policies for the internal areas of southern Italy: general principles for the valorization of rural areas in Calabria region. Valori e Valutazioni. **33**, 105–116 (2023)

Angrisano, M.: The canvas model to support the circular urban regeneration projects. In: Lecture Notes in Computer Science, LNCS, vol. 14108, pp. 483–497. Springer, Cham (2023)

Zhang, J.: Developing a dynamic evaluation framework for rural heritage regeneration from the perspective of rural revitalization (2025)

Cyber Intelligence and Applications (CIA 2026)

Enabling AI-Driven Cybersecurity Through Reproducible Cloud-Native Security Architectures

Gianni D' Angelo(✉), Carmine Ambrosino, Simona Ettari, Alfonso Petrosino, and Francesco Palmieri

Department of Computer Science, University of Salerno, Fisciano, Salerno, Italy
{giadangelo,fpalmieri}@unisa.it,
{c.ambrosino11,s.ettari,a.petrosino39}@studenti.unisa.it

Abstract. Artificial intelligence is increasingly used in cybersecurity for anomaly detection, behavioral modeling, and automated defense. However, the reproducibility of AI-based cybersecurity experiments depends not only on model code and datasets, but also on the infrastructure used to generate security data. In cloud-native environments, reverse proxies, storage backends, monitoring pipelines, timeout policies, and network configurations can shape traffic behavior and introduce hidden biases into detection models. This paper presents *local-lab*, an Infrastructure as Code-based architecture for building isolated and repeatable cloud-native security testbeds. The platform uses OpenTofu and `libvirt` to instantiate a stack composed of Nextcloud, a Caddy reverse proxy, MinIO object storage, and a monitoring pipeline based on Netdata, Prometheus, and Grafana. We use this environment to execute controlled Slow HTTP Denial-of-Service attacks, collect packet-level and system-level telemetry, construct the labeled SLOWDOS-2025 dataset, and evaluate neural detection models, including multilayer perceptrons, long short-term memory networks, and autoencoder-based configurations. The results show that the proposed infrastructure supports controlled observation of SlowDoS behavior across multiple layers and provides a structured setting for evaluating AI-based detection models. Beyond classification performance, the main contribution of this work is to treat the experimental infrastructure itself as a version-controlled scientific artifact, improving transparency, repeatability, and comparability in AI-driven cybersecurity research.

Keywords: Cloud-native security · Reproducibility · Local lab · Intrusion detection · Slow DoS attacks · Cybersecurity dataset

1 Introduction

Artificial Intelligence (AI) is widely used in cybersecurity for anomaly detection, behavioral modeling, and automated defense. Yet the reliability of AI-based

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 575–591, 2027.
https://doi.org/10.1007/978-3-032-30518-3_36

cybersecurity experiments depends not only on learning algorithms, model code, or released datasets, but also on the infrastructure in which security data are produced. In cloud-native environments, components such as reverse proxies, object storage backends, monitoring services, timeout policies, and network configurations can directly affect traffic behavior. Even small infrastructure changes may therefore alter the generated data and influence the evaluation of detection models. This raises a reproducibility problem. Many cybersecurity studies release datasets or source code, while the experimental infrastructure that generated the data is only partially described, manually configured, or left implicit. As a result, the environment becomes a hidden experimental variable. Differences in service deployment, connection handling, storage integration, and telemetry collection can introduce biases that are difficult to identify and make results harder to compare across studies. This paper investigates how cloud-native security experiments can be made more transparent and repeatable by treating the infrastructure itself as a scientific artifact. We present *local-lab*, an Infrastructure as Code-based architecture that enables consistent deployment and reconstruction of an isolated cloud-native stack composed of Nextcloud, a Caddy reverse proxy, MinIO object storage, and a monitoring pipeline based on Netdata, Prometheus, and Grafana. We evaluate the architecture through Slow HTTP Denial-of-Service (SlowDoS) experiments. SlowDoS attacks are well suited to this setting because they exploit timing, connection persistence, and resource management rather than high traffic volume. Their behavior is therefore strongly influenced by infrastructure-level factors such as proxy timeout policies, connection handling, and service configuration. In our experiments, benign and malicious traffic are collected from the deployed platform, transformed into flow-based features, and used to evaluate neural detection models, including multilayer perceptrons, long short-term memory networks, and autoencoder-based representations.

The study is guided by the following research questions:

- **RQ1:** Can cloud-native security architectures be formally specified and instantiated as version-controlled artifacts for reproducible cybersecurity experiments?
- **RQ2:** Can such an architecture support controlled observation of application-layer SlowDoS attacks across multiple infrastructure layers?
- **RQ3:** Can it provide a structured experimental environment for evaluating AI-based SlowDoS detection models?

The paper differs from prior work on cyber ranges, SlowDoS detection, stealthy cloud denial-of-service strategies, and machine learning-based intrusion detection. Its primary goal is not to introduce a new attack strategy, training scenario, or classifier. Instead, it focuses on the infrastructure conditions under which security data are generated and models are evaluated. By specifying the complete experimental environment as version-controlled Infrastructure as Code artifacts, the proposed approach makes the data-generation process more transparent, repeatable, and easier to inspect.

The main contributions of this paper are as follows:

1. We define a cloud-native security architecture based on Infrastructure as Code, where virtual machines, network topology, service deployment, monitoring components, and attack execution are specified as version-controlled artifacts.
2. We instantiate the architecture through the open-source *local-lab* project and use it to execute controlled SlowDoS experiments against a realistic Nextcloud-based application stack integrating Caddy and MinIO.
3. We generate the SLOWDOS-2025 dataset from packet captures collected during benign and malicious scenarios, using a documented pipeline for flow-based feature extraction, labeling, preprocessing, and balancing.
4. We evaluate multiple neural models, including MLP, LSTM, and autoencoder-based configurations, showing how the proposed environment supports AI-based SlowDoS detection studies under controlled experimental conditions.

The remainder of the paper is organized as follows. Section 2 discusses related work. Sections 3 and 4 describe the methodology and architecture. Section 5 presents the experimental results, limitations, and threats to validity. Section 6 concludes the paper.

2 Related Work

This work builds on several research areas, including cyber range infrastructures, low-rate and Slow HTTP denial-of-service attacks, machine learning-based intrusion detection, and dataset generation for security experimentation. Although these research directions have evolved significantly, existing studies rarely integrate reproducible infrastructure deployment, controlled attack execution, telemetry collection, dataset generation, and machine learning evaluation into a single end-to-end workflow.

2.1 Cyber Range Infrastructures and Reproducible Security Testbeds

Cyber ranges have emerged as important instruments for cybersecurity experimentation, training, and evaluation. Surveys of cyber range platforms highlight significant heterogeneity in architecture, deployment strategies, and experimental objectives [25]. Modern platforms increasingly rely on virtualization and automated orchestration to support dynamic scenario creation and infrastructure scalability [26]. Container-based approaches further improve portability and reproducibility by enabling lightweight deployment of complex service architectures through container orchestration frameworks [27]. Recent work has explored model-driven and automatically generated cyber ranges capable of synthesizing realistic infrastructures from high-level descriptions of systems and attack scenarios [28]. Other studies have investigated scalable cyber range architectures based on cloud technologies, enabling the deployment of large-scale experimental environments for cybersecurity research [29]. Despite these advances, most

cyber range systems are primarily designed for training, education, or red–blue team exercises. In contrast, our work treats the experimental infrastructure itself as a scientific artifact whose configuration can be versioned, reconstructed, and reused for machine learning-based cybersecurity experimentation.

2.2 Low-Rate and Slow HTTP Denial-of-Service Attacks

Denial-of-service attacks remain one of the most widely studied threats affecting networked systems. Low-rate denial-of-service attacks are particularly challenging because they exploit protocol dynamics and server resource management rather than high traffic volumes. Early work on low-rate TCP-targeted attacks demonstrated that carefully crafted periodic traffic patterns can degrade service availability while remaining difficult to detect through conventional monitoring mechanisms [30]. Slow HTTP attacks constitute a relevant subclass of these techniques. Slowloris maintains large numbers of partially completed HTTP requests that force servers to keep connections open indefinitely [31]. Similarly, slow HTTP POST attacks send request bodies at extremely low transmission rates, exhausting server-side connection pools. Surveys on low-rate denial-of-service detection highlight the difficulty of identifying these attacks because of their low traffic footprint and similarity to legitimate client behavior [32]. Experimental studies further show that SlowDoS effectiveness depends on infrastructure configuration, including reverse proxies, load balancers, and web application gateways [33]. Stealthy denial-of-service strategies in cloud computing have also been studied from the perspective of resource exhaustion under low-profile traffic patterns. Ficco and Rak propose a strategy for orchestrating slowly increasing attack patterns against cloud applications, with the objective of degrading service and increasing resource consumption while delaying detection [42]. This line of work is closely related to SlowDoS attacks because both exploit the interaction between application-layer behavior, resource allocation, and infrastructure configuration. Our work differs in its primary objective. Rather than proposing a new stealthy attack strategy or a dedicated detection algorithm, it focuses on the reproducibility of the infrastructure used to generate and evaluate security data. By combining Infrastructure as Code deployment, multi-layer telemetry collection, packet capture, dataset generation, and machine learning evaluation, the proposed framework enables SlowDoS behavior to be analyzed as an infrastructure-dependent phenomenon within a reproducible experimental workflow.

2.3 Machine Learning for Intrusion Detection

Machine learning techniques have been widely adopted for intrusion detection systems. Traditional IDS research initially relied on classical algorithms such as decision trees, support vector machines, and clustering-based anomaly detection [34]. More recently, deep learning models have shown strong potential

for detecting complex attack patterns in network traffic data. Recurrent neural networks and Long Short-Term Memory architectures can capture temporal dependencies in network flows [35], while autoencoders have been used for anomaly detection by learning compact latent representations of normal traffic behavior [36]. Several studies have specifically investigated deep learning approaches for detecting slow HTTP denial-of-service attacks. Flow-based classification models using statistical traffic features have shown promising results in distinguishing slow HTTP attacks from legitimate traffic [37]. Other work has explored feature selection and hybrid machine learning techniques for improving detection performance [38]. However, many of these studies rely on datasets whose generation processes are not fully documented, making it difficult to assess whether classification performance is influenced by the model itself or by characteristics of the dataset generation environment.

2.4 Dataset Generation and Feature Extraction for IDS Research

The reliability of machine learning-based intrusion detection depends heavily on the quality and realism of the datasets used during training and evaluation. Widely used benchmark datasets have been shown to contain limitations such as outdated traffic patterns, unrealistic attack scenarios, or incomplete documentation of the experimental setup used during data collection [39]. Recent work has therefore focused on improving dataset generation methodologies and traffic feature extraction pipelines. Flow analysis frameworks such as NTLFlowLyzer enable the extraction of statistical features from packet captures, supporting expressive representations of network traffic [23]. Other work has explored realistic network dataset generation through controlled cyber range environments [40]. Nevertheless, realistic IDS dataset generation remains challenging because it requires experimental environments capable of orchestrating complex service infrastructures, generating representative benign workloads, executing controlled attack scenarios, and collecting telemetry across multiple infrastructure layers.

2.5 Research Gap and Relation to the Research Questions

Existing research has addressed cyber range infrastructures, SlowDoS detection, stealthy cloud DoS strategies, machine learning-based intrusion detection, and dataset generation largely as separate research directions. Cyber range platforms mainly focus on training or red–blue team exercises. SlowDoS and stealthy DoS studies often analyze attack behavior or detection strategies without providing a sufficiently specified experimental infrastructure. Machine learning-based IDS studies frequently evaluate models on datasets whose generation process is only partially specified. This fragmentation creates a gap in experimental cybersecurity research: the lack of frameworks that jointly support infrastructure deployment, controlled SlowDoS experimentation, telemetry collection, dataset generation, and machine learning evaluation in a repeatable workflow. The proposed

Table 1. Qualitative positioning of related work with respect to the reproducible experimental pipeline addressed in this paper.

Reference	Infra.	Stack	SlowDoS	ML
Munaiah et al. [25]	P	–	–	–
Yamin et al. [26]	P	–	–	–
Nakata and Otsuka [27]	✓	P	–	–
Ficco and Rak [42]	–	P	✓	–
Aladaileh et al. [37]	–	–	✓	✓
Sharafaldin et al. [40]	P	P	P	✓
This work	✓	✓	✓	✓

Legend. A check mark (✓) indicates that the dimension is explicitly implemented and represents a central part of the work. The symbol *P* indicates partial support: the dimension is discussed, indirectly supported, or only partially implemented, but not specified in enough detail to allow direct reconstruction in the context considered here. A dash (–) indicates that the dimension is not addressed.
Dimensions. *Infra.* refers to an explicitly specified and reproducible experimental infrastructure, such as deployment artifacts, configuration files, or sufficiently detailed setup descriptions. *Stack* refers to a realistic multi-component application stack, rather than a single isolated service or a purely synthetic scenario. *SlowDoS* refers to explicit support for Slow HTTP, low-rate, or stealthy denial-of-service experimentation. *ML* refers to the training or evaluation of machine learning models on the generated or analyzed data. Assignments were made conservatively, based only on aspects explicitly described in the cited papers.

local-lab architecture addresses this gap by treating the experimental infrastructure as a version-controlled scientific artifact. Compared with existing cyber range platforms, the proposed work places less emphasis on training scenarios and greater emphasis on reproducible experimentation. Compared with prior SlowDoS and stealthy DoS studies, it analyzes attack behavior while explicitly specifying the infrastructure conditions under which traffic is generated. Compared with ML-based IDS studies, the contribution is not a new classifier, but a controlled and reproducible data-generation and evaluation pipeline. These aspects directly support RQ1, RQ2, and RQ3.

To improve transparency, Table 1 reports a qualitative positioning of related work according to four dimensions: reproducible infrastructure, realistic application stack, SlowDoS experimentation, and machine learning evaluation. The legend and assignment criteria are reported in the table notes. The table is therefore intended as a qualitative positioning of research scope, not as a ranking of scientific quality or effectiveness.

Overall, Table 1 shows that existing works usually address only part of the experimental pipeline. Cyber range studies emphasize deployment and training scenarios, SlowDoS studies focus on attack behavior or detection, and IDS dataset works mainly address data generation or model evaluation. In contrast, the proposed architecture integrates reproducible infrastructure deployment, controlled SlowDoS execution, telemetry collection, dataset construction, and neural model evaluation within a single reproducible workflow.

3 Materials and Methods

3.1 Reproducible Experimental Infrastructure

The experimental evaluation was conducted in a fully virtualized and isolated cybersecurity testbed instantiated through an Infrastructure as Code (IaC) workflow. This approach supports controlled experimentation by allowing the entire infrastructure configuration to be automatically deployed and consistently reconstructed. OpenTofu [8] orchestrates provisioning through the `libvirt` provider in a NAT-based private network (`10.0.2.0/24`). The infrastructure reproduces a simplified cloud-native deployment scenario while keeping the experimental environment fully contained.

Our `local-lab` includes three nodes:

1. **Target node:** hosts the application stack composed of Nextcloud behind a Caddy [1] reverse proxy with MinIO [2] as S3-compatible object storage.
2. **Monitoring node:** collects telemetry and system metrics.
3. **Attacker node:** generates malicious traffic for the attack scenarios.

This architecture enables simultaneous observation of application behavior, network traffic, and system-level resource consumption.

3.2 Infrastructure as Code for Reproducibility

Traditional cybersecurity testbeds are often assembled manually, for example by creating virtual machines, configuring networks, installing services, and applying security settings step by step. This process can introduce configuration drift and limit portability across research settings. To mitigate these issues, the proposed architecture adopts an Infrastructure as Code workflow in which topology, VM resources, networking, and service composition are explicitly described and version-controlled. The declarative OpenTofu-based workflow enables automated deployment and consistent reconstruction of the experimental platform.

3.3 Attack Scenarios and Execution Methodology

To evaluate the cloud-native service stack under resource exhaustion, SlowDoS attacks were simulated after deployment. Two tools were used:

1. Slowloris, which maintains multiple incomplete HTTP header requests to saturate server sockets.
2. slowhttptest [22], which generates Slow POST attacks by transmitting HTTP body data at very low rates.

The attack execution follows a workflow inspired by the Penetration Testing Execution Standard (PTES) [41], enabling controlled and systematic generation of malicious traffic.

3.4 Telemetry Collection and Traffic Capture

Two complementary telemetry channels were collected during the experiments:

1. **System and service monitoring:** Netdata collected fine-grained indicators, including TCP sockets, open connections, CPU usage, and memory usage, which were periodically scraped by Prometheus and visualized through Grafana dashboards.
2. **Packet capture:** packet-level traffic was captured on the target node using Wireshark or `tcpdump` under three operational conditions: benign usage, Slowloris, and Slow POST. These captures were used to generate the dataset for the machine learning analysis.

3.5 Dataset Construction

Packet captures were converted into flow-based tabular datasets using NTLFlowLyzer [23], which extracts statistical features from network flows. For each PCAP file, a configuration specifies the input file, output dataset, traffic label (Normal or SlowDoS), and processing parameters. The resulting labeled CSV files were merged into a unified dataset. Duplicate flows were removed, and balanced instances of benign and malicious traffic were combined to build the final dataset, SLOWDOS-2025.

3.6 Preprocessing and Transformation

The dataset contained missing and infinite values across multiple features.

Columns with more than 15% missing values and low relevance were removed, while the remaining missing values were imputed using the median. Infinite values were replaced with the median computed after excluding infinities. Moreover, the presence of `"not a complete handshake"` values in handshake-related features was leveraged through feature engineering by introducing a boolean indicator, `is_complete_handshake`. The preprocessed dataset was exported as `slowdos_2025_preprocessing.csv`.

A second dataset, `slowdos_2025_mix_preprocessing.csv`, was generated to evaluate a reduced feature subset inspired by prior work on slow DoS feature behavior. Finally, numerical features were standardized with `StandardScaler`, producing `std_slowdos_2025_preprocessing.csv`.

3.7 Neural Models

Three neural approaches were evaluated for binary classification (Normal vs. SlowDoS): an autoencoder, two MLP classifiers, and an LSTM sequence model. The autoencoder compresses the input feature space into latent representations of size 128, 64, or 16, which are subsequently used as inputs for supervised classifiers. Two MLP architectures were implemented: a compact baseline with 64–32 neurons and a deeper architecture with 128–64–32 neurons, both using ReLU activation, batch normalization, and dropout. To capture temporal dependencies in traffic flows, the LSTM model processes sequences constructed through a sliding window of length 50 and step 5. The architecture includes an LSTM layer with 32 units followed by dense layers and a sigmoid output. All models were implemented in TensorFlow/Keras and trained using the Adam optimizer.

3.8 Evaluation Protocol and Metrics

The experimental workflow integrates automated architecture deployment, SlowDoS attack execution, telemetry collection, and machine learning analysis within a reproducible cloud-native environment. Packet captures collected during the experiments are converted into flow-based datasets and preprocessed before being used to train and evaluate neural models under multiple feature configurations, including the full feature set, the FeatureMix subset, and autoencoder embeddings. To avoid potential data leakage, the train, validation, and test partitions were defined at the capture-session level, using PCAP files as the splitting unit. In particular, all flows extracted from the same PCAP capture were assigned to the same partition before constructing LSTM sliding windows. This protocol prevents overlapping or near-duplicate temporal windows originating from the same capture session from appearing in different splits. After the session-level split, LSTM sequences were generated using a sliding window of length 50 and step 5. The resulting partitions were used consistently for all evaluated models. Model performance was evaluated using Accuracy, Precision, Recall, Specificity, F1-score, Matthews Correlation Coefficient (MCC), and ROC-AUC, enabling comparison of neural architectures within the controlled experimental environment.

4 Proposed Architecture

This section describes the architecture implemented by the *local-lab* project [16]. The implementation is publicly available and version-controlled as part of the *local-lab* artifact. The architecture integrates four functional components: automated infrastructure deployment, controlled SlowDoS attack simulation, telemetry collection and dataset generation, and machine learning-based traffic analysis. By defining these components as version-controlled Infrastructure as Code artifacts, the environment can be reconstructed in a consistent manner for reproducible cybersecurity experimentation.

4.1 Local-Lab as a Reproducible Cloud-Native Security Architecture

Our *local-lab* defines the experimental infrastructure through an Infrastructure as Code workflow. Virtual machines, network topology, service deployment, and monitoring components are specified as version-controlled configuration artifacts and instantiated using OpenTofu [8] with the `libvirt` provider [18].

This design reduces configuration drift, supports consistent re-instantiation, and allows the infrastructure itself to be treated as a reusable scientific artifact. The deployed environment consists of three isolated nodes: a target node hosting the application stack, a monitoring node collecting telemetry, and an attacker node generating controlled malicious traffic. Figure 1 summarizes the architecture.

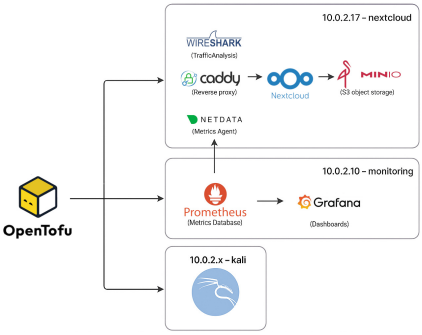

Fig. 1. Reproducible cloud-native architecture deployed using *local-lab* and OpenTofu.

4.2 Architectural Components and Node Roles

The architecture is composed of three nodes deployed within an isolated NAT-based network.

Target Node. The target node hosts Nextcloud behind a Caddy reverse proxy, with MinIO providing S3-compatible object storage. Packet-level traffic is captured using Wireshark or `tcpdump` [13], while Netdata collects system-level metrics such as CPU usage, memory consumption, open connections, and TCP socket states.

Monitoring Node. The monitoring node hosts Prometheus and Grafana. Prometheus scrapes metrics exposed by Netdata, while Grafana provides dashboards for observing system behavior during benign and malicious scenarios.

Attacker Node. The attacker node runs Kali Linux [15] and generates controlled malicious traffic against the target. Separating attacker, target, and monitoring roles enables realistic interaction patterns while preserving isolation from external networks.

4.3 Controlled Simulation of SlowDoS Attacks

Once the experimental platform is instantiated, controlled attack scenarios can be executed within the isolated network. Two representative SlowDoS techniques were evaluated:

- *Slowloris*, which maintains multiple incomplete HTTP header requests to exhaust server connection pools.
- *Slow POST*, generated using slowhttptest [22], which sends HTTP request bodies at extremely low transmission rates to keep server connections open for extended periods.

The attack execution process follows a methodology inspired by PTES [41], ensuring structured and systematic generation of malicious traffic. These attacks exploit protocol-level behavior rather than high bandwidth consumption, making them difficult to detect using traditional traffic-volume monitoring techniques. Within the proposed architecture, the attacker node generates malicious traffic toward the target node while the monitoring infrastructure records system behavior and network activity.

4.4 Telemetry Collection and Dataset Generation

A central goal of the proposed architecture is to enable systematic generation of cybersecurity datasets within a controlled and reproducible environment. Telemetry is collected through two complementary mechanisms. First, system and service metrics are captured using Netdata on the target node. Metrics such as TCP socket states, open connections, CPU utilization, and memory usage are periodically scraped by Prometheus and visualized through Grafana dashboards. Second, network traffic is captured at packet level using Wireshark or `tcpdump` under three operational conditions: normal application usage, Slowloris attack, and Slow POST attack. The resulting PCAP traces are processed using NTLFlowLyzer [23], which extracts statistical flow-level features from captured traffic. The extracted flows are labeled according to the corresponding operational scenario and exported as CSV datasets. The final dataset, SLOWDOS-2025, is obtained by merging balanced instances of benign and malicious traffic. Preprocessing includes duplicate removal, missing-value handling, feature engineering related to TCP handshake completion, and feature normalization using `StandardScaler`. By generating the dataset within a formally specified infrastructure, the resulting data remain tightly coupled with the experimental environment that produced them.

4.5 AI-Based Traffic Analysis

To evaluate intelligent detection within the proposed architecture, multiple neural network models were trained and evaluated on the generated dataset. An autoencoder was used to learn compressed latent representations of the traffic features, with embedding sizes of 128, 64, and 16 dimensions. These embeddings

were then used as inputs for supervised MLP classifiers. Finally, an LSTM network was introduced to capture temporal dependencies in network traffic flows. Unlike feed-forward models that operate on independent feature vectors, the LSTM processes sequences of flows using sliding windows, allowing the model to identify temporal patterns characteristic of slow denial-of-service attacks. The models were evaluated under different feature configurations, including the full feature set, a reduced feature subset, and autoencoder-generated embeddings. This setup enables analysis of the relationship between feature representation and detection performance.

4.6 Integrated Experimental Workflow

The overall experimental workflow enabled by the proposed architecture can be summarized as follows:

1. automated deployment of the cloud-native architecture using *local-lab*;
2. controlled execution of SlowDoS attacks within the isolated environment;
3. real-time monitoring and packet capture across infrastructure layers;
4. flow-based feature extraction and dataset generation;
5. data preprocessing and feature transformation;
6. training and evaluation of machine learning models.

By integrating reproducible infrastructure deployment, attack simulation, telemetry collection, and AI-based analysis, the proposed architecture provides a structured experimental framework for studying slow denial-of-service attacks in cloud-native environments.

5 Discussion and Limitations

5.1 Interpretation of the Experimental Results

Dataset Construction and Attack Effectiveness. The final dataset was obtained from PCAP captures collected under three operating conditions: benign traffic generated by a Nextcloud instance [9], and two SlowDoS attack scenarios produced using Slowloris [31] and slowhttptest [22]. Network traces were converted into tabular flow-based features through NTLFlowLyzer [23], resulting in a balanced binary dataset composed of Normal and SlowDoS traffic.

An important experimental observation concerns the different operational impact of the two slow-attack generators within the tested cloud-native stack. Slowloris connections were largely mitigated by the reverse proxy behavior, as Caddy [1] actively closes idle connections. In contrast, Slow POST attacks generated through slowhttptest were able to maintain partial requests and eventually produced service degradation. These observations show how the proposed architecture enables controlled observation of SlowDoS behavior across multiple infrastructure layers, including reverse proxy handling, application response, and network traffic characteristics. In this sense, the experimental setup demonstrates that a reproducible cloud-native security architecture can support systematic observation of application-layer SlowDoS dynamics and their operational impact, providing empirical evidence related to **RQ2**.

Role of Temporal Modeling with the Full Feature Set. When the full preprocessed feature set is used, the LSTM model achieves the best overall performance, with Accuracy 0.985, Precision 1.000, Recall 0.971, F1 0.985, and MCC 0.971, producing virtually no false positives while maintaining high sensitivity.

This result is consistent with the intrinsic characteristics of SlowDoS attacks. Rather than generating high traffic volumes, slow attacks exploit timing and persistence effects, such as incomplete requests or long-lived connections. Sequence-aware models are therefore able to capture dependencies between consecutive observations and recognize temporal attack patterns that static classifiers may partially overlook [35]. At the same time, the MLP configurations remain competitive, with MLP2 reaching AUC 0.992 and F1 0.966 on the full feature set. This indicates that aggregated flow statistics already encode a large portion of the discriminative signal extracted from network traffic [23]. The results therefore confirm that the dataset generated within the reproducible experimental environment enables systematic evaluation of AI-based detection models, supporting **RQ3**.

Impact of Feature Reduction. A consistent performance degradation is observed across all models when the FeatureMix dataset is used. For example, the LSTM accuracy decreases to 0.809 and the MCC drops to 0.618. This result suggests that the reduced feature subset discards information needed to represent the slow-attack dynamics within the evaluated cloud-native stack. SlowDoS attacks typically influence multiple correlated aspects of the communication process, including request completion patterns, connection duration, throughput behavior, and timing characteristics. When the representation becomes too sparse, classifiers lose complementary signals that jointly describe these behaviors.

Autoencoder Embeddings and Dimensionality Trade-Offs. The autoencoder-based dimensionality reduction provides insight into the amount of information required for reliable detection [36]. The results show a trade-off between compression and discriminative power. Latent representations of size 128 maintain high predictive performance, with the LSTM achieving an accuracy of 0.963 and an MCC of 0.927. A latent size of 64 offers an effective balance between compactness and accuracy, with MLP2 performing close to the best full-feature results. When the latent space is reduced to 16 dimensions, however, the loss of information becomes evident, indicating that aggressive compression removes fine-grained traffic characteristics required to avoid false positives and capture subtle attack behaviors.

Implications for Reproducible Cybersecurity Experimentation. Beyond the classification results, the experimental workflow demonstrates how cloud-native security architectures can be formally specified and instantiated as reproducible experimental environments.

The experimental infrastructure is defined through an Infrastructure as Code workflow based on OpenTofu and the `libvirt` provider [8,18]. The *local-lab* project [16] enables automated deployment of the same cloud-native stack used

throughout the experiments, allowing the infrastructure to be reconstructed in a consistent and version-controlled manner. Within this environment, monitoring and observation mechanisms operate across multiple layers. System and service metrics are collected through Netdata and exported to Prometheus and Grafana for time-series analysis [3,11,12], while network traffic is captured at packet level using Wireshark or `tcpdump` [13]. The combination of Infrastructure as Code provisioning, controlled traffic generation, and repeatable dataset extraction demonstrates that cloud-native security architectures can be instantiated as version-controlled experimental artifacts. This finding provides practical evidence related to **RQ1**.

Overall, the findings support the three research questions. The IaC-based deployment provides evidence for **RQ1**, the multi-layer observation of SlowDoS behavior supports **RQ2**, and the generated dataset enables the controlled evaluation of neural detection models required by **RQ3**.

5.2 Limitations and Threats to Validity

Operating System Scope. At present, the *local-lab* infrastructure [16] has been validated only on Linux-based hosts and guests. Extending the architecture to Windows or mixed operating-system environments would require additional base images, OS-specific configuration layers, and dedicated validation.

Scope and Generalizability. The evaluation focuses on a specific cloud-native stack composed of Nextcloud [9], Caddy [1], and MinIO [2]. Alternative reverse proxies, timeout policies, application stacks, or storage backends may affect cross-layer behavior, feature distributions, and model performance.

Attack Diversity. The experiments consider Slowloris [31] and slowhttptest-based Slow POST attacks [22]. Although these tools capture representative SlowDoS behaviors, additional application-layer and network-layer attacks should be considered in future work.

Dataset Realism and Workload Diversity. Benign traffic was generated through typical Nextcloud interactions, such as file uploads, downloads, and account operations. This workload is sufficient for controlled experimentation, but production environments exhibit greater variability in user behavior and traffic intensity.

Feature Extraction and Observability Limits. The detection pipeline relies on flow-based features extracted from PCAP traces using NTLFlowLyzer [23]. This supports network-level detection, but provides limited visibility into host-level and application-level signals such as process activity, application logs, or internal service states.

Statistical Robustness. The reported results correspond to a fixed experimental configuration and evaluation campaign, as shown in Table 2. Future work should exploit the repeatability of the platform to quantify run-to-run variability, perform additional repetitions, and report confidence intervals.

Table 2. Model performance for each feature configuration.

	ALL			FeatureMix			AE128			AE64			AE16		
Metric	MLP1	MLP2	LSTM	MLP1	MLP2	LSTM	MLP1	MLP2	LSTM	MLP1	MLP2	LSTM	MLP1	MLP2	LSTM
Accuracy	0.964	**0.965**	**0.985**	0.836	0.828	0.809	0.951	0.951	0.963	0.964	**0.965**	0.941	0.964	0.873	0.926
Precision	0.944	0.944	**1.000**	0.949	0.915	0.795	0.917	0.921	**0.985**	0.944	0.947	0.984	0.942	0.804	0.954
Recall	0.985	**0.988**	0.971	0.709	0.724	0.841	**0.991**	0.985	0.942	0.985	0.985	0.899	**0.988**	**0.988**	0.899
Specificity	0.942	0.942	**1.000**	0.962	0.724	0.776	0.910	0.915	**0.985**	0.942	0.945	**0.985**	0.938	0.758	0.955
F1-Score	0.964	**0.966**	**0.985**	0.812	0.808	0.817	0.953	0.952	0.963	0.964	**0.966**	0.939	0.965	0.887	0.925
AUC	0.990	**0.992**	0.987	0.936	0.911	0.894	**0.992**	**0.992**	0.984	0.989	**0.991**	0.982	0.987	0.990	0.957
MCC	0.928	**0.931**	**0.971**	0.722	0.6715	0.618	0.904	0.903	0.927	0.928	**0.931**	0.886	0.928	0.7672	0.855

Sequence Construction and Potential Data Leakage. The LSTM model uses overlapping temporal windows, which may introduce data leakage if windows originating from the same capture session are randomly distributed across training and test sets. To mitigate this risk, the split was performed at the PCAP capture-session level before window generation. All flows extracted from the same capture were assigned to a single partition, and sliding windows were then constructed independently within each partition. This protocol prevents overlapping or near-duplicate windows from the same capture session from being shared across different evaluation splits.

6 Conclusion and Future Work

This work examined how cloud-native security infrastructures can be specified, reconstructed, and reused to support controlled cybersecurity experimentation and AI-based detection model evaluation. The proposed *local-lab* architecture integrates Infrastructure as Code deployment, controlled SlowDoS attack generation, multi-layer telemetry collection, packet capture, dataset construction, and machine learning analysis within a unified experimental pipeline. The architecture was implemented using OpenTofu and the `libvirt` virtualization backend to instantiate a cloud-native stack composed of Nextcloud, Caddy, and MinIO, with monitoring provided by Netdata, Prometheus, and Grafana. Packet captures collected under benign, Slowloris, and Slow POST scenarios were processed through NTLFlowLyzer to generate the labeled SLOWDOS-2025 dataset. The results show that the proposed environment enables controlled observation of SlowDoS behavior across multiple infrastructure layers and provides a structured setting for evaluating AI-based detection models. In particular, the LSTM model achieved the best performance on the full feature set, confirming the relevance of temporal modeling for low-rate attacks characterized by persistent connections and delayed request completion. More broadly, the study demonstrates that cloud-native security architectures can be formally described and instantiated as version-controlled artifacts through Infrastructure as Code workflows. This design supports transparent, repeatable, and comparable cybersecurity experimentation, addressing the reproducibility limitations of studies that release only model code or datasets. Future work will extend the architecture to public cloud

platforms, multi-platform deployments, additional attack classes, and container-based or Kubernetes-based environments. These extensions will allow the proposed methodology to be evaluated under larger, more heterogeneous, and more realistic cloud-native scenarios.

References

1. Caddy Server: Caddy Reverse Proxy. https://github.com/caddyserver/caddy
2. MinIO: MinIO Object Storage. https://github.com/minio/minio
3. Netdata: Netdata Monitoring. https://github.com/netdata/netdata
4. Ylönen, T.: Secure Shell Protocol. RFC 4250 (1995). https://www.rfc-editor.org/rfc/rfc4250
5. Tenable Inc.: Nessus Vulnerability Scanner (2023). https://www.tenable.com/downloads/nessus
6. CVE Details (2025). https://www.cvedetails.com/vendor/15913/
7. OpenCVE (2025). https://app.opencve.io/cve/?vendor=nextcloud
8. Linux Foundation: OpenTofu (2025). https://opentofu.org/
9. Nextcloud GmbH.: Nextcloud (2025). https://nextcloud.com
10. Ambrosino, C.: Nextcloud with S3 Object Storage (2025). https://github.com/carmine-ambrosino/nextcloud-s3
11. Prometheus Monitoring System (2025). https://prometheus.io
12. Grafana Labs: Grafana (2025). https://grafana.com
13. Wireshark Team: Wireshark Network Analyzer (2025). https://www.wireshark.org
14. Oracle Corporation: VirtualBox (2025). https://www.virtualbox.org
15. Kali Linux (2025). https://www.kali.org
16. Ambrosino, C.: local-lab: Cloud Lab Locally (2025). https://github.com/carmine-ambrosino/local-lab
17. Lubuntu Project: Lubuntu (2025). https://lubuntu.me
18. dmacvicar: Terraform Provider for Libvirt (2025). https://github.com/dmacvicar/terraform-provider-libvirt
19. Red Hat, et al.: Fedora Workstation (2025). https://fedoraproject.org
20. Sarti, M.: How to run virtual machines with virt-manager. Fedora Mag. (2025). https://fedoramagazine.org/full-virtualization-system-on-fedora-workstation-30/
21. Ettari, S.: System Monitoring Stack (2025). https://github.com/simona2606/system-monitoring-stack
22. Shekyan, S.: SlowHTTPTest (2025). https://github.com/shekyan/slowhttptest
23. Shafi, M.M., Lashkari, A.H., Roudsari, A.H.: NTLFlowLyzer: towards generating an intrusion detection dataset and intruders behavior profiling through network and transport layers traffic analysis and pattern extraction. Comput. Secur. **148** (2025). https://doi.org/10.1016/j.cose.2024.104160
24. Fayyad, U., Piatetsky-Shapiro, G., Smyth, P.: From data mining to knowledge discovery in databases. AI Mag. **17**(3), 37 (1996)
25. Munaiah, N., Meneely, A., Eriksson, B., Williams, L.: A survey of cyber ranges and testbeds. Comput. Secur. **88**, 101636 (2019)
26. Yamin, M.M., Katt, B., Gkioulos, V.: Cyber range: the next generation cyber security training platform. Comput. Secur. **97**, 101961 (2020)
27. Nakata, T., Otsuka, A.: Cyber range architecture using container-based virtualization. In: Proceedings of the IEEE International Conference on Cyber Conflict (2018)

28. Tavakoli, M., Smith, P.: Automatic generation of cyber ranges for security experimentation. J. Cybersecur. **7**(1) (2021)
29. Kyriakopoulos, K., Theocharidou, M., Gritzalis, S.: Cyber range platforms and testbeds for cybersecurity research: a survey. IEEE Access **10**, 109512–109534 (2022)
30. Guirguis, M., Bestavros, A., Matta, I., Zhang, Y.: Low-rate TCP-targeted denial of service attacks and counter strategies. In: Proceedings of the IEEE International Conference on Network Protocols (2005)
31. Russell, R.: Slowloris HTTP Denial of Service Attack (2010). https://ha.ckers.org/slowloris/
32. Sivaguru, R., Nithyanand, R.: Detection and mitigation of low-rate denial-of-service attacks: a survey. IEEE Access **10**, 89764–89790 (2022)
33. Shekyan, S.: Experimental evaluation of slow HTTP denial-of-service attacks. Int. J. Inf. Secur. **17**(5), 497–510 (2018)
34. Javaid, A., Niyaz, Q., Sun, W., Alam, M.: A deep learning approach for network intrusion detection system. In: Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (2016)
35. Yin, C., Zhu, Y., Fei, J., He, X.: A deep learning approach for intrusion detection using recurrent neural networks. IEEE Access **5**, 21954–21961 (2017)
36. Sakurada, M., Yairi, T.: Anomaly detection using autoencoders with nonlinear dimensionality reduction. In: Proceedings of the MLSDA Workshop at ACM SIGKDD (2014)
37. Aladaileh, M., Anbar, M., Hasbullah, H.: Deep learning based detection of slow HTTP denial of service attacks. J. Netw. Comput. Appl. **150**, 102485 (2020)
38. Yu, S., Zhou, W., Doss, R., Jia, W.: Traceback of DDoS attacks using entropy variations. IEEE Trans. Parallel Distrib. Syst. **22**(3), 412–425 (2011)
39. Tavallaee, M., Bagheri, E., Lu, W., Ghorbani, A.: A detailed analysis of the KDD CUP 99 data set. In: IEEE Symposium on Computational Intelligence for Security and Defense Applications (2009)
40. Sharafaldin, I., Lashkari, A., Ghorbani, A.: Toward generating a new intrusion detection dataset and intrusion traffic characterization. In: Proceedings of the International Conference on Information Systems Security and Privacy (ICISSP) (2018)
41. Cloudflare, Inc.: Penetration Testing Execution Standard. http://www.pentest-standard.org/index.php/PTES_Technical_Guidelines
42. Ficco, M., Rak, M.: Stealthy denial of service strategy in cloud computing. IEEE Trans. Cloud Comput. **3**(1), 80–94 (2015). https://doi.org/10.1109/TCC.2014.2325045



LADDM: A Flexible Tool for Log Anomaly Detection

Gianni D'Angelo(✉), Raffaele Esposito, Mariano Caiazzo, and Francesco Palmieri

University of Salerno, Fisciano, Italy
giadangelo@unisa.it

Abstract. Log files keep track of everything that happens when services are used, and looking through them is important for finding any threats or problems that might come up because of mistakes. This paper looks at different cutting-edge ways to find log anomalies and presents a new one called LADDM (Log Anomaly Detection as a Docker Microservice). LADDM was made with a focus on the importance of logs at the enterprise level. To see how well it works, a synthetic dataset simulating enterprise logs was made. The solution is set up in a Dockerized environment, which makes sure it runs on its own and doesn't depend on any other services.

This approach not only highlights the importance of efficient log analysis in safeguarding enterprise operations but also demonstrates the practicality and scalability of containerized solutions in addressing security challenges.

Keywords: Anomaly Detection · Log Analysis · Machine Learning · Deep Learning · Log Parser

1 Introduction

The Clusit report [1] - Report on Cybersecurity in Italy and in the world - shows that in 2024 incidents worldwide had increased by 36% compared to 2023. The global trend for the first half of 2025 shows a further growth, equal to +36% compared to the previous half of the year.

In the first half of 2025, not only did the frequency of incidents increase, but also their average severity. In 2024, incidents with a "Critical" or "High" impact accounted for 77% of the total. In the first half of 2025, the impact increased further compared to 2024 (82% of incidents with a critical or high severity). The sector with the highest number of cyber incidents in Italy is the Government/Military/Law Enforcement sector (38% of the total), with a 279% increase compared to the entire previous year.

Check Point Research, in January 2026, released the Global Threat Intelligence Report [2] which shows that Italy records 2,400 cyber attacks per week, significantly higher than the global average. Europe is highly vulnerable from a cyber perspective, paying a very high price. New research conducted by Howden [3], a global insurance broker, has estimated the losses suffered by the four major European economies—Italy, France, Germany, and Spain—due to cyber attacks between 2020 and 2025 at €307

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 592–606, 2027.
https://doi.org/10.1007/978-3-032-30518-3_37

billion. The resulting situation is alarming and we must do everything possible to prevent and block these attacks by limiting and reducing the risks and resulting damage.

In recent years, cybersecurity systems based on log anomaly detection have been proliferating. These systems identify logs that highlight unexpected and potentially dangerous activity, allowing timely intervention to prevent or reduce system damage and avoid negative cascading effects.

Logs are chronological and sequential records of operations performed by a user or computer system. Log files are stored in systems and remain available for monitoring and administration. They are essential for IT security as they allow the tracking and recording of activities performed by users or other systems, and are often used in the event of litigation.

Each type of log plays a specific role in managing and monitoring IT systems, enabling system administrators to identify and resolve problems, ensure data security, and monitor user and application activity.

Specifically, applications also generate logs when the system enters an unexpected or undesirable state, there are availability issues, or security incidents. We can use these logs, which highlight unexpected and potentially dangerous activities, to take timely action to prevent or mitigate system damage and avoid negative cascading effects.

A complex system typically produces a huge amount of logs. For example, Alibaba Inc.'s cloud computing systems produce approximately 30–50 gigabytes (about 120–200 million lines) of tracking logs per hour, typically in the form of unstructured text messages [4].

It is essential to adopt anomaly detection techniques capable of automatically learning patterns describing normal system behavior, so that any deviations from these patterns can be identified and flagged as potential malicious activity requiring operator intervention. Anomaly detection is a methodology used to identify atypical events or data within a data set.

2 Background

2.1 Anomaly Detection

Anomalies are patterns in data that do not conform to a well defined notion of normal behavior [5] (Chandola et al., 2009).

Anomaly detection is a process used to identify behaviors that significantly deviate from expectations [5] (Chandola et al., 2009). In other words, it involves finding exceptions or deviations from the norm within a dataset. This process is fundamental in many fields and applications, such as cybersecurity, fraud detection, predictive maintenance, and system health monitoring.

Identifying real anomalies is not easy because false positives or data noise are often found. A balance must be struck between identifying too many anomalies (false positives) and failing to identify real anomalies (false negatives). This activity is important for preventing system security attacks and serves to optimize decision-making processes by anticipating and mitigating potential risks, improving operational efficiency.

There are two main ways to find anomalies: outlier detection and novelty detection.

In the first instance, identifying extreme or atypical values that deviate from the normal distribution of data frequently signifies measurement inaccuracies, extraordinary variations, or potential security risks.

In the second case, novelty detection focuses on the discovery of patterns or instances that have not been previously observed in the training dataset and may indicate system evolution or the presence of new scenarios or threats.

2.2 Log Anomaly Detection Techniques

Log anomaly detection techniques can vary depending on the specific needs and characteristics of the system under consideration. However, they generally involve analyzing patterns in log data using machine learning algorithms, statistical analysis, or rule-based anomaly detection techniques. By applying advanced data analytics, artificial intelligence, and machine learning techniques, cybersecurity experts and system administrators can promptly identify significant deviations from normal system behavior and take appropriate action.

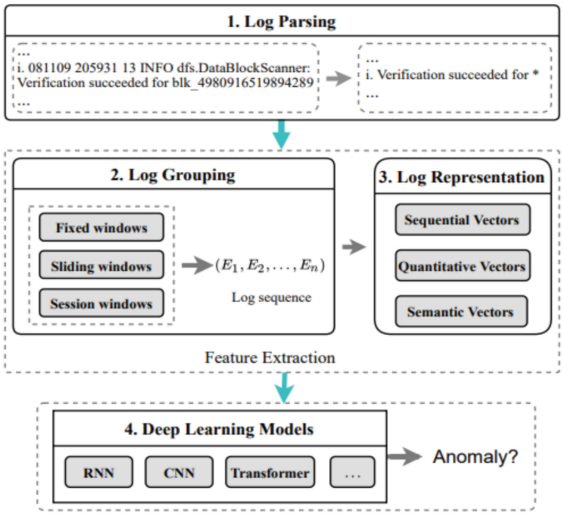

Fig. 1. Workflow Deep Learning models [4]

Figure 1 illustrates the typical workflow adopted by the most widely used deep learning approaches for log anomaly detection [4]. This process is structured into four stages:

1. **Log Parsing**. Log parsers process raw log entries at first to get log templates. This step makes the data uniform by turning each log message into a structured template. This lets you tell the difference between constant elements and variable parameters. The logs are now easier to read and better for further analysis.
2. **Log Grouping**. After being parsed, the logs are put into groups or sessions using methods like fixed windows, sliding windows, or session-based windows. This step

organizes the data so that it shows event sequences or operational sessions. This is important for finding temporal relationships between log entries.

3. **Log Representation**. The log sequences are turned into numbers after they are grouped. These representations can hold sequential, quantitative, or semantic information, depending on what you want to learn from the analysis. To make the data work with deep learning algorithms, you need to change logs into vector form.
4. **Deep Learning Models.** In the last step, the numerical vectors are used to teach deep learning models to tell the difference between normal and abnormal log sequences. These models learn about patterns and trends in the data through training. This lets them find problems by using their ability to model complex relationships.

3 Related Work

Several deep learning–based approaches have been proposed for detecting anomalies in system logs [4]. One of the most notable methods is DeepLog, which employs an LSTM model to learn the normal execution patterns of a system by analyzing sequences of log events. The model predicts the next log event based on previous ones, and an anomaly is detected when an incoming event does not match the predicted outcome. This approach has demonstrated high effectiveness on the HDFS dataset, achieving an F-measure of 0.96.

DeepLog [6] were the first to propose an approach based on deep learning that enables detection of anomalous event sequences in log data.

Another related approach is LogAnomaly [7], which combines log count vectors with an LSTM model. In this method, log templates are represented as semantic vectors generated through a technique based on synonym and antonym relationships. The model predicts the next log event and flags any event that deviates from the predicted result as anomalous. The experimental results are very good, with F-measures of 0.95 on the HDFS dataset and 0.96 on the BGL dataset.

PLELog [8] deals with the problem of having too little labeled data. It uses a probabilistic label estimation method to make up for missing annotations and an attention-based GRU neural network to find strange patterns in log sequences. The model sorts sequences into two groups: normal and strange. It has shown to work better than other semi-supervised methods on both the HDFS and BGL datasets.

LogRobust [9] is another method that focuses on making logs more stable and less prone to change. It uses TF-IDF weighting and a pre-trained Word2Vec model (FastText) to learn vector representations of log templates. An attention-based Bi-LSTM model then processes these vectors, which makes the system better at finding problems even when log events are different.

There are methods that use Convolutional Neural Networks (CNNs) [10]. In this case, logs are put into sessions and turned into a trainable matrix that the CNN model can use as input. The network learns to tell the difference between normal and strange log sequences. This method worked especially well on the HDFS dataset, getting an F-measure of 0.98.

The difficulties in enhancing the efficacy of deep learning models for log anomaly detection indicate potential avenues for future research and development.

The study of these typical models of the contemporary state of the art has elucidated the obstacles in enhancing their efficacy in detecting log anomalies:

- Training data selection: The way you choose your training data can have a big effect on how well your model works. Random selection can cause data leakage problems, where the training set may have data from the future, making the models less useful in real life.
- Log data clustering: the right window size for log clustering is critical. Wrong sizes can make it difficult to capture anomalies
- Unbalanced class distribution: An unbalanced class distribution, where anomalous log sequences make up a small part of the dataset, can make it much harder to find anomalies
- Data quality: Noise in the data, like false positives and false negatives that happen during the manual labeling process, can make the model work worse. Also, errors in log parsing can add wrong log events, which makes it harder to find anomalies.

Deepcase [11] is an innovative system designed to enhance cybersecurity monitoring. It uses deep learning techniques to analyze security events within their context, with the goal of reducing the workload of security operators by automating the correlation of security events.

This approach allows the system to filter a large portion of events, decreasing the need for manual analysis while improving the accuracy of threat detection. Additionally, it provides detailed explanations of why certain events are classified as malicious.

The deep learning model is made to find and study complicated relationships between sequences of events, both new and old. The model uses an advanced similarity function to group event sequences and show why different sequences are thought to be similar. This method lets security operators focus on a small number of representative sequences for each group. This makes it easier for them to decide what to do and cuts down on their work by a large amount.

LogAnomaly and Deeplog are just two of the many tools used to find strange things in log data. But the current methods have a big problem: they use predefined log patterns to train their models, and these patterns usually only work well in environments that are stable and predictable.

Isolation Forest [12], Random Forest [13], K-Nearest Neighbors [14], density-based clustering [6], and Template-to-Vec [7] are some of the most popular clustering-based methods. These methods frequently encounter difficulties associated with restricted accuracy, elevated false positive rates, and significant computational expenses.

Du et al. [4] present a technique that integrates a log parser with an LSTM network to autonomously detect patterns during standard system operations. This method is adaptable and can be used in a variety of situations. However, it has a big problem: it has trouble finding anomalies when log patterns change a lot from what it learned during training.

Guo et al. [15] introduced BERT, a model that learns sequence patterns using two new self-supervised learning methods. This model does a good job of finding differences between log sequences that are expected and those that are not. But it can't generalize well when it sees data that is very different from what it saw during training.

Liu et al. [16] use a self-attention neural network to find problems and help people make decisions. One of the best things about this method is that it uses data augmentation techniques to set decision thresholds using only the normal training data that is already available.

Nipa et al. [17] performed a comparative analysis of current machine learning and deep learning methodologies for anomaly detection by assessing their efficacy in identifying anomalies across various metrics and efficiency criteria.

The results show the best way to find problems in this system, which lets you do maintenance ahead of time and makes the system more reliable overall.

4 Our Solution

We came up with the name LADDM (Log Anomaly Detection as Docker Microservice) for our system. Because it is built on containers, it is a log anomaly detection system that is made for business settings.

Containers are standardized packages that hold an application's code and all the parts it needs to work. This makes it possible for the app to run quickly and reliably in any setting, no matter what operating system, libraries, or runtime versions are installed.

Virtualization focuses on the operating system instead of the hardware when using containers. Each container has only the application and the libraries it needs to work, so it doesn't need a full guest operating system. Because of this, containers are much lighter than regular virtual machines, which makes them faster and easier to move around.

We made a container-based architecture using Docker [18], which is an open-source platform. Docker lets applications run separately from the infrastructure they are built on, which makes it easy and quick to deploy software. Docker makes sure that development, testing, and production environments are all the same by putting an application and all of its dependencies into a standard container that can run on any system with Docker installed. The client-server architecture, in which the client interacts with the Docker daemon, which manages containers, images, and volumes, makes resource management more efficient and keeps processes separate.

In this case, one container acts like the business environment by making logs at regular intervals, while the other container uses and tests the anomaly detection models, separating and reporting logs that are out of the ordinary. LADDM's modular design gives users a lot of freedom. They can easily add or remove anomaly detection models by just adding the new model definition to the system.

Furthermore, LADDM supports integration with **DeepCase** and leverages models such as **IsolationForest** with semantic vector transformation and **RandomForest** with pattern recognition. To validate LADDM, a synthetic dataset designed to reflect an enterprise environment was used, where each record follows the format:

$$\text{"Action"}\ \textit{generated}\ \text{"}x\text{"}\ \textit{bytes from User}\ \text{"}Y\text{"}$$

In this context, "Action" indicates the operation performed by user "Y," resulting in the generation of "x" bytes. To identify anomalous logs, each action in the dataset produces a quantity of bytes that varies randomly within a predetermined range for that specific action. A log is considered anomalous when the volume of data generated by an

action is N times greater than the norm, thus providing a clear and quantifiable criterion for distinguishing anomalies within the dataset.

LADDM Structure. The operation of LADDM is described in Fig. 2.

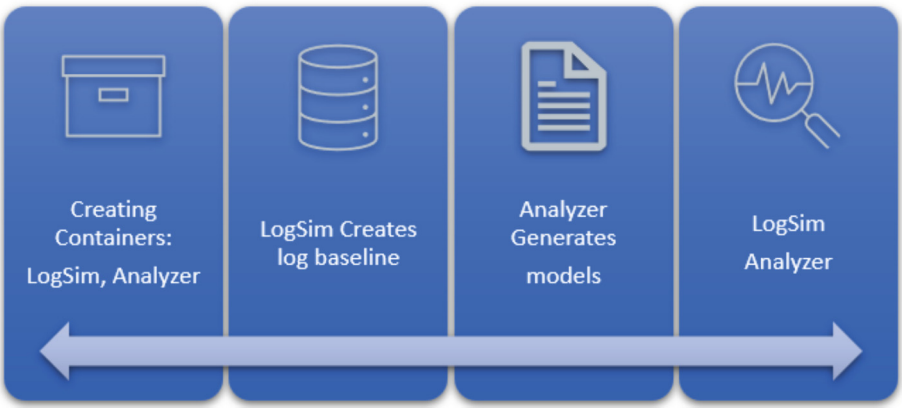

Fig. 2. LADDM Structure.

LADDM operates through a well-structured process that begins with the establishment of two Docker containers, named **LogSim** and **Analyzer**, each configured with a dedicated dockerfile detailing the dependencies required for their operation.

The **LogSim** container performs two crucial tasks: first, it generates an initial set of logs that will serve as a baseline for training anomaly detection models, and then it simulates the operation of an application by producing periodic logs.

LADDM takes on the role of model management: it acquires the log baseline, trains the anomaly detection models, and, if necessary, creates new ones. Next comes the monitoring phase: LADDM allows you to select specific log ranges to analyze and decide which models to use for the scan. If all available models are used, the analysis proceeds in parallel, optimizing the efficiency and effectiveness of the anomaly detection process.

5 Experiments

A Log Anomaly Detection Tool can become a key element for cybersecurity. Experimental evaluation is a crucial part of the development and validation of such systems and allows us to understand their effectiveness and validity, assessing their reliability, especially when dealing with complex and highly heterogeneous data. In this work, the evaluation is performed using two complementary types of datasets: a synthetic one, generated through log simulation, and a real one, HDFS_v1. This methodological choice combines the advantages of experimental testing with the benefits of analyzing realistic scenarios.

The synthetic dataset allows for a controlled environment where anomalies can be introduced explicitly and parameterizably. This approach allows for systematic variation in the type, frequency, and distribution of anomalies, thus offering the ability to analyze system behavior under specific conditions and validate design assumptions.

The HDFS_v1 dataset is a standard dataset in the literature for evaluating log anomaly detection techniques, and is extracted from Hadoop Distributed File System (HDFS) logs. It contains millions of log lines generated by a real cluster and organized into sequences associated with specific operations (e.g., read or write requests to data blocks). The dataset has an important feature: it has sequence-level labels that identify each session as normal or anomalous based on the outcome of the operation. Anomalies are typically caused by system failures or operational errors, so the dataset is highly representative of real-world scenarios. Furthermore, HDFS_v1 exhibits high structural variability, noise in log messages, and a highly skewed distribution between normal and anomalous samples, which increases its complexity and makes it a valuable benchmark for evaluating the robustness and generalization of models.

The use of synthetic and real datasets in the experimental evaluation phase allows for a more complete and reliable analysis, helping to highlight the strengths and limitations of the proposed log anomaly detection system.

The first step is to generate two containers: **LogSim**, which is oriented towards business operations, and **Analyzer**, which focuses on anomaly detection models and the ability to interact with **LogSim**. The docker file used is shown Fig. 3 .

```
1. FROM      python:3.10-slim
2. RUNapt-get update
3. && apt-get install -y --no-install-recommends \
4. build-essential \
5. libssl-dev \
6. libffi-dev \
7. python3-dev \
8. bash \
9. && rm -rf /var/lib/apt/lists/*
10. COPY      ./requirements.txt .
11. RUN       pip install --upgrade pip && \
12. pip install -r requirements.txt
13. RUN       rm ./requirements.txt
```

Fig. 3. Docker file.

To optimize the management and control of these containers, Docker Compose was adopted. This YAML configuration file allows you to orchestrate the startup and coordination of services with a single command, significantly simplifying the administration of the associated components, networks, and volumes.

After the containers have been created and launched, the next step is to develop a log baseline, which will serve as the training dataset for the anomaly detection models. A dedicated log generation function produces random logs following the structure outlined

in the introduction. Specifically, one thousand logs were generated incorporating four distinct action types

—"LOGIN," "LOGOUT," "CREATE," and "EDIT"—

performed by a set of twenty-six different users, ranging from "User_A" to "User_Z."

Each action type is associated with a defined range of generated bytes. Normal logs are marked with a "1," while anomalous logs, created with a 10% probability, are marked with a "-1." A log is considered anomalous when the number of bytes generated exceeds N times the typical range for that specific action.

The generated data is organized in a CSV file, structured into four main columns:

- "event" describing the log,
- "machine" indicating the log source (in this case, a single machine),
- "timestamp" recording the time the log was generated,
- "label" classifying the log as normal or abnormal.

Fig. 4 shows an example of what a generated file looks like.

```
1. event,machine,timestamp,label
2. Login generate 73 bytes performed by USER_C,1,1705949473,1
3. Edit generate 2000 bytes performed by USER_P,1,1707677473,1
4. Login generate 56 bytes performed by USER_V,1,1706122273,1
5. Edit generate 6528 bytes performed by USER_X,1,1707591073,1
6. Create generate 54290 bytes performed by USER_R,1,1707677473,1
7. Create generate 89030 bytes performed by USER_M,1,1706554273,1
8. Edit generate 2451 bytes performed by USER_T,1,1707331873,1
9. Logout generate 279 bytes performed by USER_V,1,1707159073,1
10. Edit generate 5159 bytes performed by USER_U,1,1706381473,1
11. Logout generate 819 bytes performed by USER_E,1,1706813473,1
12. Login generate 51 bytes performed by USER_B,1,1707331873,1
13. Logout generate 808 bytes performed by USER_V,1,1706035873,1
14. Logout generate 56700 bytes performed by USER_N,1,1705863073,-1
15. Login generate 21 bytes performed by USER_W,1,1705690273,1
16. Login generate 89 bytes performed by USER_W,1,1705603873,1
17. Login generate 4 bytes performed by USER_P,1,1706554273,1
```

Fig. 4. Example of baseline.

After generating the baseline file in CSV format, the next step is to develop the anomaly detection models. For each selected model, a dedicated directory is created to store the developed models. In this particular analysis, three separate directories were set up for the various models:

- "deepcase" for the DeepCase model,
- "if" for Isolation Forest, and
- "rf" for Random Forest.

Once the model generation is complete, you can activate the **Analyzer** scanning function.

All models available in their respective directories are run simultaneously, analyzing one log per second. If a model identifies an anomaly, it is logged in a file called "anomalies.csv" (Fig. 5), which documents the affected log and the name of the model that detected the anomaly.

```
1. event,machine,timestamp,label,models
2. Edit generate 959800 bytes performed by USER_G,1,1705525458,-1,IF
3. Create generate 13280 bytes performed by USER_Z,1,1705266259,1,IF
4. Login generate 5400 bytes performed by USER_F,1,1706907927,-1,RF
5. Edit generate 362100 bytes performed by USER_V,1,1706907935,-1,RF
6. Logout generate 75300 bytes performed by USER_J,1,1706130343,-1,RF
7. Edit generate 659400 bytes performed by USER_G,1,1705611976,-1,RF
8. Edit generate 384500 bytes performed by USER_V,1,1707599185,-1,RF
9. Login generate 3600 bytes performed by USER_H,1,1705957591,-1,RF
10. Logout generate 57500 bytes performed by USER_L,1,1706735195,-1,RF
11. Login generate 9600 bytes performed by USER_J,1,1706389598,-1,RF
12. Edit generate 603400 bytes performed by USER_Y,1,1705439249,-1,RF
```

Fig. 5. File anomalies.csv.

6 Results

For this synthetic dataset, the DeepCase model did not perform satisfactorily, demonstrating significant uncertainty in anomaly detection. Conversely, the Isolation Forest model, enriched with semantic transformation, performed better than DeepCase, although not without some ambiguities (Fig. 6). This latter finding is supported by the confusion matrix, which reveals a significant number of false negatives, indicating that there were several cases where the model did not correctly identify anomalies in the dataset.

The Random Forest model was the most effective, excelling both in the training evaluation phase and in real-time practical application, proving particularly suitable for analyzing this dataset. It demonstrated a high ability to accurately identify anomalous patterns in the logs, confirming its optimality as an anomaly detection tool in this specific context (Fig. 7).

6.1 HDFS_v1 Dataset

HDFS_v1 is a collection of logs produced in a private cloud environment using benchmark workloads and manually labeled using rules to identify anomalies. The dataset consists of 575,062 log sequences, of which 558,319 are normal (97.1%) and 16,743 are anomalous (2.9%). Pre-processing steps included grouping the log entries by block ID and filtering incomplete tracks. Other than the provided models, no further feature engineering technique has been applied.

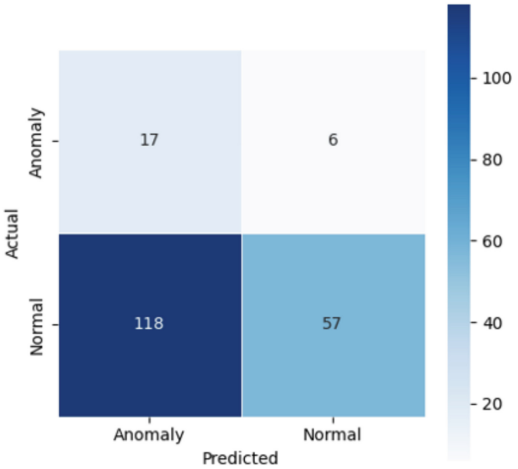

Fig. 6. Confusion matrix IF.

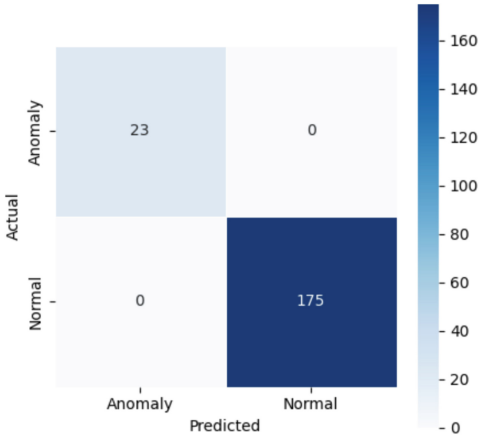

Fig. 7. Confusion matrix RF.

The logs are divided into traces based on block IDs. Each trace associated with a specific block ID is assigned a ground truth label: normal/anomaly.

For example

Record dataset HDFS_v1:

BlockId, Label, Type, Features, TimeInterval, Latency

blk_-3102267849859399193, Fail, 16, "[E5,E5,E22,E7,E11,E10,E14,E7]", "[0.0, 0.0, 60.0, 1.0, 0.0, 657.0, 0.0]", 718

where:

BlockId = "blk_-3102267849859399193"

Label = "Fail"

Type = "16"

Features = "[E5,E5,E22,E7,E11,E10,E14,E7]"
TimeInterval = "[0.0, 0.0, 60.0, 1.0, 0.0, 657.0, 0.0]"
Latency = 718

We used LADDM with the HDFS_v1 dataset. Deepcase's performance significantly outperforms the synthetic dataset (Fig. 8). This was expected, considering that HDFS is a frequently cited and recommended dataset in the Deepcase documentation, indicating potential optimization or specific adaptation of the model for this type of data.

	precision	recall	f1-score	support
0	0.6090	0.1832	0.2817	50178
1	0.8451	0.5529	0.6685	178358
2	0.5898	0.8079	0.6818	149575
3	0.8195	0.9997	0.9006	729051
4	0.6411	0.9057	0.7507	1124
5	0.9950	0.9873	0.9911	727926
6	0.9106	0.9366	0.9234	727926
7	0.8625	0.4791	0.6160	1100
8	0.9628	0.8474	0.9014	1101
9	0.0000	0.0000	0.0000	97
10	0.9990	0.9997	0.9993	593633
11	0.9994	0.3402	0.5077	242643
12	0.8905	0.9529	0.9207	592612
13	0.7429	0.3333	0.4602	1101
14	0.9382	0.9188	0.9284	730014
15	0.0000	0.0000	0.0000	2
16	0.0000	0.0000	0.0000	5
accuracy			0.9039	4726446
macro avg	0.6944	0.6026	0.6195	4726446
weighted avg	0.9110	0.9039	0.8946	4726446

Fig. 8. Classification report Deepcase - HDFS dataset

The confusion matrices for the Isolation Forest (Fig. 9) and Random Forest (Fig. 10) models show similar results, characterized by high levels of performance, reflecting the effectiveness of these machine learning techniques in correctly identifying anomalies in the dataset under study.

To improve readability and comparability, we report precision, recall, F1-score, and accuracy for Isolation Forest and Random Forest on both the synthetic dataset and HDFS_v1. These metrics are derived from the confusion matrices shown.

In Table 1 it is clear that the Random Forest algorithm outperforms Isolation Forest. In fact, Random Forest demonstrates a greater ability to detect anomalies, achieving significantly fewer false positives, which indicates improved accuracy in identifying irregularities.

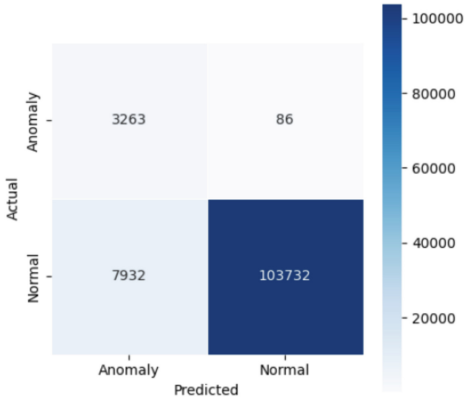

Fig. 9. Confusion matrix IF - HDFS dataset

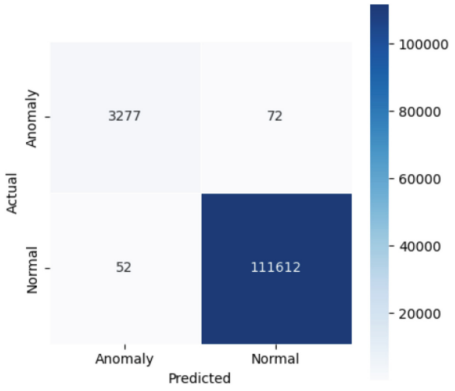

Fig. 10. Confusion matrix RF - HDFS dataset

Table 1. Results

Dataset	Model	Precision	Recall	F1-score	Accuracy
Synthetic	Isolation Forest	0,13	0,74	0,22	0,37
Synthetic	Random Forest	1	1	1	1
HDFS_v1	Isolation Forest	0,29	0,97	0,45	0,93
HDFS_v1	Random Forest	0,98	0,98	0,98	1

7 Conclusions

LADDM was conceived as a flexible tool that can be integrated into business processes to monitor services, leveraging a modular architecture of anomaly detection models. Analyzing the use of the "synthetic" dataset, the Random Forest model showed the best

performance, with similar results also obtained with the "hdfs" dataset. However, unlike the first "hdfs" dataset, it also showed good confidence with the deepcase model. It is therefore essential to recognize the possibility of selecting the most suitable algorithm or algorithms for each specific dataset.

LADDM's unique feature is its ability to interchange or replace anomaly detection models, allowing for tailored adaptation to specific enterprise datasets or the introduction of new, more advanced methodologies.

This flexibility opens the door to many future developments, such as the testing of new models or the handling of larger datasets that show the unusual behavior or real-world attacks that businesses deal with every day. As for future plans, the goal is to make LADDM a fully autonomous tool that runs in the cloud and is hosted in a distributed way. This would make it possible to handle a lot of data and create and host proprietary anomaly detection models, which would make sure that everything works well and is reliable.

To come up with solutions that are not only cutting-edge but also sustainable and scalable, LADDM development must be open and proactive, welcoming new challenges and adapting to change. These solutions must be able to change as the market and social needs change. In conclusion, LADDM could be both a new technology for finding anomalies and a way to make business processes more efficient and change things. LADDM could become a support for data security and analytics in the modern business world if it keeps looking into and using new technologies.

References

1. Rapporto Clusit: (2026). https://clusit.it/rapporto-clusit/. Accessed 15 Mar 2026
2. Check Point Research (CPR). https://www.checkpoint.com/
3. Howden's 2025 cyber insurance report https://www.howdengroupholdings.com/sites/default/files/2025-09/howden-2025-cyber-report-rebooting-growth.pdf
4. Le, V.-H., Zhang, H.: Log-based anomaly detection with deep learning: how far are we? In: Proceedings of the 44th International Conference on Software Engineering (ICSE '22), pp. 1356–1367. Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3510003.3510155
5. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: a survey. ACM Comput. Surv. **41**(3), 1–58 (2009)
6. Du, M., Li, F., Zheng, G., Srikumar, V.: Deep log: anomaly detection and diagnosis from system logs through deep learning. In: Proceedings ACM SIGSAC Conference Computer and Communication Security, pp. 1285–1298 (2017)
7. W. Meng, et al.: LogAnomaly: Unsupervised detection of sequential and quantitative anomalies in unstructured logs
8. Lin, Y., et al.: Semi-supervised log-based anomaly detection via probabilistic label estimation. In: ICSE 2021, pp. 1448–1460. IEEE (2021)
9. Zhang, X., et al.: Robust log-based anomaly detection on unstable log data. In: ESEC/FSE 2019, pp. 807–817 (2019)
10. Lu, S., Wei, X., Li, Y., Wang, L.: Detecting anomaly in big data system logs using convolutional neural network. In: 2018 IEEE 16th Intl Conference on Dependable, Autonomic and Secure Computing, 16th Intl Conference on Pervasive Intelligence and Computing, 4th Intl Conference on Big Data Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech), pp. 151–158. IEEE (2018)

11. Ede, T.V., et al.: DEEPCASE: semi-supervised contextual analysis of security events. In: 2022 IEEE Symposium on Security and Privacy (SP), pp. 522–539, San Francisco, CA, USA (2022)
12. Ding, Z., Fei, M.: An anomaly detection approach based on isolation forest algorithm for streaming data using sliding window. IFAC Proc. Vol. **46**(20), 12–17 (2013)
13. Breiman, L.: Random forests. Mach. Learn. **45**, 5–32 (2001)
14. Wang, B., Ying, S., Yang, Z.: A log-based anomaly detection method with efficient neighbor searching and automatic k neighbor selection. Sci. Program. **2020**, 1–17 (2020)
15. Guo, H., Yuan, S., Wu, X.: LogBERT: log anomaly detection via BERT. In: 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8, Shenzhen, China (2021). https://doi.org/10.1109/IJCNN52387.2021.9534113
16. Liu, Z., Qin, T., Guan, X., Jiang, H., Wang, C.: An integrated method for anomaly detection from massive system logs. IEEE Access. **6**, 30602–30611 (2018)
17. Nipa, N.A., et al.: A comparative study of log-based anomaly detection methods in real-world system logs. In: International Conference on Internet of Things, Big Data and Security (2025)
18. DOCKER https://www.docker.com

Analyzing Ensemble Stacking Technologies for Supervised Anomaly Detection

Arcangelo Castiglione[1], Gianni D' Angelo[1(✉)], Francesco Palmieri[1], and Antonio Scarfó[2]

[1] Department of Computer Science, University of Salerno, Fisciano, Salerno, Italy
{arcastiglione,giadangelo,fpalmieri}@unisa.it
[2] Fortinet, Rome, Italy
ascarfo@fortinet.com

Abstract. Network Anomaly Detection (NAD) has emerged as one of the most effective approaches for timely identifying network-based cyber-attacks, offering clear advantages over traditional signature-based intrusion detection systems. The integration of advanced machine learning (ML) techniques can further enhance the adaptability and automation of detection mechanisms, improving their capacity to identify evolving and polymorphic threats. This study evaluates several supervised detection schemes employing ensemble-based ML models, specifically, bagging, boosting, and stacking, to assess their relative strengths and limitations. The analysis aims to develop novel architectures that exploit ensemble learning's capacity to combine diverse models, thereby improving detection performance even under highly imbalanced data conditions. Experimental results indicate that ensemble methods based on decision tree weak learners, such as Random Forest and XGBoost, achieve superior performance compared to both conventional classifiers and stacked ensembles, effectively reducing variance through bagging, mitigating bias via boosting, and enhancing overall prediction accuracy.

Keywords: Ensemble Learning · Network Anomaly Detection · Stacking · XGBoost · Random Forest · Imbalanced Datasets

1 Introduction

In modern cybersecurity, network intrusion and attack detection have become central to the protection and management of organizational infrastructures. Effective detection requires processing large-scale data, making big data analytics and machine learning (ML) essential for improving event classification and detection accuracy. Since no single ML model universally optimizes key performance indicators (KPIs) such as accuracy and robustness, ensemble learning offers a viable strategy by integrating multiple models to enhance performance [23] [28] [14]. However, existing ML-based detection methods still face

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 607–623, 2027.
https://doi.org/10.1007/978-3-032-30518-3_38

critical limitations: they often rely on hand-crafted feature selection, struggle with highly imbalanced traffic distributions, and lack robustness against evolving attack patterns. Moreover, many conventional classifiers suffer from high variance or bias when trained on noisy or incomplete network flow data, leading to degraded generalization in real-world deployments.

In supervised intrusion detection, ensemble methods mitigate overfitting and bias while increasing sensitivity and reducing false alarms. This study evaluates detection performance, measured through standard KPIs, of several homogeneous and heterogeneous weak learners combined via bootstrap aggregation, gradient boosting, or stacking. The analysis prioritizes integrated algorithmic frameworks but also explores hybrid stacked architectures using base-learners (level-0) and meta-learners (level-1). The proposed architectures are benchmarked on widely adopted labeled datasets to establish reliable performance references and model-selection guidelines for supervised network anomaly detection.

The core classifiers used in stacked models include Gaussian Naive Bayes, Decision Trees, Support Vector Machines, and Logistic Regression. Meanwhile, Random Forest and eXtreme Gradient Boosting (XGBoost) represent state-of-the-art ensemble approaches that leverage bagging and boosting, respectively.

Although stacking multiple heterogeneous learners improves performance, the integrated tree-based ensembles, Random Forest and XGBoost, consistently outperform traditional classifiers and stacked combinations. Bagging and boosting reduce variance and bias, respectively, resulting in superior predictive stability. Moreover, gradient boosting further minimizes error rates through regularization, pruning, and cross-validation, thereby combining the advantages of its base learners while mitigating their drawbacks. Both Random Forest and XGBoost exhibit scalability across diverse attack scenarios and varying class balance conditions.

The remainder of this paper is organized as follows: Sect. 2 reviews related work; Sect. 3 introduces the ensemble and baseline learning methods; Sect. 4 details the experimental architectures and datasets; Sect. 5 discusses results; and Sect. 6 presents the final remarks.

2 Related Work

Ensemble Learning is a machine learning methodology that combines several baseline models to create a model that outperforms them. Ensemble learning within a meta-layered architecture [34] has demonstrated significant improvements in base model performance across several domains [18], particularly in medical image and sentiment analysis. However, its adoption in network traffic and Network Anomaly Detection (NAD) remains limited, despite its suitability for supervised binary classification problems. According to [22], ensemble techniques rank among the top ten methods employed, with hybrid combinations frequently used for anomaly detection. Nonetheless, [35] notes that ensemble-based NAD studies represent only a small fraction compared to conventional machine learning approaches.

Within the meta-learning class, stacking is relatively uncommon yet predominant [18]. For instance, [13] proposes a stacking framework combining KNN, AdaBoost, and Random Forest, optimized via Logistic Regression. Similarly, [4] contrasts an SVM meta-classifier with nine base classifiers on the NSL-KDD dataset [8], showing that only some multi-classifier models outperform SVM across key performance metrics. In [9], XGBoost surpasses various supervised and semi-supervised methods for in-vehicle CAN Bus NAD using the Avante CN7 dataset.

Other works, e.g., [1], employ Naïve Bayes, ANN, J48, and REPTree in bagging, boosting, and stacking schemes on UNSW-NB15 [19], showing stacking achieves the best accuracy. According to [21], a collective prediction is achieved through majority voting among ANN, SVM, and MARS models. Likewise, [36] demonstrates that combining one-class classifiers (LGP, ANFIS, RF) on KDD-Cup 1999 [33] enhances detection accuracy by leveraging model complementarity. Using CICDDoS 2019 [16,30] shows that stacking-based ensembles outperform both traditional and other ensemble methods in detection accuracy and false alarm rate.

In web anomaly detection, [31] compares strong ensemble base learners (GBM, RF, XGB) on multiple datasets (CSIC-2010v2, CICIDS-2017, NSL-KDD, UNSW-NB15), revealing superior performance of stacked ensembles. Similarly, [26] confirms that meta-learning improves classification accuracy and reduces false positives using real traffic data. Finally, [11] highlights that ensemble models better detect complex attacks and minimize false alarms through multi-source correlation.

Overall, the cited works confirm that ensemble learning enhances the performance of both weak and strong base learners across various application domains.

3 Background

Algorithms and techniques adopted in our analysis are concisely introduced to provide the essential methodological background.

3.1 Ensemble Learning

Ensemble learning [34] combines multiple base models to enhance prediction accuracy. Two main strategies exist: **sequential** (models trained in series, each correcting previous errors) and **parallel** (models trained independently, aggregated via voting or averaging). Ensembles can be *homogeneous* or *heterogeneous* depending on whether the base learners are of the same type.
The main techniques relevant to this work are bagging, boosting, and stacking, defined as follows:

– **Bagging (Bootstrap Aggregation)**: a parallel approach that trains multiple models on random bootstrap samples of the dataset and aggregates their predictions. It effectively reduces variance and is well-suited for unstable learners such as Random Forests or Extra Trees.

– **Boosting**: a sequential technique that iteratively combines weak learners to minimize bias and residual errors. Misclassified instances receive higher weights in subsequent iterations. Variants include AdaBoost, Gradient Boosting Machine (GBM), XGBoost, LightGBM, and CatBoost. In gradient boosting [5], the model is updated iteratively to reduce prediction error:

$$F_m = F_{m-1} - \mu_m \cdot \nabla \varepsilon_m \tag{1}$$

where μ_m is the learning rate and ε_m denotes the prediction error at step m.
– **Stacking (Stacked Generalization)**: an ensemble method that integrates the predictions of several base-level (level-0) models using a meta-model (level-1). The meta-model, often a simple regressor or classifier (e.g., Linear or Logistic Regression), learns from the outputs of heterogeneous base learners to improve generalization and reduce bias or variance.

3.2 Logistic Regression

Logistic Regression [15] is a statistically interpretable model for binary classification. It employs the logistic (sigmoid) function to estimate the probability of the positive class, with parameters optimized via Maximum Likelihood Estimation to maximize the likelihood of the observed outcomes.

3.3 Support Vector Machines

Support Vector Machines (SVMs) [6] are supervised learning models designed to find the optimal separating hyperplane that maximizes the margin between classes in a high-dimensional feature space. Kernel functions (e.g., linear, polynomial, radial-basis, sigmoid) enable nonlinear mapping. The optimization process minimizes the norm of the weight vector subject to margin constraints, relying on *support vectors*, i.e., samples near the decision boundary that define the classifier.

3.4 Decision Tree and Random Forest

A Decision Tree [12] partitions the feature space through hierarchical decisions, where internal nodes represent feature tests and leaves denote class labels. Features appearing at upper tree levels typically have greater discriminative importance. Pruning reduces overfitting by eliminating weak branches.
Random Forest [2] extends this principle through a bagging-based ensemble of multiple decision trees, each trained on a random subset of data/features. Aggregating these predictors via majority voting (classification) or averaging (regression) lowers variance, enhances generalization, and mitigates overfitting. This was selected as the most effective bagging-based solution in our analysis.

3.5 Gaussian Naive Bayes

The Gaussian Naive Bayes (GNB) classifier assumes conditional independence among features and Gaussian-distributed variables. Based on Bayes' theorem, the predicted class is determined by:

$$pred = \arg\max_{class} p(class|features) = \frac{p(features|class) \cdot p(class)}{p(features)} \quad (2)$$

Despite its simplicity, GNB is computationally efficient and performs effectively on high-dimensional data.

3.6 eXtreme Gradient Boosting

eXtreme Gradient Boosting (XGBoost) [5] extends gradient boosting by incorporating second-order optimization (Newton-Raphson method) and advanced regularization to enhance scalability and convergence. It constructs an ensemble of regularized decision trees that sequentially correct previous residuals, thereby improving stability and preventing overfitting. XGBoost's support for parallel and distributed computing makes it suitable for large-scale datasets, motivating its adoption as the reference gradient-boosting solution in this study.

4 The Analysis Framework

Our analysis framework aims to assess the capabilities of different ensemble-based ML architectures to inspect network traffic flows in multiple flow-based anomaly detection scenarios and identify attack-related flows that occur among normal ones. It essentially consists of a supervised binary classification framework in which training is typically performed on a highly unbalanced baseline, as in a typical traffic observation/capture scenario, where the number of normal traffic flows (negatives in the supervisory signal) significantly exceeds that of anomalous ones (positives).

4.1 Architectural Layout

The implemented architectures are designed to address distinct analysis tasks through tailored compositions of Stacking, Bootstrap Aggregation, and Boosting techniques. These configurations aim to evaluate both individual learners and their combinations across four detection scenarios, each tested on different datasets. The overall structure is illustrated in the four schemes of Fig. 1, each corresponding to a specific evaluation task.

In particular, Fig. 1a presents the comparison among the three primary ensemble approaches-bagging, boosting, and stacking. For this purpose, Random Forest and XGBoost were selected as representative models for bootstrap aggregation and optimized boosting, respectively. Their integration within a unified algorithmic framework ensures a balance between optimization efficiency

and architectural simplicity, yielding stable and effective solutions. The stacking ensemble adopts a two-level configuration: Logistic Regression, Gaussian Naive Bayes, Decision Tree, and SVM serve as level-0 learners, whose predictions are combined by a Logistic Regression meta-learner at level 1. This evaluation not only benchmarks stacking against bagging and boosting but also investigates the benefits of aggregating heterogeneous learners with varying complexity and predictive accuracy. Given the difficulty of identifying the optimal single model for a dataset, stacking offers a straightforward alternative that integrates multiple learners, often achieving superior overall performance.

The second configuration (Fig. 1b) compares standalone bagging (Random Forest) and boosting (XGBoost) models with their stacked combination, where a Logistic Regression meta-learner aggregates their outputs to assess whether this hybridization enhances predictive performance.

The third scheme (Fig. 1c) focuses exclusively on stacking, comparing (i) pure stacking, where level-0 learners are trained independently and the meta-classifier learns from their predictions, and (ii) passthrough stacking, where the meta-model is trained on both base model predictions and original features, providing contextual information for improved combination. Both stacking variants are evaluated for the configurations in Fig. 1a and Fig. 1b.

Finally, the fourth scheme (Fig. 1d) explores the use of a boosted meta-classifier (XGBoost) within a passthrough stack comprising heterogeneous learners (e.g., Logistic Regression, Naive Bayes) and ensemble methods (e.g., Random Forest), to assess its effectiveness relative to pure bagging or boosting architectures.

4.2 Datasets

In all conducted experiments, the objective was to ensure comparability with existing state-of-the-art studies; therefore, we employed the most widely adopted benchmark datasets in the field [27]. Although numerous datasets provide labeled flows and raw packet data, only a few are recognized as standard references in the field of anomaly and intrusion detection research. This is primarily due to limitations in traffic generation methods, testing environments, and data realism, which must align with the evolving complexity of modern networks and emerging cybersecurity threats.

The first dataset used was NSL-KDD [32] [8] [33], developed by the Canadian Institute for Cybersecurity as an enhanced and de-duplicated version of the DARPA Lincoln Labs KDD'99 dataset [7]. It contains labeled flows extracted from several weeks of simulated U.S. Air Force LAN traffic using the Bro IDS [24], including 24 attack types and 41 features per flow, divided into training and testing subsets. Despite criticism regarding its artificial nature [3], KDD99 remains among the most frequently used datasets in intrusion detection research [22].

The second dataset, UNSW-NB15 [19], was created at UNSW Canberra to address major deficiencies in KDD99 and NSL-KDD [20], such as outdated traffic patterns and unbalanced data distributions. Generated in a controlled environment in 2015, it comprises over two million labeled flows, each with 49 features,

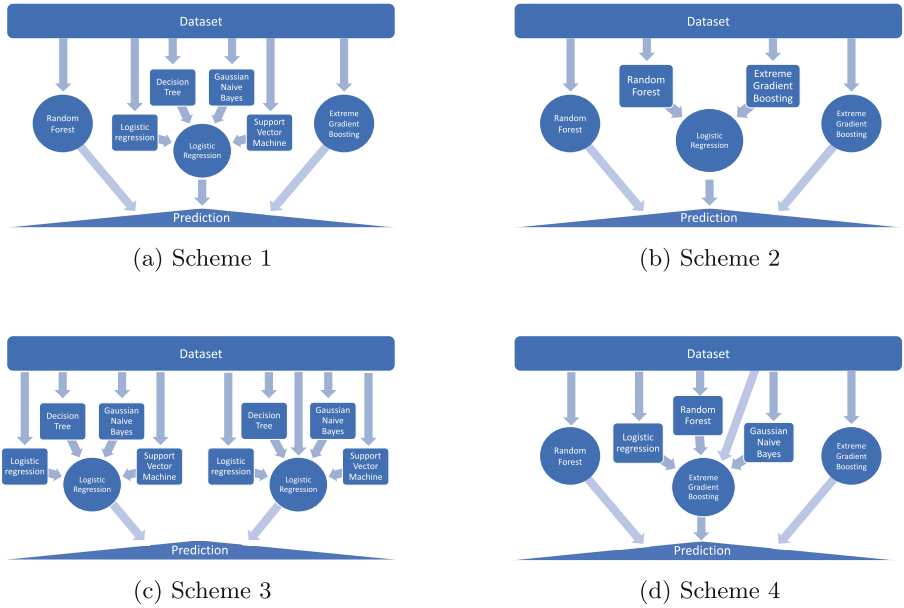

Fig. 1. Layout of the architectures used in the analysis.

extracted using the Bro IDS. Approximately 87.3% of the records correspond to normal traffic and 12.7% to anomalous activity.

Lastly, the CIC-IDS-2017 dataset [29], also from the Canadian Institute for Cybersecurity, encompasses five days of realistic network traffic across a full-scale topology including firewalls, switches, routers, and heterogeneous hosts (Windows, Ubuntu, macOS). Using CICFlowMeter [10] [17], more than 80 flow features were extracted from both intra-LAN and Internet communications, covering multiple attack scenarios mixed with regular network behavior. The dataset comprises nearly three million flows, with 80.3% labeled as normal and 19.7% labeled as anomalous.

In our experiments, for NSL-KDD, we used the already provided split. For both UNSW-NB15 and CIC-IDS-2017, 70% of the entire dataset was used for training, with the remaining part serving as the testing reference.

4.3 Features

All the features that are obviously unnecessary (source and destination IP addresses of flows, timestamps, and attack labels) for classification were not considered by preliminarily removing them from the feature set. In addition, to optimize the effectiveness and responsiveness of the overall ML architecture, it is necessary to reduce the problem dimensions by minimizing the number of features used for both training and testing, according to a strong diversification

principle that prioritizes the choice of highly discriminative and mutually independent features. This implies the determination of a more compact knowledge representation, also improving predictiveness through the reduction of redundancy in feature space, since a set of fully independent features, when considered together, is more effective in performing binary classification. To provide an example at a glance of feature correlation, we report in Fig. 2 the correlation heatmaps related to the CIC-IDS-2017 and UNSW-NB15 datasets, respectively.

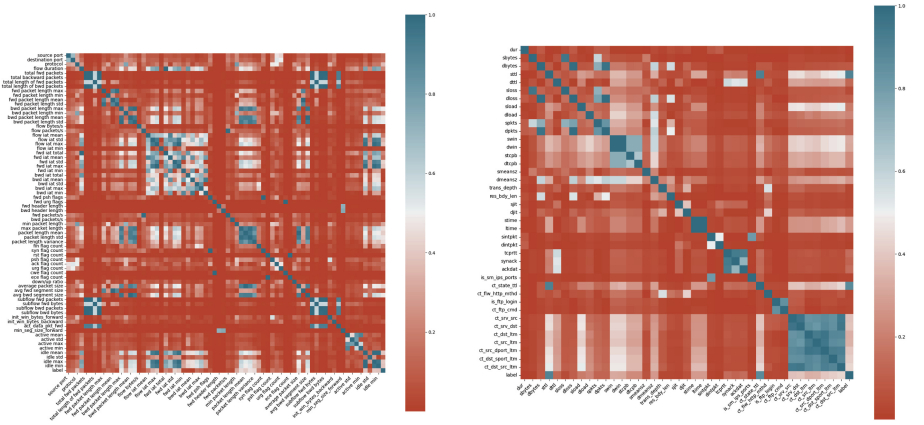

(a) CIC-IDS-2017 (b) UNSW-NB15

Fig. 2. Correlation heatmaps.

According to the previous considerations, some feature engineering has been performed on the original datasets. More precisely, in the presence of two or more highly correlated features (characterized by a mutual correlation greater than 0.95), we only kept one of them and removed all the others. The features removed by the datasets of interest are summarized in Table 1.

Finally, all numeric feature values have been normalized to the range $[0, 1]$ through standard scaling, and categorical features have been remapped into ordered numeric values, preserving the original order information.

5 Performance Evaluation and Results Analysis

In this section, we describe the results of the experiments performed in our analysis to evaluate the effectiveness of the available ensemble-based detection approaches by also comparing them to determine the best options in terms of configuration as well as the more interesting properties characterizing their behavior, together with their weaknesses, depending on the specific usage scenario.

Table 1. Feature Engineering Details

Dataset	Removed Features
NSL-KDD	'su attempted' (15), 'serror rate' (25), 'rerror rate' (27), 'dst host srv diff host rate' (37), 'dst host serror rate' (38), 'dst host rerror rate' (40), 'dst host srv rerror rate' (41)
CIC-IDS-2017	'flow id', 'source ip', 'source port', 'destination ip', 'destination port', 'timestamp', 'attack cat', 'total backward packets', 'total length of bwd packets', 'fwd packet length std', 'bwd packet length mean', 'bwd packet length std', 'fwd iat total', 'fwd iat max', 'fwd packets/s', 'packet length std', 'syn flag count', 'cwe flag count', 'ece flag count', 'average packet size', 'avg fwd segment size', 'avg bwd segment size', 'subflow fwd packets', 'subflow fwd bytes', 'subflow bwd packets', 'subflow bwd bytes', 'idle mean', 'idle max', 'idle min'
UNSW-NB15	'srcip', 'sport', 'dstip', 'dsport', 'attack cat', 'sloss', 'dloss', 'dpkts', 'dwin', 'ltime', 'ct srv dst', 'ct src dport ltm', 'ct dst src ltm'

A simple testbed has been implemented in `Python` by using the `scikit-learn` machine learning framework on an Intel Xeon-based server equipped with double 2.2Ghz CPUs and 160 GB RAM. The most significant hyperparameters of all the learners used within the assessed ensemble have been tuned by using grid search (i.e., exhaustive exploration of all the combinations in specific ranges or sets of values) 5-fold cross-validation for selecting the best values to be used in all the runs.

To ensure reproducibility, we report the optimal hyperparameters obtained via grid search with 5-fold cross-validation:

- **Random Forest**: `n_estimators=200`, `max_depth=20`, `min_samples_split=5`.
- **XGBoost**: `learning_rate=0.1`, `n_estimators=150`, `max_depth=6`, `gamma=0`, `reg_alpha=0.1`, `reg_lambda=1`.
- **Logistic Regression**: `C=1.0`, `solver='lbfgs'`.
- **SVM**: `C=1.0`, `kernel='rbf'`, `gamma='scale'`.
- **Gaussian Naive Bayes** and **Decision Tree** (`max_depth=10`): default scikit-learn parameters were used unless otherwise tuned.

All other ensemble-specific parameters were kept at their default values as provided by the adopted machine learning frameworks [25].

Common performance metrics used to evaluate machine learning-based classification models have been considered in all experiments performed within the

analysis framework. In particular, Precision, Recall, and F1-score, determined from the confusion matrix, have been examined for both the negative and positive classes.

5.1 Scheme 1 Experiments

In the analysis scenario depicted in Fig. 1a, the results presented in Figs. 3, 4, and 5 demonstrate that ensemble-based approaches consistently outperform conventional classifiers (SVM, Gaussian Naive Bayes, Logistic Regression, Decision Tree). Among them, Random Forest and XGBoost achieved the highest accuracy across all evaluated datasets, with XGBoost showing a slight advantage on larger and more imbalanced datasets (UNSW-NB15 and CIC-IDS-2017). The superior performance of these ensemble methods arises from their inherent algorithmic characteristics. Random Forest mitigates variance and overfitting by averaging multiple decorrelated decision trees generated via bootstrap aggregation, while feature randomization enhances robustness to missing data and improves generalization. Conversely, XGBoost constructs successive shallow trees in a boosting framework, iteratively correcting previous misclassifications. Its pruning strategy and the exploitation of feature distribution statistics optimize split selection and reduce computational complexity. Regularization terms-γ (minimum loss reduction for split creation), α (L1), and λ (L2)-further balance bias and variance by penalizing model complexity, which were tuned in this study to enhance predictive performance.

The stacking ensemble, although marginally inferior to Random Forest and XGBoost, produced results comparable to those of the strongest base learner (typically the Decision Tree) while surpassing the performance of individual models. This improvement stems from the complementary learning of level-0 models trained on distinct data partitions, which effectively combines diverse prediction patterns and reduces correlated errors.

Finally, all models exhibited higher accuracy on the negative class, except for NSL-KDD. This behavior is attributed to the pronounced class imbalance in the other two datasets, where negative samples comprise approximately 80% of the training data.

5.2 Scheme 2 Experiments

Regarding the analysis scenario reported in Fig. 1b, from the results provided in Figs. 6, 7, and 8, we can immediately appreciate that no performance gain is introduced by stacking the XGBoost and Random Forest models on all three reference datasets. The reason is that both algorithms are already highly optimized and directly achieve a performance that is extremely difficult to improve. Hence, the benefits introduced by stacking are either imperceptible or negligible. Indeed, the main aim of stacking is to complement the strengths of one learner with the weaknesses of another and vice versa, in a logic of mutual support between the involved individual learners. However, in our case, both the

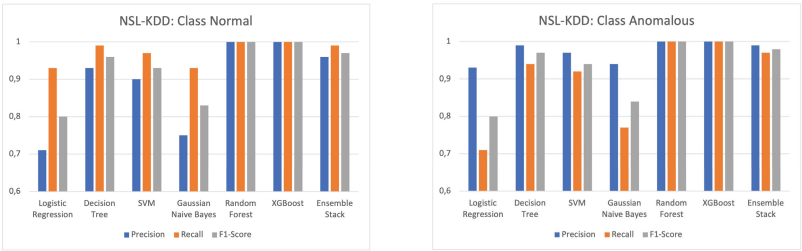

(a) Negative Class. (b) Positive Class.

Fig. 3. NSL-KDD dataset Per-class precision, recall, and F1-score for Scheme 1 (bagging: Random Forest; boosting: XGBoost; stacking vs. conventional classifiers: SVM, GNB, LR, DT).

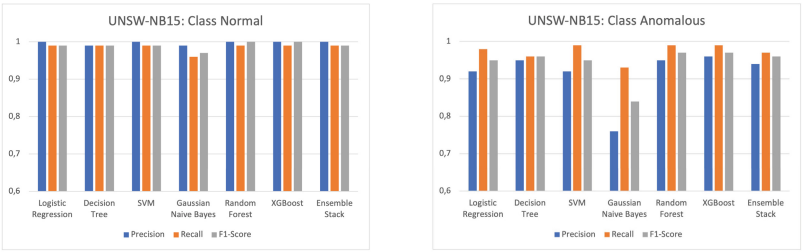

(a) Negative Class. (b) Positive Class.

Fig. 4. UNSW-NB15 dataset Per-class performance comparison for Scheme 1. Ensemble methods (RF, XGB) outperform individual baselines across both classes.

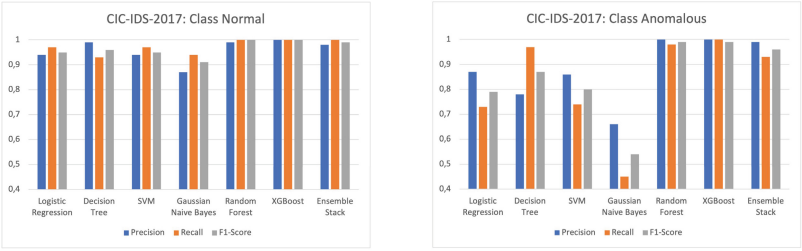

(a) Negative Class. (b) Positive Class.

Fig. 5. CIC-IDS-2017 dataset Scheme 1 results. XGBoost shows a slight advantage over RF on the imbalanced positive class.

XGBoost and Random Forest algorithms achieve nearly perfect accuracy on all the datasets, leaving no room for mutual support.

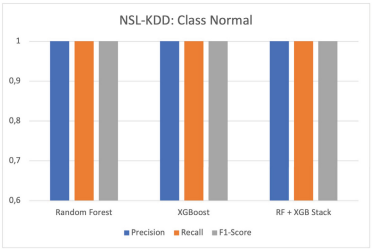

(a) Negative Class.

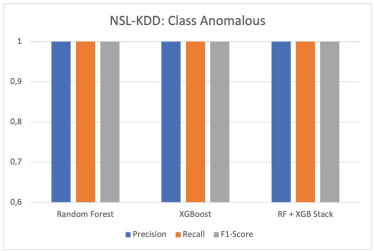

(b) Positive Class.

Fig. 6. NSL-KDD: Per-Class Scheme 2 performance.

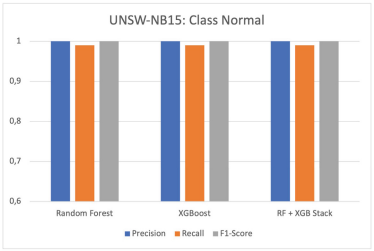

(a) Negative Class.

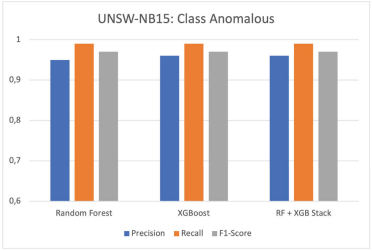

(b) Positive Class.

Fig. 7. UNSW-NB15: Per-Class Scheme 2 performance.

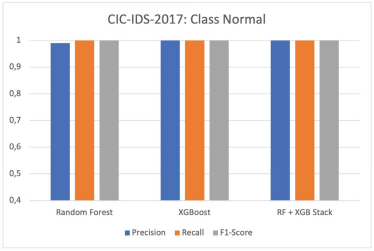

(a) Negative Class.

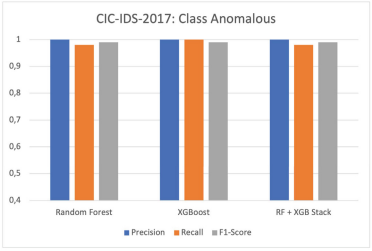

(b) Positive Class.

Fig. 8. CIC-IDS-2017: Per-Class Scheme 2.

5.3 Scheme 3 Experiments

Regarding the analysis scenario reported in Fig. 1c, from the results provided in Fig. 9, Fig. 10 and Fig. 11 we can see that training the metamodel in the stacking architectures on the predictions made by the base models merged with the training original data, leads in several cases to worse results, mainly when the weaker learners are used in the level-0 (this effect is less evident, but also present, when the XGBoost and Random Forest models are stacked. This may

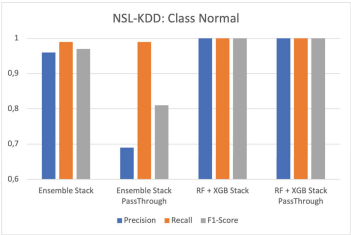

(a) Negative Class.

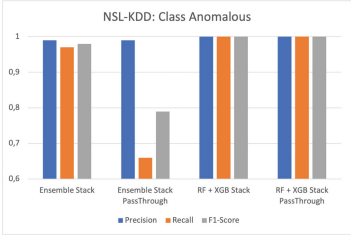

(b) Positive Class.

Fig. 9. NSL-KDD: Per-Class Scheme 3 performance.

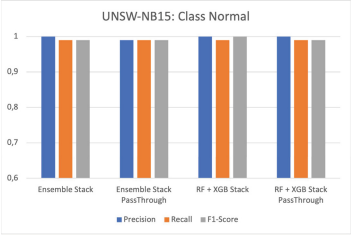

(a) Negative Class.

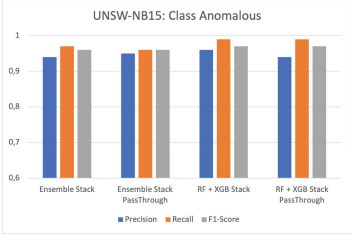

(b) Positive Class.

Fig. 10. UNSW-NB15: Per-Class Scheme 3 performance.

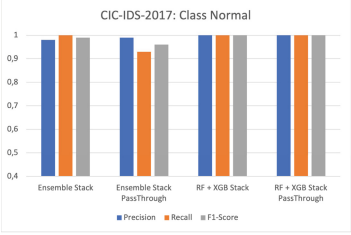

(a) Negative Class.

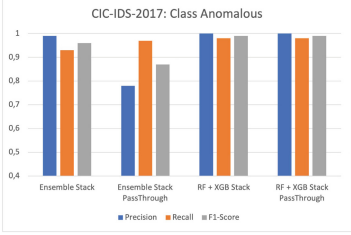

(b) Positive Class.

Fig. 11. CIC-IDS-2017: Per-Class Scheme 3 performance.

be due to the heterogeneity of data in the original datasets compared to those from the individual learners' predictions.

5.4 Scheme 4 Experiments

Finally, regarding the analysis scenario reported in Fig. 1d, the same considerations made in the two previous cases can be extended to the results provided in Fig. 12, where, for simplicity, the data relative to the passthrough stacking ensemble have been aggregated on a single chart for all three datasets. Here, it is immediate to see how the use of XGBoost as a meta-classifier in the stacked

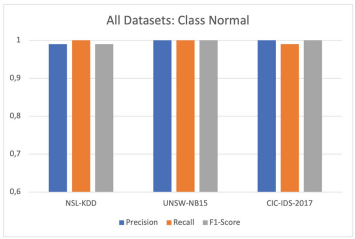

(a) Negative Class.

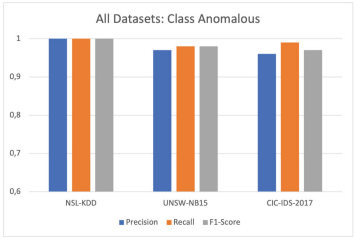

(b) Positive Class.

Fig. 12. All Datasets: Per-Class Scheme 4 performance.

configuration, taking as its input the predictions of the level-0 classifiers (Logistic Regression, Random Forest, Gaussian Naive Bayes) merged with the original dataset, does not introduce any advantage and, in some cases, its effectiveness is also slightly perturbated by the effect of the weaker level-0 learners (Gaussian Naive Bayes).

6 Conclusions and Future Work

This study conducts a comparative evaluation of various ensemble-based machine learning (ML) architectures for analyzing network traffic flows across multiple flow-based anomaly detection scenarios, to identify attack-related flows concealed among legitimate ones. The analysis primarily focused on selecting optimal ensemble configurations, favoring approaches integrated within single algorithmic frameworks, while also exploring their use as base (level-0) or meta (level-1) learners in stacked architectures. The resulting reference detection models were assessed using widely recognized labeled datasets in the field, i.e., NSL-KDD, UNSW-NB15, and CIC-IDS-2017, to provide practical guidance for selecting suitable supervised learning models for network anomaly detection.

The performance evaluation confirmed the superior effectiveness of ensemble methods, even under conditions of significant class imbalance. Most ensemble-based classifiers enhanced the predictive performance of their constituent models, particularly when the latter exhibited limited standalone accuracy. Despite the capacity of stacked ensembles to integrate complementary learning strategies, the results demonstrated that well-designed single-algorithm ensembles, specifically, Random Forest (bagging-based) and XGBoost (boosting-based), outperformed both traditional classifiers and their stacked counterparts. These models achieved optimal detection performance by reducing variance through bagging, mitigating bias via boosting, and thereby improving overall prediction quality. However, it is worth noting that stacking may still prove beneficial in contexts where the level-0 models are weak learners (e.g., simple classifiers with high bias or variance), rather than already-optimized ensembles like XGBoost or Random Forest.

Our evaluation has several limitations. First, experiments were conducted on simulated, though widely used, datasets (NSL-KDD, UNSW-NB15, CIC-IDS-

2017); real-world network traffic may exhibit different distributions and concept drift. Second, we only considered supervised binary classification; multi-class attack identification and semi-supervised scenarios were not tested. Third, computational efficiency and inference latency, that is critical for real-time deployment, were not benchmarked. Finally, hyperparameter values were optimized per dataset and may not generalize to other environments.

Addressing these aspects is part of our future work. Also, future research could explore semi- or self-supervised ensemble frameworks to reduce dependence on large labeled datasets. Incorporating automated feature engineering and deep representation learning may enable the extraction of more discriminative patterns from raw network flows. Efforts should also focus on improving computational efficiency via model distillation or online learning for real-time applications. Strengthening robustness against adversarial evasion and concept drift, alongside integrating explainable AI for interpretability, remains essential. Extending these approaches to cross-domain and multimodal detection tasks would further enhance generalizability and practical utility.

Acknowledgments. This work was partially supported by project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU.

References

1. Belouch, M., El Hadaj, S.: Comparison of ensemble learning methods applied to network intrusion detection. In: Proceedings of the 2nd International Conference on Internet of Things, Data and Cloud Computing (ICC 2017), pp. 1–4. ACM, New York, NY, USA (2017)
2. Breiman, L.: Random forests. Mach. Learn. **45**, 5–32 (2001)
3. Brugger, S.T., Chow, J.: An assessment of the DARPA ids evaluation dataset using snort, UC Davis, Department of Computer Science (2007). Tech. rep
4. Chand, N., Mishra, P., Krishna, C.R., Pilli, E.S., Govil, M.C.: A comparative analysis of SVM and its stacking with other classification algorithm for intrusion detection. In: 2016 International Conference on Advances in Computing, Communication, & Automation (ICACCA), pp. 1–6. Dehradun, India (2016)
5. Chen, T., Guestrin, C.: XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016), pp. 785–794. ACM, New York, NY, USA (2016)
6. Cortes, C., Vapnik, V.: Support-vector networks. Mach. Learn. **20**(3), 273–297 (1995)
7. DARPA: KDD cup 1999 data set (1999). http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html
8. Mohi-ud din, G.: NSL-KDD. IEEE Dataport (2018)
9. Dong, Y., Chen, K., Peng, Y., Ma, Z.: Comparative study on supervised versus semi-supervised machine learning for anomaly detection of in-vehicle can network. In: IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), pp. 2914–2919. Macau, China (2022)

10. Drapper Gil, G., Lashkari, A.H., Mamun, M., Ghorbani, A.A.: Characterization of encrypted and VPN traffic using time-related features. In: 2nd International Conference on Information Systems Security and Privacy (ICISSP 2016), pp. 407–414 (2016)
11. Folino, G., Sabatino, P.: Ensemble based collaborative and distributed intrusion detection systems: a survey. J. Netw. Comput. Appl. **66**, 1–16 (2016)
12. Furnkranz, J.: Decision tree. In: Sammut, C., Webb, G.I. (eds.) Encyclopedia of Machine Learning. Springer, Boston, MA (2011). https://doi.org/10.1007/978-0-387-30164-8
13. Hai, T., Hoang, L., Huh, E.: Network anomaly detection based on late fusion of several machine learning algorithms. Int. J. Comput. Netw. Commun. **12**, 117–131 (2020)
14. Hansen, L.K., Salamon, P.: Neural network ensembles. IEEE Trans. Pattern Anal. Mach. Intell. **12**, 993–1001 (1990)
15. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning. Springer, New York (2001). https://doi.org/10.1007/978-0-387-84858-7
16. Khoei, T.T., Aissou, G., Hu, W.C., Kaabouch, N.: Ensemble learning methods for anomaly intrusion detection system in smart grid. In: 2021 IEEE International Conference on Electro Information Technology (EIT), pp. 129–135. Mt. Pleasant, MI, USA (2021)
17. Lashkari, A.H., Draper-Gil, G., Mamun, M.S.I., Ghorbani, A.A.: Characterization of tor traffic using time based features. In: 3rd International Conference on Information System Security and Privacy, Porto, Portugal (2017)
18. Mohammed, A., Kora, R.: A comprehensive review on ensemble deep learning: opportunities and challenges. J. King Saud Univ. Comput. Inf. Sci. **35**(2), 757–774 (2023)
19. Moustafa, N., Slay, J.: UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). In: 2015 Military Communications and Information Systems Conference (MilCIS), pp. 1–6 (2015)
20. Moustafa, N., Slay, J.: The evaluation of network anomaly detection systems: statistical analysis of the UNSW-NB15 data set and the comparison with the KDD99 data set. Inf. Secur. J. A Global Perspect. **25**(1–3), 18–31 (2016)
21. Mukkamala, S., Sung, A.H., Abraham, A.: Intrusion detection using an ensemble of intelligent paradigms. J. Netw. Comput. Appl. **28**(2), 167–182 (2005)
22. Nassif, A.B., Talib, M.A., Nasir, Q., Dakalbab, F.M.: Machine learning for anomaly detection: a systematic review. IEEE Access **9**, 78658–78700 (2021)
23. Oza, N.C., Tumer, K.: Classifier ensembles: select real-world applications. Inf. Fus. **9**(1), 4–20 (2008)
24. Paxson, V.: Bro: a system for detecting network intruders in real-time. Comput. Netw. **31**(23), 2435–2463 (1999)
25. Pedregosa, F., Varoquaux, G., Gramfort, A., et al.: Scikit-learn: machine learning in Python. J. Mach. Learn. Res. **12**, 2825–2830 (2011)
26. Possebon, I.P., Silva, A.S., Granville, L.Z., Schaeffer-Filho, A., Marnerides, A.: Improved network traffic classification using ensemble learning. In: IEEE Symposium on Computers and Communications (ISCC), pp. 1–6. , Barcelona, Spain (2019)
27. Ring, M., Wunderlich, S., Scheuring, D., Landes, D., Hotho, A.: A survey of network-based intrusion detection data sets. Comput. Secur. **86**, 147–167 (2019)
28. Schapire, R.E.: The strength of weak learnability. Mach. Learn. **5**(2), 197–227 (1990)

29. Sharafaldin, I., Lashkari, A.H., Ghorbani, A.A.: Toward generating a new intrusion detection dataset and intrusion traffic characterization. In: 4th International Conference on Information Systems Security and Privacy (ICISSP) (2018)
30. Talukder, M.A., Uddin, M.A.: CIC-ddos2019 dataset. Mendeley Data, V1 (2023)
31. Tama, B.A., Nkenyereye, L., Islam, S.M.R., Kwak, K.S.: An enhanced anomaly detection in web traffic using a stack of classifier ensemble. IEEE Access **8**, 24120–24134 (2020)
32. Tavallaee, M.: NSL-KDD dataset (2012). http://www.iscx.ca/NSL-KDD
33. Tavallaee, M., Bagheri, E., Lu, W., Ghorbani, A.A.: A detailed analysis of the KDD cup 99 data set. In: 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications, pp. 1–6. ON, Canada, Ottawa (2009)
34. Valiant, L.G.: A theory of the learnable. Commun. ACM **27**, 1134–1142 (1984)
35. Vanerio, J., Casas, P.: Ensemble-learning approaches for network security and anomaly detection. In: Proceedings of the Workshop on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA 2017), pp. 1–6. ACM, New York, NY, USA (2017)
36. Zainal, A., Maarof, M.A., Shamsuddin, S.M., Abraham, A.: Ensemble of one-class classifiers for network intrusion detection system. In: 4th International Conference on Information Assurance and Security, pp. 180–185. Naples, Italy (2008)



Toward Secure and Intelligent Subject Authentication Using Brainprints and Enhanced AMEEGNet

Adnan Elahi Khan Khalil, Jesus Arturo Perez-Daz(✉), Jos Antonio Cantoral-Ceballos, and Javier M. Antelis

School of Engineering and Sciences, Tecnológico de Monterrey, Monterrey, NL 64700, Mexico
jesus.arturo.perez@tec.mx

Abstract. Neural based subject authentication has emerged as a promising paradigm in biometric security, with the P300 event related potential (ERP) gaining particular attention due to its reliability and individual specificity. The proposed scheme introduces a personalized brainprint recognition framework that utilizes an enhanced AMEEGNet architecture to accurately classify symbol evoked P300 responses. A visual oddball paradigm featuring directional arrows and a stop sign was employed to elicit P300 activity with minimal cognitive demand. 16 subjects (participants) completed the task across three separate sessions conducted on different days. EEG data from the session 1 and 2 were used to train the deep learning model, while the third session was utilized for cross-session validation and testing. The proposed framework achieved strong performance, with an average classification accuracy of 89.25% and average Cohen's kappa score of 0.88. These findings demonstrate that symbol-evoked P300 responses carry consistent, individual-specific features, underscoring their potential for use in practical EEG-based subject authentication systems.

Keywords: Brainprint biometrics · Cognitive biometric authentication · Neural signature recognition · Privacy-preserving EEG user authentication

1 Introduction

In current period of knowledge and advanced technology, the importance of safe user authentication schemes are growing with fast pace. The lack of typical conventional approaches such as, passwords and physical tokens have illustrated considerable issues due to their vulnerability to be forgotten, spoofed, or bypassed, that leads to a considerable security breaches and financial losses [1]. Even though facial recognition and fingerprint based user authentication systems have illustrated a very high accuracy, yet they are not immune to security breaches, including spoofing attacks, data breaches, and identity fraud [2]. Subsequently,

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 624–640, 2027.
https://doi.org/10.1007/978-3-032-30518-3_39

scholars have pursued way more secure and tamper-resistant user authentication schemes and methods. In this context, electroencephalogram (EEG) signals have been getting attention as a viable source due to its unique, user specific patterns that are difficult to replicate [3]. Experimental research have shown that EEG-based biometrics exhibit important key properties, including universality, ease of acquisition, and the ability to be revoked or reissued [4]. As a result, EEG based user authentication has emerged a useful approach for future generation biometric system [5]. For example, an EEG based experiment was conducted [6] that involves 30 subjects which involves resting state EEG signals with closed and open eyes for user identification. By employing a similarity based classification approach, the system achieved an average identification accuracy of 85%. However, since the system relies on resting state activity, there was a lack of active mental activity, increasing the surrounding noise, single session experiment leading to potential spoofing attacks. In a given study [7] an EEG authentication framework was developed utilizing rapid serial visual presentation (RSVP) to capture neural responses during login attempts. The given framework employed shrinkage discriminant analysis to classify EEG signals from 16 subjects, which yields a mean validation accuracy of 87%. Although, RSVP based approaches often depend on repetitive passive visual stimulation, that leads to a lower inter session robustness with limited user interaction, which affects a long term stability and user experience for real time applications. In a another investigation, researchers developed an user authentication system using steady state auditory evoked potentials collected from 40 subjects. A canonical correlation analysis was implemented to match the EEG template, they obtained competitive recognition performance [8]. Despite these good results, SSVEP-based methods typically require a long periods for stimulation with deep electrode arrays implementation, which may restrict it for real time practice with fast, and user friendly user authentication scenarios.

Some of the current studies have investigated symbol evoked P300 responses as a robust biometric markers for user authentication systems [9]. Unlike traditional visual evoked potentials, these event related potentials are elicited by different meaningful symbol stimuli that produces a distinctive waveforms which exhibits a strong inter individual variability [10]. The P300 component, particularly when evoked by different symbolic patterns demonstrates an enhanced signal-to-noise characteristics as compared to the conventional flash producing stimuli [11]. Most recent studies indicate that symbol evoked P300 paradigms achieve high classification accuracy in EEG based biometric applications, reporting user authentication in controlled experiments [12]. These studies have established symbol evoked P300 responses as a robust EEG based biometric verifier. Rathi [13] proposed a P300 based brain signal user authentication scheme using picture based visual stimuli instead of conventional alphanumeric characters. The recorded P300 responses from users were tested with three different classifiers by presenting object images that achieved a an accuracy of 97%, highlighting the method's robustness and usability for secure user authentication. However, this work primarily focused using compact EEG setups, no evaluation of temporal

generalization. A recent study by Akram [14] presented a smart home control with symbol based P300 BCI paradigm that achieved an improved flexible command and control classification accuracy of 92.25%. The system helped users to manage various appliances and even to make phone calls, demonstrating the real time use of symbol evoked P300 responses in assistive technologies demonstrating robustness of the EEG based authentication. But, the study was based on control applications rather than universal biometric user authentication, and also did not address the longitudinal stability or security specific evaluations of the biometric system.

Neural cognitive tasks has improved EEG based authentication by utilizing unique neural patterns like P300 ERPs [15]. This method is widely employed for capturing individual active neural responses. For instance, [16] utilized SSVEPs for user identification from a character recognition task. A causal connectivity analysis to EEG data from 15 subjects was applied achieving a high user identification accuracy of 95.5%. Despite of having high performance, this approach is based on large number electrode configurations and extended stimulation, which limited to employ this system in a compact and efficient user authentication environments. Similarly, [17] proposed a user authentication method using rapid serial visual presentation (RSVP) with similar and non similar faces (face RSVP). To support the effectiveness of cognitive task for enhanced EEG biometrics, these stimuli achieved an average verification accuracy of 88.88%. However, but this paradigm is based on controlled laboratory settings and large number electrodes, and cross session robustness with lack of model interpretability exploration which is crucial for scalable user authentication frameworks.

EEG data processing for user authentication employs several key methods and techniques. Like, temporal averaging of EEG epochs enhances the signal to noise ratio of event related potentials (ERPs), improving user biometric performance [18]. Many studies have demonstrated effectiveness of this approach [19]. Deep learning methods have emerged as a prominent classification method in EEG based biometrics, with novel convolutional neural networks (CNNs) being particularly valuable in these approches. In [20] the authors employed CNNs for EEG template matching, while Roy et al. [21] applied CNN based feature extraction to motor imagery data. The specialized CNN architecture knowns as EEGNet [22] has been widely adopted in various user authentication systems demonstrating its effectiveness for EEG signal processing.

EEG based user authentication systems faces four major limitations: (1) reliance on large numbers of electrodes placement, (2) large duration for user authentication process, (3) lacking of standardized naturalistic neural cognitive paradigms, and (4) unverified longitudinal stability. To address these limitations, a novel framework that combines a symbol evoked P300 responses with an enhanced AMEEGNet-CAM architecture is proposed [23]. The system employs (1) a low burden on neural cognitive activity using familiar symbols to elicit distinct P300 ERPs responses, and (2) AMEEGNet-CAM, a convolutional network with channel attention mechanisms using key input electrodes, less EEG data processing and improved interpretability. While active participation is required,

this approach enhances resistance to spoofing attacks and increases user biometric specificity, making it ideal for future critical security applications. The key contribution of the proposed study is as follows:

1. The introduction of a novel symbol evoked P300 paradigm designed to elicit robust and user friendly event related responses for EEG based user authentication.
2. This study proposes AMEEGNet-CAM, an enhanced deep learning architecture with channel attention mechanisms, enabling efficient less EEG data processing and use of key electrode inputs.
3. The proposed study conducts a comprehensive evaluation across performance metrics, temporal robustness a cross session generalization (EEG data across different sessions), and an efficient system with key electrode setup, demonstrating the framework's practicality, stability, and spoof resistance for real time applications.

This scheme was evaluated and tested across three dimensions: (i) performance metrics such as accuracy, ROC, FAR, and FRR; (ii) robustness through test and retest reliability across sessions (training on session 1 and session 2, testing on session 3); and (iii) the efficiency of the scheme using discriminative feature extraction from a compact EEG setup. The adaptive feature learning along with explainable attention mechanisms was integrated which advances EEG based user authentication for more practical and scalable real world scenarios. To highlight the distinctions between the proposed framework and existing EEG based user authentication approaches, a comprehensive comparison of representative studies is provided in Table 1 with detail explanation.

Table 1. Summary of previous EEG based user authentication studies and the proposed approach highlighting methodology, paradigms, classifiers, and performance

Author	Year	Subjects	Channels	Methodology	Protocol	Classifier
Maiorana et al. [6]	2015	50	32	Resting-state EEG	Eyes-open/closed	SVM
Chen et al. [7]	2016	29	31	RSVP features	RSVP	Shrinkage LDA
Seha et al. [8]	2020	40	32	SSVEP analysis	Auditory evoked	CCA
Rathi et al. [13]	2021	10	10	PCA	Picture-based P300	QDA
Akram et al. [14]	2022	10	32	Signal Processing	P300	RF
Santamaria et al. [16]	2019	20	8	ORSD	P300 speller	CCA
Zeng et al. [17]	2018	45	16	HDCA spatiotemporal	RSVP paradigm	HDCA + GA
Proposed Work	**2026**	**16**	**8**	**Processed Raw EEG**	**Symbol-based P300**	**AMEEGNet**

2 Methodology

2.1 Schematic Overview of the Proposed Approach

The proposed study investigates subject authentication through symbol evoked P300 ERP responses, implemented with two discrete stages: subject identification and authentication. During the subject identification stage, the

AMEEGNet-CAM architecture processes EEG signals that are elicited by symbol based stimuli to generate individual active neural profiles. These active neural profiles captures a discriminative P300 features using the model's integrated channel attention mechanisms (CAM) and explainable feature extraction method. In the subject authentication stage, the system verifies subjects comparing the EEG responses to the stored subject profiles, combing the high inter subject variability of symbol evoked P300 neural waveforms. As illustrated in Fig. 1, this framework ensures robust and efficient subject authentication, combining task specific cognitive neural signatures with AMEEGNet-CAM's adaptive data processing.

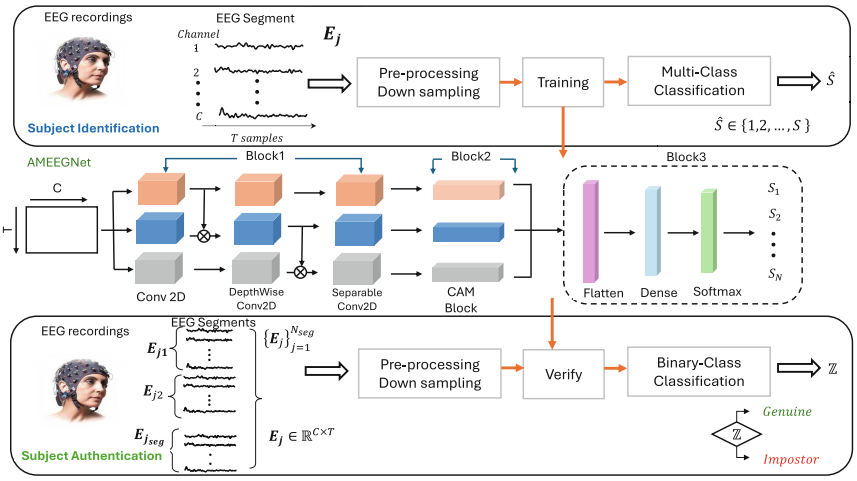

Fig. 1. The proposed EEG-based authentication framework uses symbol-evoked P300 responses in two stages: **(a)** subject identification, where AMEEGNet-CAM employs channel attention to extract unique neural subject profiles emphasizing on key P300 ERP features; and **(b)** subject authentication, where new EEG data is matched against these neural profiles for verification. This scheme integrates the deep learning features with ERP based cognitive neural signatures for robust and efficient subject authentication.

2.2 Problem Statement and Objectives

The given dataset consists of various EEG trials represented as $\mathcal{E} = \{(\mathbf{E}_j, c_j, s_j)\}_{j=1}^{J}$, where each trial $\mathbf{E}_j \in \mathbf{R}^{C\times T}$ corresponds to a segment of the signal recorded from C electrodes over T time points. The variable c_j demonstrates the experimental condition, and $s_j \in \{1, \ldots, S\}$ denotes the subject identity associated with EEG trial j. The proposed method addresses both subject identification and authentication by transforming input EEG signals into discriminative features that capture subject and task related neural signatures.

While the dataset includes condition specific information of the EEG data, the primary focus is to obtain a stable, subject specific EEG patterns for subject related biometric recognition.

The objective of the subject identification task is to determine the true identity s_j from a set of S enrolled subjects with EEG trial $\mathbf{E}_j$. This task is defined, mapping $\mathcal{I} = \mathcal{C} \circ \mathcal{H}$, where $\mathcal{H}$ denotes the feature extraction function which maps EEG epoch into a space representation, and $\mathcal{C}$ is a classifier that allocates identity label.

The objective of the subject authentication task is to verify whether a given EEG trial $\mathbf{E}_j$ relates to a claimed subject identity $s_c \in \{1, \ldots, S\}$. This process is implemented using a binary classifier $\mathcal{A} = \mathcal{V} \circ \Phi$: $\mathbf{R}^{C \times T} \rightarrow \{0, 1\}$, where 1 denotes a trial related to genuine and 0 to impostor. A reliable subject authentication requires robust extraction of subject specific features that minimize false acceptance rates (FAR) and false rejection rates (FRR), thereby ensuring strong EEG biometric environment.

2.3 EEG Data Preprocessing

Filtering and Baseline Correction. To remove non-cognitive relevant brain activity, the EEG data were filtered between (418 Hz) using FIR filter in MNE-Python [24]. This removal of low frequency shifts and noise with high frequency in order to preserve neural EEG signals. This procedure is expressed as:

$$\mathbf{Y}_{i,\mathrm{f}}(t) = \mathcal{B}_{[min,max]} \{\mathbf{X}_i(t)\} \tag{1}$$

where $\mathbf{X}_i(t)$ is the raw epoch, $\mathbf{Y}_i(t)$ is the filtered signal, and $\mathcal{B}_{[4,18]}$ denotes band-pass filtering.

The inherence of EEG signals is low amplitude (μV range) which requires complete preprocessing proedure to reduce interference from both physiological sources (such as eye or muscle movement) and non-physiological artifacts. These EEG recordings were windowed into stimulus locked EEG epochs $\mathbf{E}_j \in \mathbf{R}^{C \times T}$ (200ms pre-stimulus to 800ms post stimulus) which was followed by baseline correction using the prestimulus interval $t \in [-200, 0]$ ms. The baseline corrected epoched signal was computed as:

$$\bar{\mathbf{E}}_j(t) = \mathbf{E}_j(t) - \frac{1}{\Delta\tau} \int_{\tau_b}^{0} \mathbf{E}_j(\tau)\, d\tau \tag{2}$$

where τ_b = -200 ms marks the baseline start, and $\Delta\tau$ = 200 ms is the baseline duration. Here C represents the number of EEG channels, and T denotes the temporal samples per epoch.

EEG Epoch Rejection, Cropping and Downsampling. In order to have EEG data reliability, epochs that contained artifacts of high amplitudes (e.g., from motion or transient noise) were removed. Specifically, those trials where the absolute voltage exceeded from$50\mu V$ on any channel was discarded as well:

$$\max_{t \in [\tau_1, \tau_2]} |\mathbf{E}_j(t)| > 50, \mu\mathrm{V} \tag{3}$$

This strong rejection policy preserved only artifact free EEG segments thus, enhancing the integrity of subsequent neurophysiological analysis of the EEG data.

To remove those EEG trials affected by artifacts such as electromyographic interference or surrounding noise, all EEG segments exceeding that predefined amplitude were excluded. The remaining epochs were cropped to focus on the relevant time window spanning EEG epochs $t \in [0, 800]$ ms are defined by:

$$\bar{\mathbf{X}}_j(t) = \mathbf{X}_j(t) \quad \forall t \in [0, 800] \text{ ms}. \tag{4}$$

Subsequent to this temporal selection, the EEG data was resampled in order to standardize the temporal resolution. These EEG signals were reduced from their normal sampling rate f_s to a resonated rate of $f_t = 64Hz$. This reduction in dimensionality reduced the sample count by a factor of $\frac{f_s}{f_t}$ while it maintained all physiologically relevant frequency components below the $\frac{f_t}{2} = 32Hz$ cutoff.

EEG Vectorization. Each preprocessed EEG epoch $\mathbf{E}_j \in \mathbf{R}^{C \times T}$ was vectorized into a one dimensional feature vector $\mathbf{x}_j \in \mathbf{R}^{C \cdot T}$ using flattening process across channels and time samples. This operation ensures compatibility with conventional machine learning pipelines and is formally expressed as:

$$\mathbf{x}_j = \text{vec}(\mathbf{E}_j) \in \mathbf{R}^{C \cdot T} \tag{5}$$

The data of all the subjects were aggregated into one unified dataset, with each one observated data labeled according to its corresponding subject identity. This standardization of the data supports consistent analysis across classification and evaluation procedures.

2.4 Classification Model

To enhance the feature extraction procedure and improve the ability of the model to focus on the features that are discriminative in symbol evoked based EEG signals, an attention-enhanced multiscale EEGNet architecture, represented as the AMEEGNet-CAM model is proposed. This framework accommodates the classification and subject authentication capabilities. This framework consists of three main components. Block 1, known as multiscale feature extractor which employs three parallel branches inspired by EEGNet with temporal kernel sizes of 32, 64, and 128, dropout rate of 0.5, and L2 regularization ($\lambda = 0.001$) that helps to extract rich temporal spatial patterns from raw EEG data. Block 2 consists of an attention mechanism having a channel attention module (CAM) and a spatial attention module that arbitrarily increase weights and refine the feature learned based on channel and spatial temporal features. Block 3 is the classification stage, includes a global average pooling layer and a dense layer with 256 units and *Swish* activation. This classifier is trained using the adamW optimizer with a learning rate of 5×10^{-5}. The overall AMEEGNet-CAM framework is illustrated in Fig. 1, along with the associated subject identification and authentication scheme.

This classifier employes three parallel EEGNet inspired branches to extract temporal and spatial features from EEG signals containing P300 ERPs for subject identification and authentication scheme. Each branch processes the same 2D input (channels × time samples) and includes: (1) a temporal convolutional layer to capture P300 specific time features using size of various filters (32, 64, and 128), (2) a depthwise convolutional layer to learn spatial patterns across channels, and (3) a separable convolutional layer that combine features efficiently. The depthwise layer uses twice as many filters as the temporal layer, and the separable layer applies filters of size (1, 16). Batch normalization and ELU activation follow both convolutional layers to stabilize learning. Features from all branches are concatenated and passed to a spatial attention module that highlights informative EEG channels. This multi-scale, attention based design accelerates the model ability to decode subject specific ERP patterns for robust subject recognition.

In order to enhance the feature representation for subject identification and authentication based P300 ERPs, the model accomodates a Channel Attention Module (CAM) after the depthwise convolutional layers. CAM adaptively recalibrates channel-wise features by intergrating both average and max. pooled descriptors and move them via a shared MLP that is followed by sigmoid activation to generate attention weights. Following the parallel EEGNet blocks, the resulting feature map X is passed to CAM block which then applies global average pooling (GAP) and global max. pooling (GMP) throughout the spatial dimensions to generate to extract meaningful channel wise information as illustrated in equation:

$$s = \sigma\left(\mathrm{MLP}(\mathrm{GAP}(X)) + \mathrm{MLP}(\mathrm{GMP}(X))\right) \tag{6}$$

where X denotes the input feature map, σ is the sigmoid activation function, and $s \in \mathbf{R}^C$ represents the channel attention weights. These weights are then applied to each channel of the input feature map using element-wise multiplication to recalibrate the features as shown in equation:

$$\tilde{X}_c = s_c \cdot X_c, \quad \text{for } c \in 1, \ldots, C \tag{7}$$

This mechanism allows the model to concentrate on EEG channels most relevant to subject specific P300 ERP responses, hence improving both its discriminative power and robustness while maintaining computational efficiency.

2.5 Implementation of the Subject Authentication Framework

In this research, EEG based subject recognition framework is implemented employing the Enhanced AMEEGNet model to perform subject identification and authentication. The objective of this framework is to validate the effectiveness of the proposed work by recording EEG data from 16 subjects, sessions 1 and 2 was used for training and session 3 for testing. EEG data epochs from all subjects were preprocessed and then analyzed to perform classification based on the learned features.

The EEG based subject identification framework employs a raw EEG segment $\mathbf{E}_j \in \mathbf{R}^{C \times T}$ as input. This undergoes preprocessing steps which includes bandpass filtering, baseline correction, epoch rejection, cropping, and downsampling in order to standardize and clean the EEG signal. The resulted preprocessed EEG segment is then passed through the AMEEGNet-CAM model to integrates channel attention module in order to enhance feature interpretability of the EEG data. EEG data from sessions 1 and 2, the model is trained to learn EEG patterns with specific subject identities. During testing stage the unseen EEG data from session 3 is utilized so the model predicts the subject identity, and labels are evaluated to compare the performance.

The EEG based subject authentication framework is employed to verify either an input EEG segment that originates from a claimed identity is part of the dataset or not. Number of EEG segments $\mathbf{E}_j \in \mathbf{R}^{C \times T}$ undergoes a list of preprocessing stages, including bandpass filtering, baseline correction, artifact rejection, and downsampling, to improve EEG signal quality and reduce artifacts. The preprocessed EEG signal is then fed into the AMEEGNet-CAM model that incorporates channel attention module to improve the interpretability of the features that are extracted. The model work as a binary decider $\mathbf{Z} \in (0, 1)$ where 1 represents a genuine subject and 0 represents an imposter. To assess the generalization ability of the subject authentication system, a leave-one-subject-out (LOSO) cross-validation strategy was implemented, meaning that in each fold the model is trained on all subjects except one and tested on the left-out subject.

2.6 Validation Procedure

Dataset Overview. The proposed scheme analyzes an EEG dataset [25] containing recordings from 16 visually normal, BCI-inexperienced individuals. EEG Neural activity was performed during visual oddball paradigms eliciting P300 potentials. Neural Signals were obtained from 8 optimized electrode sites (Fz, Cz, P3, Pz, P4, PO7, PO8, Oz) according to the 10–20 placement system, with a 256 Hz sampling rate. Sixteen subjects (10 males, 6 females; age 1933, mean 25) completed three sessions with distinct visual paradigms. Facial stimuli elicited stronger P300 responses and higher classification accuracy, while target set size had little effect. These results highlight the value of stimulus optimization and establish the dataset as a benchmark for ERP based subject authentication system.

2.7 Quantitative Metrics

The performance of the proposed methodology was evaluated using accuracy, cohen's kappa, FAR, FRR, and ROC curves. Accuracy measures the overall correctness of predictions, while Kappa accounts for chance agreement between predicted and actual labels:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{8}$$

$$\kappa = \frac{p_o - p_e}{1 - p_e} \tag{9}$$

Here, p_o denotes the observed agreement (accuracy), and p_e represents the expected agreement by chance.

A confusion matrix was employed to provide insights into classification outcomes. Additionally, system security and usability were analyzed using flse acceptance rate (FAR) and false rejection rate (FRR), defined as follows:

$$\text{FAR} = \frac{\text{Accepted Imposters}}{\text{Unauthorized Attempts}} \tag{10}$$

$$\text{FRR} = \frac{\text{Rejected Genuine Users}}{\text{Authorized Attempts}} \tag{11}$$

Finally, ROC curves were used to visualize the trade-off between the true positive rate (TPR) and false positive rate (FPR) under varying decision thresholds.

3 Results and Performance Analysis

This section outlines the evaluation approach used to assess the effectiveness of the proposed framework and explores the impact of EEG channel utilization on subject identification and authentication performance. Results from experiments conducted across multiple subjects are reported, with various performance metrics analyzed to validate the method's capability.

3.1 Implementation

The proposed classifier was trained using the adamW optimizer with default parameters with 100 epochs and a batch size of 64. EEG data were normalized, shuffled, and preprocessed using MNE-Python. Implemented in Python 3.8 with Keras 3.8.0 and TensorFlow 2.18.0, the model features three parallel EEGNet inspired branches with varying kernel sizes to extract multi-scale temporal spatial features. Residual connections and a channel attention module, integrating spatial and channel attention, refine feature representations. This enhances performance for P300-based EEG subject identification and authentication.

3.2 Results and Outcomes of the Subject Recognition Scheme (Identification and Authentication)

This section outlines the evaluation procedure of the proposed approach and examines the impact of EEG channels on recognition performance. A P300 ERP-based authentication test was conducted on EEG data collected from 16 subjects using the enhanced AMEEGNet model. The framework's performance was assessed for both subject identification and authentication tasks. Results

demonstrate the effectiveness of AMEEGNet in accurately capturing discriminative EEG patterns.

Subject identification is approached as a multi-class classification task, assigning a unique class label to each individual to distinguish them within the system. Table 2 presents the individual classification results for each subject following the application of the multi-class identification method. The performance of the model for subject identification was assessed using 5-fold stratified cross-validation, with precision-recall AUC and cohen's kappa as metrics to evaluate its ability to differentiate between various classes. The model obtained an average precision-recall AUC of 0.90- $SD = 0.0125$, p-value $= 0.0076$, demonstrating the ability to distinguish between various classes. The effectiveness of AMEEGNet was evaluated using the multi-class AUC-ROC metric. The AUC values, presented in Fig. 2 highlight the model's strong capability in distinguishing between subjects. The model's performance was assessed for each individual subject, as shown in Fig. 3, to identify areas of strength and challenge. Although the model achieved almost of 90% accuracy for most subjects, subjects 2 and 14 had accuracies below 80%. Nevertheless, with an overall accuracy of 89.25%, the model effectively captures P300 EPRs from EEG signals.

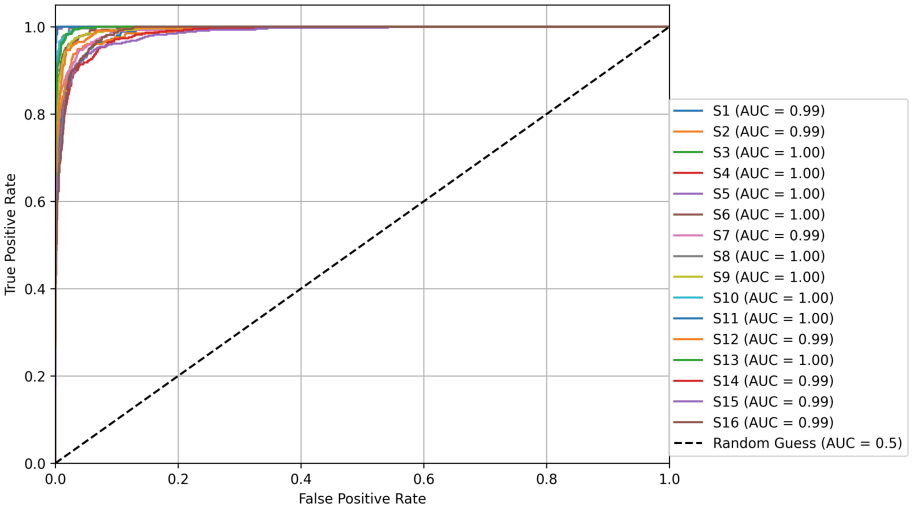

Fig. 2. ROC curves for the AMEEGNet-CAM model across multiple subjects, showing the true positive rate (TPR) versus false positive rate (FPR). AUC values indicate the model's classification performance, with the dashed line representing random guessing (AUC = 0.5).

Subject authentication was assessed using AMEEGNet-CAM by analyzing P300 ERPs from 16 subjects, making the task as a binary classification in which one subject represented the genuine user (1) and the other 15 were grouped as impostors (0). Mean accuracy of 89.25% was obtained for all subjects. The

Table 2. Subject identification performance of AMEEGNet-CAM on 16 subjects

Subject	Precision	Recall	F1-score	k-value
S01	0.72	0.93	0.81	0.80
S02	0.92	0.78	0.85	0.84
S03	1.00	0.99	0.99	0.99
S04	1.00	0.98	0.99	0.99
S05	0.98	0.95	0.97	0.96
S06	0.89	0.95	0.92	0.91
S07	0.81	0.94	0.87	0.86
S08	0.96	0.97	0.96	0.96
S09	0.95	0.89	0.92	0.92
S10	0.96	0.95	0.96	0.95
S11	0.93	0.98	0.95	0.95
S12	0.93	0.91	0.92	0.91
S13	0.94	0.96	0.95	0.95
S14	0.96	0.72	0.83	0.81
S15	0.85	0.87	0.86	0.85
S16	0.85	0.84	0.84	0.83

highest accuracy was obtained for the subject S03 and S04 (100%), while the lowest levels of accuracy were observed for the subject S01 (80%), which seems to be an outlier. The performance of the model for each user was evaluated using the proposed method, as outlined in Table 3. Figure 4 illustrates the classification performance of the proposed model, AMEEGNet-CAM, across 16 subjects using accuracy, precision, and f1 Score. The model achieved an average accuracy of 89.25%, demonstrating strong performance in EEG-based subject authentication. Individual accuracies ranged from 0.81 to 1.00, with precision values mostly above 0.90, indicating effectiveness in reducing false positives. Subjects S03 and S04 showed perfect classification, while S05 and S14 had slightly lower recall values, impacting their f1 scores. Overall, the model maintained high precision and stable performance across all subjects, showcasing its robustness and generalization capability. To further evaluate AMEEGNet's performance, sensitivity and specificity across 16 subjects were analyzed as shown in Fig. 5. Sensitivity ranged from 0.80 to 0.99, indicating consistent true positive. Specificity values were tightly clustered between 0.93 and 0.96, showing robustness in minimizing false positives. The compact distribution in the plot highlights model's reliability and stability, with minimal subject-specific variance. The high and balanced sensitivity and specificity underscore model's effectiveness for EEG-based subject authentication, ensuring accurate detection and minimal false alarms. The proposed scheme exhibited a minimal false acceptance rate (FAR), averaging 3.5% ($\pm$ 2.73%), and a false rejection rate (FRR) of 10.81% ($\pm$ 6.24%) across all

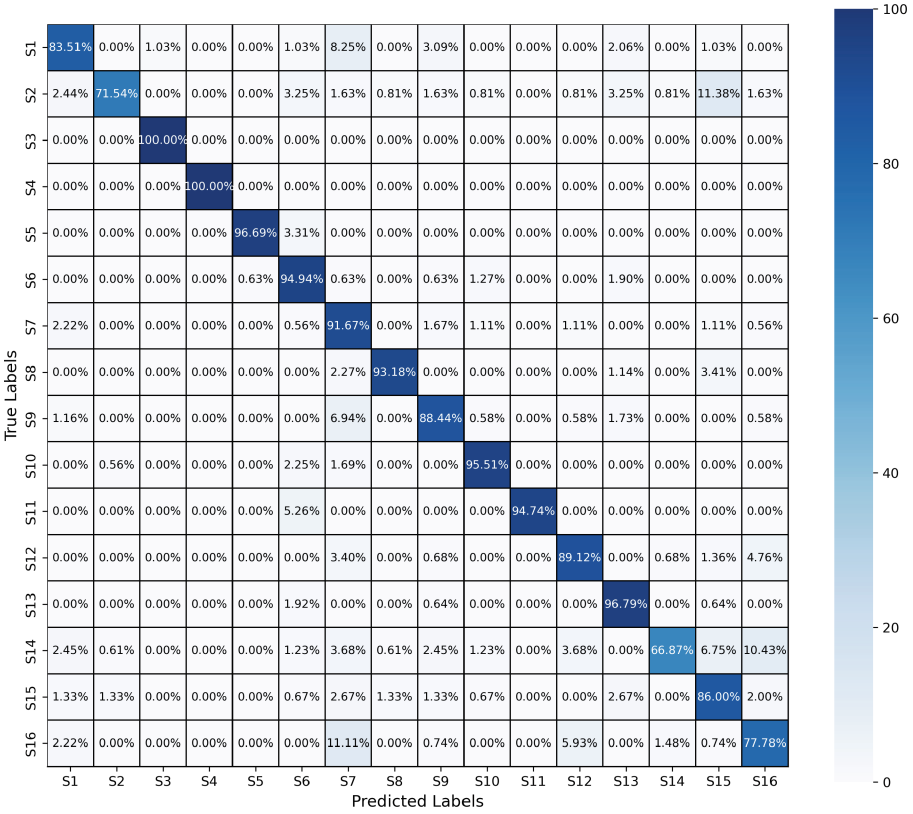

Fig. 3. Identification Accuracy of subjects. Confusion matrix illustrating the classification performance of the AMEEGNet-CAM model for subject identification, highlighting the accuracy and distribution of correct and incorrect predictions across subject classes.

subjects. This indicates its robust capability to accurately differentiate between genuine and impostors.

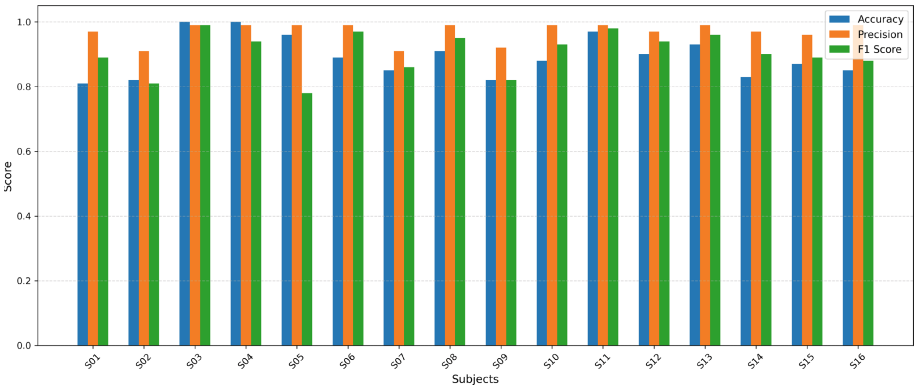

Fig. 4. Performance evaluation of the proposed model across all 16 subjects, illustrating key classification metrics including accuracy, precision, and f1 Score. The results indicate consistent and reliable performance.

Table 3. Subject authentication performance metrics (k-value, FRR, FAR, and accuracy) across 16 subjects for the 2021 samples using the AMEEGNet-CAM model

Subject	k-value	FRR	FAR	Accuracy
S01	0.18	0.03	0.03	81%
S02	0.54	0.20	0.02	82%
S03	0.99	0.01	0.01	100%
S04	0.99	0.03	0.01	100%
S05	0.76	0.06	0.01	96%
S06	0.28	0.04	0.01	89%
S07	0.23	0.16	0.09	85%
S08	0.47	0.08	0.01	91%
S09	0.17	0.19	0.08	82%
S10	0.54	0.12	0.01	88%
S11	0.42	0.03	0.01	97%
S12	0.47	0.08	0.03	90%
S13	0.65	0.07	0.07	93%
S14	0.30	0.16	0.01	83%
S15	0.33	0.16	0.03	87%
S16	0.23	0.20	0.04	85%

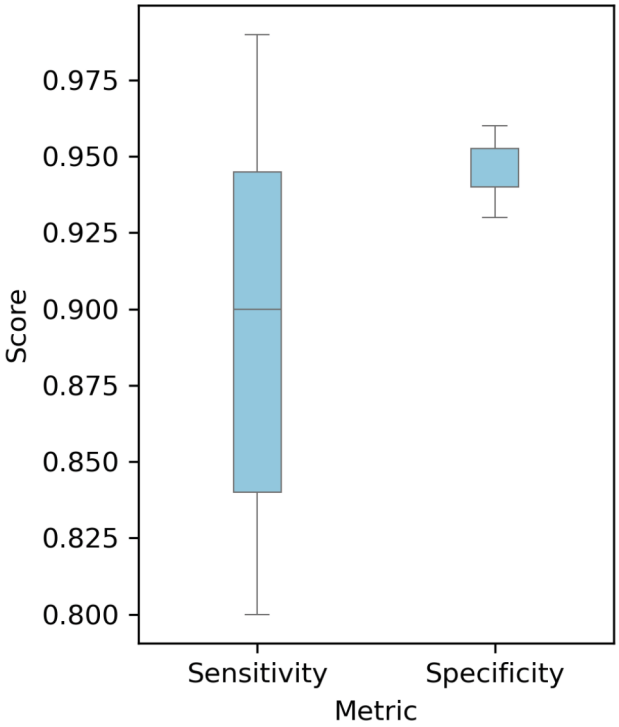

Fig. 5. Comparison of Sensitivity and Specificity values across all 16 subjects, highlighting the distribution and consistency of the model's performance.

4 Discussion and Conclusion

This study explored the use of an Enhanced AMEEGNet model for subject identification and authentication using P300 event related potentials (ERPs), achieving an overall accuracy of 89.25% and Cohen's kappa score of 0.88. As the first work to apply Enhanced AMEEGNet specifically to P300-based biometric authentication, it marks a novel contribution to EEG-based biometric research. The model showed strong generalization across users, with most individuals achieving high classification accuracy. Although some subjects were more challenging, the model remained effective, supporting its suitability for real-world application.

The proposed model demonstrates promising real-time potential for EEG-based authentication systems, thanks to its ability to capture spatiotemporal EEG dynamics efficiently. Its use of P300 ERPs as a biometric trait offers a secure and non-invasive approach for identity verification. Future research should focus on improving performance for low-performing users and exploring other ERP components, such as visual evoked potentials, to enhance biometric robustness. This work establishes a solid foundation for using Enhanced AMEEGNet in next-generation cognitive biometrics.

References

1. Wash, R., Rader, E.: Prioritizing security over usability: strategies for how people choose passwords. J. Cybersec. **7**(1), tyab012 (2021)
2. Galbally, J., Marcel, S., Fierrez, J.: Biometric antispoofing methods: a survey in face recognition. IEEE Access **2**, 1530–1552 (2014)
3. Armstrong, B.C., Ruiz-Blondet, M.V., Khalifian, N., Kurtz, K.J., Jin, Z., Laszlo, S.: BrainPrint: assessing the uniqueness, collectability, and permanence of a novel method for ERP biometrics. Neurocomputing **166**, 59–67 (2015)
4. Khalil, A.E.K., Perez-Diaz, J.A., Cantoral-Ceballos, J.A., Antelis, J.M.: Unlocking security for comprehensive electroencephalogram-based user authentication systems. Sensors **24**(24), 7919 (2024)
5. Kumari, P., Vaish, A.: Brainwave based user identification system: a pilot study in robotics environment. Robot. Auton. Syst. **65**, 15–23 (2015)
6. Maiorana, E., La Rocca, D., Campisi, P.: On the permanence of EEG signals for biometric recognition. IEEE Trans. Inf. Forensics Secur. **11**(1), 163–175 (2015)
7. Chen, Y., et al.: A high-security EEG-based login system with RSVP stimuli and dry electrodes. IEEE Trans. Inf. Forensics Secur. **11**(12), 2635–2647 (2016)
8. Seha, S.N.A., Hatzinakos, D.: EEG-based human recognition using steady-state AEPs and subject-unique spatial filters. IEEE Trans. Inf. Forensics Secur. **15**, 3901–3910 (2020)
9. Jin, J., et al.: Optimized stimulus presentation patterns for an event-related potential EEG-based brain–computer interface. Med. Biol. Eng. Comput. **49**, 181–191 (2011)
10. Combaz, A., Van Hulle, M.M.: Simultaneous detection of P300 and steady-state visually evoked potentials for hybrid brain-computer interface. PLoS ONE **10**(3), e0121481 (2015). https://doi.org/10.3886/ICPSR30122.v2
11. Combaz, A., Van Hulle, M.M.: Simultaneous detection of P300 and steady-state visually evoked potentials for hybrid brain-computer interface. PLoS ONE **10**(3), e0121481 (2015)
12. Khalil, A.E.K., Pérez-Díaz, J.A., Cantoral-Ceballos, J.A., Antelis, J.M.: Enhancing user authentication through EEG based P300 speller response. In: In: Martínez-Villaseñor, L., Ochoa-Ruiz, G. (eds) Advances in Soft Computing. MICAI 2024. LNCS, vol. 15247, pp. 124–141. Springer Nature Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-75543-9_10
13. Rathi, N., Singla, R., Tiwari, S.: A novel approach for designing authentication system using a picture based P300 speller. Cogn. Neurodyn., 1–20 (2021)
14. Akram, F., Alwakeel, A., Alwakeel, M., Hijji, M., Masud, U.: A symbols based BCI paradigm for intelligent home control using P300 event-related potentials. Sensors **22**(24), 10000 (2022)
15. Zhang, S., Sun, L., Mao, X., Hu, C., Liu, P.: Review on EEG-based authentication technology. Comput. Intell. Neurosci. **2021**(1), 5229576 (2021)
16. Santamaría-Vázquez, E., Martínez-Cagigal, V., Gomez-Pilar, J., Hornero, R.: Asynchronous control of ERP-based BCI spellers using steady-state visual evoked potentials elicited by peripheral stimuli. IEEE Trans. Neural Syst. Rehabil. Eng. **27**(9), 1883–1892 (2019)
17. Zeng, Y., et al.: EEG-based identity authentication framework using face rapid serial visual presentation with optimized channels. Sensors **19**(1), 6 (2018)
18. Chan, H.L., Kuo, P.C., Cheng, C.Y., Chen, Y.S.: Challenges and future perspectives on electroencephalogram-based biometrics in person recognition. Front. Neuroinform. **12**, 66 (2018)

19. Maiorana, E., Solé-Casals, J., Campisi, P.: EEG signal preprocessing for biometric recognition. Mach. Vis. Appl. **27**, 1351–1360 (2016)
20. Mao, Z., Yao, W.X., Huang, Y.: EEG-based biometric identification with deep learning. In: 2017 8th International IEEE/EMBS Conference on Neural Engineering (NER), pp. 609–612. IEEE (2017)
21. Roy, A.M.: An efficient multi-scale CNN model with intrinsic feature integration for motor imagery EEG subject classification in brain-machine interfaces. Biomed. Signal Process. Control **74**, 103496 (2022)
22. Lawhern, V.J., Solon, A.J., Waytowich, N.R., Gordon, S.M., Hung, C.P., Lance, B.J.: EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces. J. Neural Eng. **15**(5), 056013 (2018)
23. Wu, X., Chu, Y., Li, Q., Luo, Y., Zhao, Y., Zhao, X.: AMEEGNet: attention-based multiscale EEGNet for effective motor imagery EEG decoding. Front. Neurorobot. **19**, 1540033 (2025)
24. Gramfort, A., et al.: MEG and EEG data analysis with MNE-Python. Front. Neuroinform. **7**, 267 (2013)
25. Chailloux Peguero, J.D., Mendoza-Montoya, O., Antelis, J.M.: Single-option P300-BCI performance is affected by visual stimulation conditions. Sensors **20**(24), 7198 (2020)

Optimizing Intrusion Detection for SDN-Enabled IoT Using of Lightweight Deep Learning Architectures

Mustafa T. Saleh[1] and Ali H. Hamad[2(✉)]

[1] Informatics Institute for Postgraduate Studies, University of Information and Communication Technology, Baghdad, Iraq
ahamad@kecbu.uobaghdad.edu.iq

[2] Biomedical Application Department, College of Artificial Intelligence, University of Baghdad, Baghdad, Iraq
phd202230706@iips.edu.iq

Abstract. The rising use of the Internet of Things (IoT) has led to the use of Software-Defined Networking (SDN) as a scalable framework that organizes large amounts of heterogeneous data traffic. However, the centralization nature of SDN present sturdy security gaps, with they being distributed denial-of-service (DDoS) attacks to advanced intrusion vectors, which threaten the privacy of governmental and trade confidential information. Although Deep Learning (DL) represents a disruptive step towards the optimization of a network, as well as remediation, the sheer volume of computation often conflicts with the constrained capabilities of edge gateways and IoT endpoints. This work poses the question of the effectiveness of four canonical DL architectures, including deep neural networks (DNNs), recurrent neural networks (RNNs), convolutional neural networks (CNNs), and long short-term memory units (LSTMs) in an SDN-based intrusion detection system. The exploration, based on the SICA dataset, which contains more than 2.2 million incident records of 22 different typologies of attacks, follows three major model-compressing approaches, namely structural pruning, weight clustering, and quantization. The empirical results indicate that even though LSTM provides the best classification accuracy, over 95$\mathscr{R}$, a pruned DNN provides a good trade off in the deployment of edges with a ten times processing throughput with insignificant loss in detection confidence. Such findings provide a stringent blueprint that organizations aiming to adopt real-time, resource-saving cybersecurity defenses in the current digital ecosystem may use.

Keywords: IDS · IoT · Deep learning · Pruning · QAT · Clustering

1 Introduction

As the Internet of Things (IoT) grows at a fast rate, scores of trillions of interconnected devices are producing enormous amounts of data whilst it is demanding efficient, reliable, and adaptive communication strategies. IoT is a heterogeneous entity networks,

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 641–656, 2027.
https://doi.org/10.1007/978-3-032-30518-3_40

bandwidth constraints, and latency sensitivity, and there are significant difficulties with energy constraints of equipment to traditional paradigms of communication. The use of algorithms has been done to boost efficiency in communication in solving optimization issues. However, traditional optimization models are not always in line with the dynamic, data-oriented aspect of the Internet of Things systems. As a result, deep learning proposes new optimization procedures and takes part in their refined development. Smart learning mechanisms facilitate real time reactions and decision making that are data based. Deep optimization through learning has the potential of being transformative. IoT communication is limited by coordination of network resources, coordination of devices and intelligent provision of data. Thus, the systematic research on how deep learning methods could be applied to creating more efficient optimization algorithms of the Internet of Things is urgently required [1]. The protection of sensitive information, intellectual property, government, and trade secrets are important in the modern digital environment. Cybersecurity issues become increasingly challenging due to the technical progress. Despite some significant interest in cybersecurity education and training, researchers continue to struggle with effectively surveying different approaches. This research intends to critically evaluate frameworks and ways to develop and establish new cybersecurity training solutions to individuals and organizations [2].

SDN has become a big step in the evolution of the area of network architecture, a new era of efficiency, flexibility and programmability. At the heart of SDN is the separation of the control plane and the data plane, which is a groundbreaking idea that enables network immediate power of administrators over traffic in the network flow and management [3]. Nor can this paradigm shift, though conducive also, optimizations and agile network configurations to illuminate a range of security risks present in the SDN architecture. The SDN control is centralized coupled with the programmability and openness of plane. SDN environments are especially made by network interfaces prone to various online attacks, including Denial of Service (DDoS) attacks to advanced intrusion attempts. Considering such challenges, there is need of strong, dynamic, and proactive security is more eminent than ever. Among the bombard of the security solutions, it is IDS an essential line of defense, which is intended to identify, evaluate and control the bad practices in the network. However, the peculiarities of SDN demand a re-consideration and redesign of conventional IDS strategies in order to suitably manage the vulnerabilities peculiar to SDN infrastructures. The further tightens cyber threats, which are becoming more sophisticated this necessity, which demands novel solutions which can not only identify known attack vectors, but also predict and eliminate emerging threats [4].

2 Related Work

This part presents the latest developments in Intrusion Detection Systems (IDS) and draws particular attention to such methods as quantization, pruning, and optimization to allow the use of them on resource-restricted Internet of Things (IoT) and edge computing devices.

Idrissi, et al. (2022) [5] discussed the possibility of implementing Deep Learning-oriented Host-Intrusion Detection System (DL-HIDS) on commercial IoT devices. The authors used various optimizations that were specific to specific hardware to achieve

constrained memory and inference timing constraints. Luo (2023) [6] suggested a distributed SDN-based IDS, which makes use of a decision tree that is optimized according to the Black Hole Optimization (BHO) algorithm. Pruning of the model is done with the help of BHO to maximize the accuracy of each sub-network. Sharmila and Nagapadma (2023) [7] introduced Quantized Autoencoder (QAE) which is an unsupervised model that includes pruning, clustering, and integer quantization (uint8 and float16 versions in particular). Their model of QAE-u8 has a reduction of 92.23% in memory occupancy with high detection accuracy. Wang et al. (2023) [8] introduced a lightweight model named DL-BiLSTM, which is based on Deep Neural Networks (DNN) and Bidirectional Long Short-Term Memory (BiLSTM). It employs Incremental Principal Component Analysis (IPCA) to minimize dimensionality and dynamic quantization so as to minimize computation overheads on internet of things devices.

Acharya, et al. (2024) [9] suggested a design architecture that is based on Quantization-aware Neural Architectural Search to evolve quantized neural network models automatically. Such models are much smaller than the state-of-the-art networks and are designed with the FPGAs in mind. Huynh and Nguyen (2024) [10] proposed a better three-step training model of edge computing devices. It uses model compression and a better pruning techniqueology to compress model sizes yet maintain high accuracy on the UNSW-NB15 dataset. Bouayad, et al. (2024) [11] proposed Lightweight-Fed-NIDS that is a federated learning system with structured model pruning. One of the contributions is the single time calculation of a pruning mask to optimize local models without client access information, and pruning size is cut by 90%. Morea and Kachavimatha (2025) [12] suggested the method of Software-Defined Networking (SDN)-based intrusion detection based on meta-heuristic feature selection combined with a K-Nearest Neighbors (KNN) classifier to detect malicious actions in real-time. Rangsikunpum, et al. (2025) [13] concerned with in-vehicle networks, and the article introduces an ultra-lightweight IDS implemented with the use of Binarised Neural Networks (BNN). The authors reduced the number of parameters by 91.07% through network pruning to get a loss (accuracy) of 0.1% (Table 1).

Table 1. Comparative of previous work

Ref.	Methodology	Dataset(s)	Optimization Techniques	Key Results
Idrissi, et al. (2022)	Deep Learning-based Host-IDS (DL-HIDS)	Specific commercial IoT traffic	Hardware-specific optimizations for edge deployment	Optimized for memory consumption and inference timing on commercial IoT devices.

(*continued*)

Table 1. (*continued*)

Ref.	Methodology	Dataset(s)	Optimization Techniques	Key Results
Sharmila Nagapadma (2023)	Quantized Autoencoder (QAE)	RT-IoT2022 (Real-Time IoT 2022)	Pruning, clustering, and integer quantization (uint8 and float16)	QAE-u8 achieved a 92.23% reduction in memory size and 70.01% lower average memory utilization.
Luo (2023)	Distributed SDN-based IDS using optimized forests (Decision Trees)	NSL-KDD and UNSW-NB15	Black Hole Optimization (BHO) for tree pruning and node splitting	Achieved 99.2% accuracy on NSL-KDD and 97.2% on UNSW-NB15.
Wang, et al. (2023)	DL-BiLSTM: Combined DNN and Bidirectional LSTM	CIC-IDS2017, N-BaIoT, and CICIoT2023	IPCA for dimensionality reduction; Dynamic Quantization for cell structure	Surpassed traditional models in performance while maintaining lower complexity and model size.
Acharya, et al. (2024)	Quantization-aware Neural Architectural Search (NAS)	Not explicitly stated in abstract snippet	Automated training and evolution of quantized neural networks	Models are 1,000x smaller than state-of-the-art; 2.3x to 8.5x fewer LUTs when deployed on FPGAs.
Huynh & Nguyen (2024)	Enhanced Three-Stage Training (DSD) paradigm	UNSW-NB15	Enhanced pruning methodology and model compression	Significantly reduced model dimensions while maintaining accuracy levels of larger models.
Bouayad et al. (2024)	Lightweight-Fed-NIDS: Federated Learning with CNNs (ResNet, VGG)	UNSW-NB15, USTC-TFC2016, and CIC-IDS-2017	Structured model pruning via a one-time pruning mask	90% reduction in model size; 3X acceleration in training time; ~99% detection rate.

(*continued*)

Table 1. (*continued*)

Ref.	Methodology	Dataset(s)	Optimization Techniques	Key Results
Morea and Kachavimatha (2025)	SDN IDS using meta-heuristic feature selection	Not explicitly stated in abstract snippet	Meta-heuristic optimization for feature selection	Enabled real-time analytics and improved classification for SDN vulnerabilities.
Rangsikunpum, et al. (2025)	Pruned Binarised Neural Networks (BNN) for In-Vehicle Networks	CH and SA datasets	Binarization (1-bit weights) and network pruning	91.07% reduction in parameters with only 0.1% accuracy loss; optimized for FPGA/Jetson Orin Nano.

3 Methodology

Designing a lightweight Intrusion Detection System (IDS) for the Internet of Things (IoT) requires a methodology that balances high-security precision with the harsh hardware constraints of edge devices. Unlike cloud-based systems, an IoT-ready IDS must minimize its memory footprint and computational latency to operate on microcontrollers or gateways. The following methodology outlines the systematic transition from a complex Deep Learning (DL) model to a lightweight, deployable security solution.

3.1 Intrusion Detection System

The Intrusion Detection Systems (IDS) are complex security systems that continuously scan network traffic and system activities to identify uncharacteristic patterns or familiar threats, and alert administrators when they are detected. These can be categorized in two main dimensions namely the method used to detect and the extent of monitoring. This is based on the detection methodology into two broad categories of IDS. The former is Signature Detection (also known as Misuse based Intrusion Detection). This approach is based on a continuously maintained list of signatures, or patterns or fingerprints that identify a known attack and suspicious behavior. The IDS scans network packets or system events and compares it against this repository of archived signatures; a hit generates an alert and as a result, identifies a recognized threat. The second is the Behavioral Anomaly Detection which is a method that creates a benchmark of normal network or system behavior. Any difference with this set normal base is perceived as an anomaly and a possible intrusion. To create a dynamic profile of normal activities over time, a statistical model, machine learning algorithm or other methods are used, and the real-time activity is compared against it [14].

IDS is also bifurcated when the classification is done based on the scope of monitoring. First, the Network-Based Intrusion Detection Systems (NIDS) monitor the traffic of a network segment and analyze it in real-time to detect any threats that pass through the network. Placed at strategic points in the network, such as at the perimeter or major routing points such as switches and routers, NIDS will use network taps or switch mirror ports (SPAN) to obtain all packet data including headers and payloads on the entire watched portion of the network. Secondly, Host-Based Intrusion Detection Systems (HIDS) are software agents that are run on a specific host (servers, laptops, workstations) and thematically on that specific machine monitors activity. The agent of a HIDS is directly embedded to the operating system of every host machine [15].

3.2 Light Weight Techniques

Typically, edge devices cannot perform refined computations, making it critical to evolve and apply more efficient, lightweight models suited to their limitations. Model pruning, quantization, and clustering are the most efficient optimization methodologies that are exploited to make deep transfer learning models more lightweight [16]. These methodologies work on optimizing both the models' size and performance by eliminating any unnecessary parameters and making the computations needed during inference much easier. Decreasing the models' size represents one of the many benefits of lightweight schemes, as well as extended accessibility, improved privacy levels, and reduced costs. However, there is a reduction in the overall speed of the model. There are also downsides, like complexity, limited flexibility, low accuracy, and trade-offs between models' accuracy, speed, and size. There are three main deep learning model light weight techniques, in the following a brief description for each:

(i) Pruning

Pruning is a type of optimization method that can be used to make a model smaller eliminating irrelevant connections, neurons or parameters. Connection pruning consist of eliminating weights with values equal to or almost equal to zero in an attempt to reduce the memory needs of the whole network. The model that is pruned can be re-trained and in fact have better accuracy than the original model. Figure 1 shows a basic sample of a pre-pruned model i.e. with all neurons connected and the result simplified model with less connections. One of the most widespread pruning techniques is magnitude-based pruning and has attained significant performance with respect to decreasing memory demands of machine learning models without sacrificing the accuracy of the scheme. This technique consists of removing weights with values that are below a specified threshold of a trained model. Pruning may also be adopted by specifying the pruning ratio (PR) of the model which specifies the weight percentages to be zeroed in a whole network or the individual layers to be trimmed by the use of a pruning schedule [17]. The pruning process could be summarized as follows:

$$w'_{ij} \begin{cases} w_{ij} \ \ if \ \ w_{ij} \geq \theta_0 \\ 0 \ \ if \ \ w_{ij} < \theta_0 \end{cases} \tag{1}$$

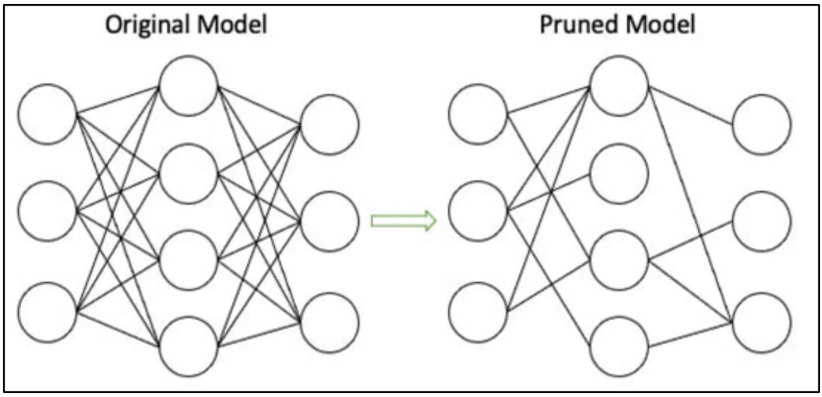

Fig. 1. Pruning light weight technique [17].

(ii) Quantization

Quantization is a model optimization method, the goal of which is to reducing the size of a DL model by quantizing weight values to a reduced precision as bits, that normally is between 32 bits and 16 bits. Moreover, it can be possible to represent binary floating-point numbers [18]. Quantization is separable into two major classes post-training quantization (PTQ) without retrain the model. Secondly, quantization-aware training (QAT) which required retrain the model. Figure 2 indicates the FP32, FP16, FP8 and INT8 value representations in a computer system. It is noted that various representations make use of varying bits count. The precision is associated with the size of bits on which a value is represented. The closer to reality the depiction is, the larger the size of the memory required to store the model. This process is executed by the application of the quantization function q, which is defined in a point wise manner.

$$quantize_{range} = \left[2^{(num\ bit-1)}, 2^{(num-1)} - 1\right] \tag{2}$$

$$scale\ factor\ \ weight = \frac{\max(weight)}{2^{(num_{bit}-1)}} \tag{3}$$

$$quantized\ value = round\left(\frac{value}{scale_{factor}}\right) \tag{4}$$

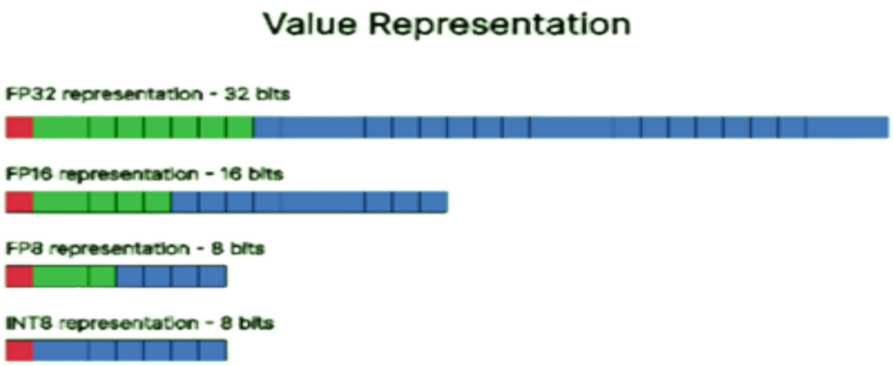

Fig. 2. Quantization principle [18].

(iii) Weight Clustering

Weight clustering or weight sharing is an optimization method which is composed of finding a representative value of a cluster or a group of weights, thereby reducing the amount of different possible values that would be represented in the network. It means that many different. The neurons within a network have the same weight. An example of simple weight clustering is shown in figure 3 that clustered sixteen weights in four categories identified by a color. Four reference values are associated with this case the average of the values of each cluster, and the weight of each is also quantized to one of the reference values mentioned above. The green cluster as an example contains five elements with different values. The average of these components is 0.15 which is the reference value of this cluster. Having used weight sharing, everything in green has a weight of 0.15 weight value. In the case of this particular cluster, the parameters are substituted by a single one. An integer that is the reference index and FP32 value. In the example, taking all the clusters into account, the network with weight clustering has instead of sixteen different FP32 values substitutes these special values with four FP32 values, therefore, optimizing the memory storage [19].

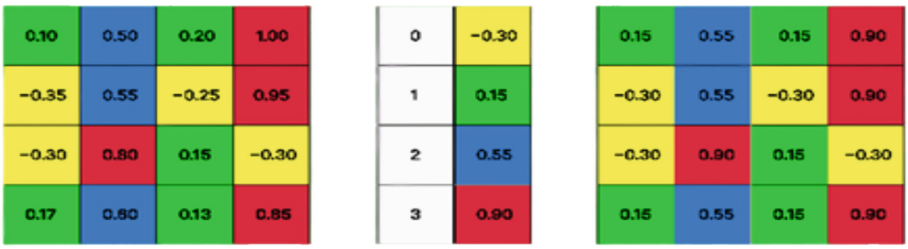

Fig. 3. Weight clustering technique [19]

4 Proposed System Design

To achieve an empirically sound and holistic dataset, we are compelled to incorporate a wide range of application services on the experimental corpus. This approach allows to simulate subtly attack vectors that are contemporary and Internet-based and are embedded in a software-defined networks (SDN) setting. Besides, the attack scenarios should include more advanced paths which will cross several SDN elements. Available academic sources cover a significant range of such situations, both in-house and external with regards to SDN ecosystem. To capture this, we outline four sub-domains we are going to deploy to Mininet operating Ubuntu 20.04 LTS, hence establishing concrete topology. These sub-domains are linked through two Ryu controllers and two switching the sub-domains through four OpenFlow vSwitch (OVS) switches. A Unix-like open-source Python controller, the Ryu framework stands out by being able to support the challenging needs of modern networks infrastructures and remain scalable at the same time. Its wide interoperability with divergent protocols of SDN and functional capabilities make it useful in disparate network designs, as well as it is susceptible to heterogeneous integration. The first two sub-domains simulate a classic network stack, which serves a broad array of services including HTTP and FTP servers. On the other hand, the other two sub-domains simulate an Internet of Things (IoT) network specifically intended to be interrogated in the course of experimentation, both indoors and outdoors. The structure of the proposed multi-controller SDN is supported in figure 4.

In this work, four deep learning models, i.e. DNN, RNN, CNN and LSTM were used. These models are determining the suggested intrusion detection system (IDS) and performance light weighting methods benchmarking. The hyperparameters of the suggested algorithms are listed in Table 2. After the training of each DL model, it was pruned, quantized, and clustered to become light weighted. The proposed algorithms were tested in terms of model sizes before and after light weighting, other training and testing measurements like accuracy and loss. An open-source software-defined networking (SDN) controller, called RYU controller, along with Mininet, a network emulator, were used to create an SDN network environment, thus a realistic SDN architecture was emulated. The simulated SDN created standard traffic as well as attack traffic of different types. It can be assumed that Metasploit and hping3 might have been used to model a wide variety of attack vectors such as exploits, denial-of-service attacks, etc., but ping utility was used to send normal traffic. An unprocessed network traffic data including various attack and normal traffic was captured by a packet analyzer like Wireshark. To create a single pcap (Packet Capture), a ubiquitous format to save network packets, all of the separate traffic captures were amalgamated into a single pcap. The statistical analysis of the suggested implementation of SICA dataset is available in the Table 3.

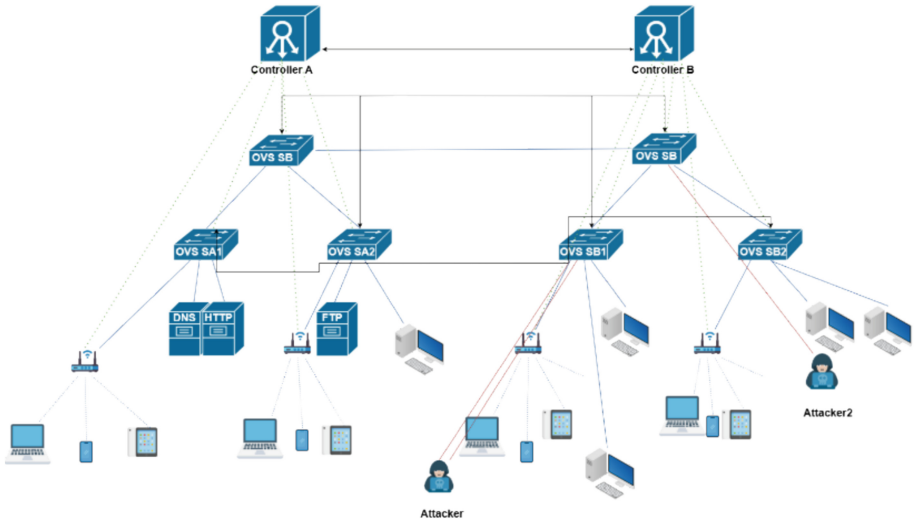

Fig. 4. Proposed multi-controller SDN

Table 2. Deep learning models hyperparmeters

Model	Layers	Neurons	Hyperparameters
LSTM	Input → LSTM (5x) → Output	256 → 128 → 128 → 64 → 32 → Output	Activations: Tanh (hidden), Softmax (output).
CNN	Input → Conv1D → MaxPool → Conv1D → MaxPool → Conv1D → MaxPool → Flatten → Dense (2x) → Output	256 → 128 → 128 → 64 → 32 → Output	Activations: ReLU (hidden), Softmax (output), Kernal size = 2,3.
RNN	Input → GRU → Dense (2x) → Output	64 → 64 → 32 → Output	Activations: ReLU (hidden), Softmax (output).
DNN	Input → Dense (5x) → Output	256 → 128 → 128 → 64 → 32 → Output	Activations: ReLU (hidden), Softmax (output).

Table 3. Statistical analysis of proposed SICA

Parameters	Description
Dataset	SICA
No. of Records	2,298,903
No. of Features	84
No. of attack classes	22
Dataset size	1262 MB
attack types	'Arp_spoofing' 'apache_range_dos' 'arp_sweep' 'avahi_portzero' 'brute_dirs' 'ddos' 'fragmentation_attack' 'http_flood' 'http_get_uri_long' 'http_get_uri_strings' 'normal' 'ntp_prot' 'ntp_protocol_fuzzer' 'slowloris_attacks' 'smb2_negotiate_corrupt' 'smb_tree_connect_corrupt' 'ssh' 'ssh_version_corrupt' 'synflood' 'teardrop' 'udp_probe' 'udp_reflection'

5 Results and Discussion

Based on the results shown in figure 5 a,b, and c from the SICA dataset across the three optimization techniques, RNN models were seen as close contenders but the LSTM had over 95% accuracy throughout the board but the RNN was seen to be much more efficient in terms of model size since it was much lighter (about 10.8 KB) than the DNN and LSTM (about 2627 KB). Conversely, the 1D-CNN showed the worst performance, with the accuracies of about 71–72% which means that it would not be suitable in the current dataset. By and large, despite some trends to use quantization as a method of hardware acceleration, pruning and clustering were found to be equally effective when it comes to saving the accuracy of the larger DNN and LSTM models, though the RNN is the best solution where extreme memory limitations are involved due to its relatively small footprint.

The training process based on quantization-awareness (QAT) generated numerical values only on the DNN model, which did not converge when using LSTM, RNN, and 1DCNN. Despite the results of the DNN, its accuracy dropped to 83.6% (QAT) as compared to 96.5% (pruned). This is an indication that the DNN was stable enough to be trained, but quantization was a big performance cost. Since LSTMs and RNNs are more sensitive than DNNs, they must have reached a zero utility or technical error at some stage of training. A number of technical reasons give rise to the common failure of QAT to deliver results - or get circumvented - on recurrent and convolutional architectures compared to plain dense networks:

- Recurrent Networks (RNN and LSTM): These networks are sequential and process data at a time, with an output of one time step serving as an input of the next stage. Fake quantization nodes are customarily inserted in QAT to simulate low-precision arithmetic. The microscopic error of the quantization process becomes magnified at every time step, often leading to gradient instability or exploding/vanishing gradients, and the training process will fail or the model will simply blow up.

- Dense Networks (DNNs): Dense feed-forward networks have a generally smaller number of depth per sample than the unrolled time-step sequences of a LSTM. They are therefore much more resistant to the added noise caused at the quantization-aware training pipeline.
- RNN/CNN Layer Complexity: Quantization of LSTM or 1DCNN layers is more complex since to quantize them, one must quantize several internal elements at once: gates, cell states, kernels and biases. In case the particularisation of these layers in the one's code does not fit in the QAT wrapper, the training will either crash, or yield invalid results (e.g., NaN).

Table 4 shows a comparative study of four canonical deep leaning frameworks (LSTM, RNN, DNN, and CNN) against their beheaded, "lightweight" versions. Through the implementation of structural pruning, weight clustering and quantization aware training (QAT) the investigation aims to determine an equilibrium between the confidence of classification and computational expediency. The empirical results highlight a strong trade-off: while optimizations greatly reduce the latency, they often introduce a phenomenon we call "label drift," where specific instances of the attack will be re-labeled into more generic categories. Baseline performance for both the LSTM and RNN architectures are quite good, especially on detection of DDoS and BFA attacks. Nonetheless, these temporally-aware models are excessively susceptible to clustering procedures. When clustering is in force both models lose the ability to distinguish BFA from a probe attacks and instead classify them as DDoS. This suggests that the shared weight enforced in clustering disrupts the fine gate mechanisms relevant to memory essential for achieving sequence-learning based inference. Moreover, while pruning is typically followed by the goal of speeding up the inference process, the LSTM/RNN mod results show a small amount of latency, from 0.06 ms to around 0.09 ms, which is often the result of failing to optimally exploit sparse matrix operations by the execution platform.

The DNN (Deep Neural Network) appears as the most "compression-friendliest" topology among the various deep network topologies in this study. Diverging from the recurrent counterparts, the DNN maintains high confidence levels involving each optimization paradigm applied. Of particular note, the values of both pruned and clustered versions of the DNN exhibit the smallest latency values shown in the whole table, down to a whopping 0.00475 ms. Like this CNN has the benefit of pruning as well but the serious decline of confidence when clustering (with the scores collapsing into the 0.35–0.41 range). This observation suggests that although CNNs can be significantly accelerated, clustering convolutional kernels in CNNs may remove spatial features that are needed for confident intrusion detection. A recurring theme throughout all experimental conditions is the difficulty of proper classification of the probe attacks. Even in baseline models, the confidence involved in probing is significantly worse than in DDoS detection. DNN using QAT the model produces agile and fine-grained labels like [*http flood*] and [*apacherange_dos*]. This pattern suggests that quantification could be an attempt to regularize the learner or that it is QAT itself that allows for the capture of more granular attack subclasses. On the other hand, most variants that are not heavy tend to resort to a "DDoS" classification in the absence of knowledge, which is a sign of loss of feature depth.

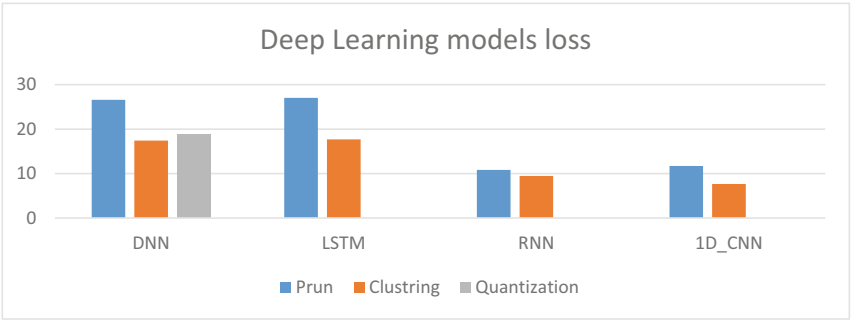

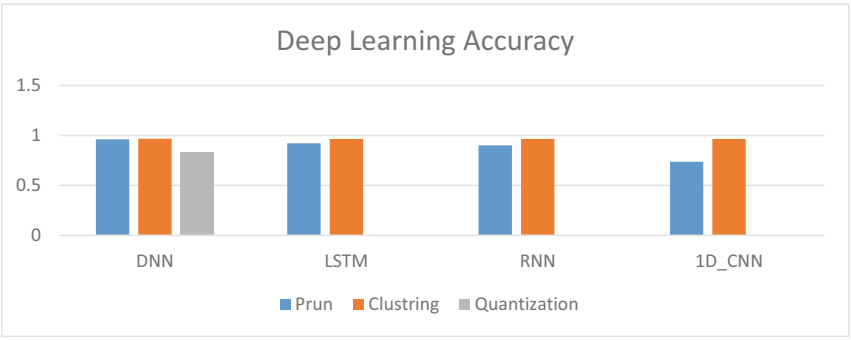

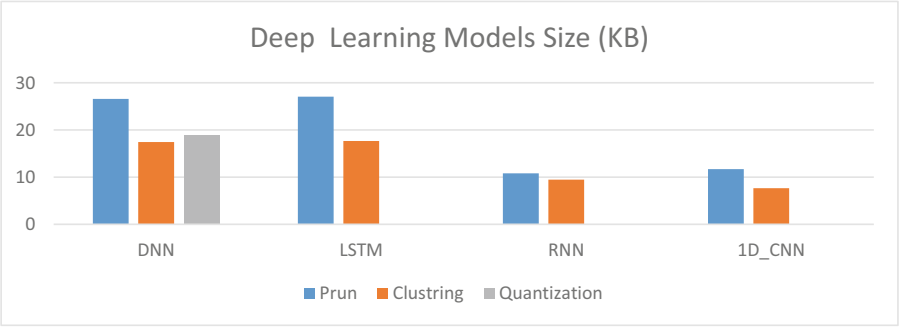

Fig. 5. Deep learning model test results. a) Loss, b) Accuracy, c) Model size

Table 4. Classification testing of real-time traffic on each dataset for 300 packets

Attack Category	Classified Label	Confidence	Time (ms)	Classified Label	Confidence	Time (ms)	Classified Label	Confidence	Time (ms)	Classified Label	Confidence	Time (ms)
Baseline model				LSTM Prun			LSTM **Cluster**			LSTM QAT		
BFA	BFA	0.8219	0.06	DDoS	0.798335	0.0925	DDoS	0.719197	0.097	-	-	-
DDoS	DDoS	0.9656	0.06	DDoS	0.78747	0.09325	DDoS	0.86546	0.10365	-	-	-
Probe	Probe	0.2956	0.26	Probe	0.41147	0.087018	DDoS	0.491566	0.093158	-	-	-
Baseline model				**RNN Prun**			**RNN Cluster**			**RNN QAT**		
BFA	BFA	0.7303	0.06	DDoS	0.626135	0.0705	DDoS	0.905968	0.1505	-	-	-
DDoS	DDoS	0.9669	0.06	DDoS	0.707992	0.07605	DDoS	0.905091	0.11185	-	-	-
Probe	Probe	0.3414	0.25	Probe	0.546567	0.072281	DDoS	0.623057	0.107193	-	-	-
Baseline model				**DNN Prun**			**DNN Cluster**			**DNN QAT**		
BFA	BFA	0.8212	0.06	DDoS	0.850372	0.005	BFA	0.849724	0.0055	[http_flood]	0.715332	0.00775
DDoS	DDoS	0.9802	0.06	DDoS	0.807679	0.0071	DDoS	0.821406	0.00475	[http_flood]	0.796465	0.00845
Probe	udp_reflection	0.7002	0.31	DDoS	0.528366	0.005439	Probe	0.703944	0.004912	[apache_range_dos]	0.37685	0.007895
Baseline model				**CNN Prune**			**CNN Cluster**			**Baseline model**		
BFA	BFA	0.8432	0.06	DDoS	0.380788	0.0115	BFA	0.750968	0.0105	-	-	-
DDoS	DDoS	0.9545	0.07	DDoS	0.379386	0.0148	DDoS	0.774739	0.0104	-	-	-
Probe	BFA	0.7728	0.24	DDoS	0.351902	0.010877	DDoS	0.416424	0.01	-	-	-

In sum, if the operating goal is maximum inference speed for edge deployed intrusion detection applications, the pruned DNN is the best option and it offers a speedup of 10x with virtually no tradeoff in confidence. On the other hand, when the most important question is the exact detection of complex attack patterns, then the baseline LSTM is still better because the interventions to make it a lightweight model (specifically clustering for consistency) adversely affect its predictive acuity. The evidence gathered here suggests that, for security critical deployments, pruning is the better optimization strategy and will always maintain the integrity of the original classification schema.

6 Conclusion

The present study fulfilled the dire need of systematic research on efficient optimization algorithms for IoT security. We used a realistic SDN test bed based on Mininet and Ryu controllers to demonstrate that standard IDS methods have to be reconfigured to manage the special vulnerabilities of the SDN architectures. Our relative assessment of light weighting measures showed that there was an apparent tradeoff between the efficiency of computation and fidelity to classification. Pruning proved to be the most stable procedure, it maintained the original schema of classification but greatly diminished the size of the models. DNN architectures were very good at compression with latencies as low as 0.00475 milliseconds. Recurrent models (LSTM and RNN) were found to be highly accurate; nevertheless, they were vulnerable to clustering and quantization components, resulting in such phenomena as label drift or training collapses due to unstable gradients. In applications where maximum inference speed is the prime consideration, pruned DNN is the best solution. On the other hand, in cases where there is a high need to detect complex patterns of attacks with a high level of accuracy, due to the involvement of protection of sensitive data and government secrets, the normal LSTM would provide better performance, regardless of the higher computational requirements. The focus of future research should be on optimization of sparse matrix operations in order to reduce the gap between the complexity of the model and performance of the edge- devices in real-time.

References

1. Guo, L., Li, X., Si, W.: Research on deep learning-based optimization algorithms for IoT communication. IEEE Access. **13**, 177724–177724 (2025). https://doi.org/10.1109/ACCESS.2025.3597098
2. Di Gennaro, F., Cucchiarelli, A., Morbidoni, C., Spalazzi, L.: A comparative analysis of datasets for intrusion detection in software-defined networks. In: Joint National Conference on Cybersecurity (ITASEC & SERICS 2025) (2025)
3. Khorseed, W.S., Hamad, A.H.: Inter and intra domain DDOS attack mitigation for software defined network based on hyperledger fabric blockchain technology. Ingénierie Des Systèmes D Inf. **29**(1), 301–311 (2024). https://doi.org/10.18280/isi.290130
4. Janabi, A. H., Kanakis, T., Mark Johnson: Survey: intrusion detection system in software-defined networking. In Barbara Guidi (Ed.), IEEE Access **12**, 164097 (2024). https://doi.org/10.1109/ACCESS.2024.3493384

5. Idrissi, I., Moussaoui, O., Azizi, M.: A lightweight optimized deep learning-based host-intrusion detection system deployed on the edge for IoT. Int. J. Comput. Digit. Syst. **11**(1), 209–216 (2022). https://doi.org/10.12785/ijcds/110117
6. Luo, K.: A distributed SDN-based intrusion detection system for IoT using optimized forests. PLoS One. **18**(8), e0290694 (2023). https://doi.org/10.1371/journal.pone.0290694
7. Sharmila, B.S., Nagapadma, R.: Quantized autoencoder (QAE) intrusion detection system for anomaly detection in resource-constrained IoT devices using RT-IoT2022 dataset. Cybersecurity. **6**(1), 41 (2023). https://doi.org/10.1186/s42400-023-00178-5
8. Wang, Z., Chen, H., Yang, S., Luo, X., Li, D., Wang, J.: A lightweight intrusion detection method for IoT based on deep learning and dynamic quantization. PeerJ Comput. Sci. **9**, e1569 (2023). https://doi.org/10.7717/peerj-cs.1569
9. Yu, R., Le Jeune, L., Mentens, N.: University of Florida, KU Leuven, & Worcester Polytechnic Institute. Quantization-aware neural architectural search for intrusion detection. arXiv. https://arxiv.org/pdf/2311.04194 (2024)
10. Huynh, T.T., Nguyen, H.T.: Effective multi-stage training model for edge computing devices in intrusion detection. Int. J. Comput. Netw. Commun. **16**(1), 13–31 (2024). https://doi.org/10.5121/ijcnc.2024.16102
11. Bouayad, A., Alami, H., Janati Idrissi, M., Berrada, I.: Lightweight federated learning for efficient network intrusion detection. In S. D. Campbell (Ed.), Access **12**, 172027 (2024). https://doi.org/10.1109/ACCESS.2024.3494057
12. More, S.A., Kachavimath, A.V.: SDN intrusion detection using meta-heuristic optimization and K-nearest Neighbors classifier. Procedia Comput. Sci. **260**, 1137–1144 (2025). https://doi.org/10.1016/j.procs.2025.03.299
13. Rangsikunpum, A., Amiri, S., Ost, L.: Ultra-lightweight and highly efficient pruned binarised neural networks for intrusion detection in in-vehicle networks. Electronics. **14**(9), 1710 (2025). https://doi.org/10.3390/electronics14091710
14. Dhirar, H., Hamad, A.: Internet of things software-defined network intrusion detection dataset: inter- and intra-domain. Al-Khwarizmi Eng. J. **21**(3), 35–47 (2025). https://doi.org/10.22153/kej.2025.08.001
15. Dhirar, H., Hamad, A.H.: Federated deep learning intrusion detection system on software defined-network based internet of things. IAES Int. J. Artif. Intell. (IJ-AI). **14**(4), 3109–3120 (2025) http://ijai.iaescore.com
16. Ferreira, G., Da Nóbrega Marinho, M.H., Severo, V., Madeiro, F.: Optimization strategies applied to deep learning models for image steganalysis: application of pruning, quantization and weight clustering. Appl. Sci. **15**(9), 4632 (2025). https://doi.org/10.3390/app15094632
17. Han, S., , Mao, H., , Dally, W. J.: Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In ICLR 2016. https://arxiv.org/pdf/1510.00149 (2016)
18. Jacob, B., et al.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2704–2713, Salt Lake City, UT, USA (2018). https://10.1109/CVPR.2018.00286
19. Zhang, M., Ma, Y., Li, J., Zhang, J.: A density connection weight-based clustering approach for dataset with density-sparse region. Expert Syst. Appl. **230**, 120633 (2023). https://doi.org/10.1016/j.eswa.2023.120633

Author Index

A
Aktas, Mehmet S. 242
Aleixo, Sandra M. 224
Alves, Maria João 73
Alves, Marta 181
Alves, Rui 344
Ambrosino, Carmine 575
Amorim, Ana Paula 192
Angrisano, Mariarosaria 506, 541, 560
Antelis, Javier M. 624
Antonia, Gravagnuolo 541

B
Bernine, Elaine 55
Bitai, Hu 3, 22
Bogdanov, Aleksandr V. 3, 22
Bol, Konuralp 242
Boquete, Yolanda Pena 387
Borges, Fábio 99
Bottero, Marta 541
Braga, Ana C. 119, 308
Buglione, Francesca 506

C
Cação, Isabel 84
Caeiro, Frederico 260
Caiazzo, Mariano 592
Cantoral-Ceballos, José Antonio 624
Carpentieri, Gerardo 387
Castiglione, Arcangelo 607
Castro, Cecilia 192
Catabriga, Lucia 55
Cavana, Giulio 541
Çayır, Sefa 242
Cerati, Davide 455
Cilgin, Onur 242
Cunha, Jorge 308

D
D'Angelo, Gianni 575, 592, 607
de Oliveira, Pedro José Costa 163
Della Morte, Piera 488
Denissova, Oxana 402
Diagilev, Mikhail 39
Dias, Cátia 144
Dik, Aleksandr 133
Dimas, Isabel Dórdio 181
Dinu, Alexandru 358

E
Esposito, Raffaele 592
Ettari, Simona 575

F
Falcão, M. Irene 84
Feoktistova, Yelena 419
Fernandes, Célia 224
Fernandez-Grela, Manuel 387
Ferreira, Artur 144
Ferreira, Thays Rocha Neri 99
Frunzete, Mădălin 358

G
Girard, Luigi Fusco 541
Gravagnuolo, Antonia 471, 506
Guida, Carmen 387
Gurskaya, Oxana 419
Guzelcay, Deniz 242

H
Hamad, Ali H. 641

I
Inozemtseva, Tatiana 419

J
Jakubócová, Lenka 437
Jovanovic, Dina 455

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 657–658, 2027.
https://doi.org/10.1007/978-3-032-30518-3

K
Khalil, Adnan Elahi Khan 624
Khvatov, Valery 3, 22
Kiyamov, Jasur 133
Konurbayeva, Zhadyra 402
Korkhov, Vladimir 39

L
Lee, Chien-Sing 289
Leitão, Ana Maria 73
Lima, José 210
Llerena, Yanelys 327
Lopes, Joana M. A. 119
Lopes, Mariana 144
Lourenço, Paulo Renato 181

M
M. Aleixo, Sandra 144
Magalhães, Alexandre 192
Malonek, Helmuth R. 84
Mateus, Ayana 260
Matos, Leonardo Nogueira 163
Mecca, Ippolita 523
Meira-Machado, Luís 344
Meireles, Paula 344
Mendes, João 210, 272
Milho, Rute 224
Mitaľ, Ondrej 437
Modan, Sara-Gabriela 358
Moreira, Carla 344
Mulinacci, Martina 455

N
Nanda, Sharini Sivadas 289
Neglia, Grazia 523
Neves, M. Manuela 372
Novais, Maria João 344
Novais, Paulo 163

O
Oreni, Daniela 455

P
Palmieri, Francesco 575, 592, 607
Pereira, Ana I. 210, 272
Perez-Díaz, Jesus Arturo 624
Petrosino, Alfonso 575
Pinto, Filipa 344
Prata Gomes, Dora 372
Prazeres, Inês 144

R
Rakhimberdinova, Madina 402
Ramires, Margarida 224
Ramos, Paulo 224
Rebelo, Teresa 181
Relvas, Marta M. M. 119
Riganti, Patrizia 471
Rocha, Miguel 344
Ručinská, Silvia 437

S
Sá, Tiago 327
Salazar-Torres, Jackiline 387
Saleh, Mustafa T. 641
Santos, Isaac P. 55
Santos, Maribel Y. 327
Santos, Martim 327
Savkov, Egor 3, 22, 133
Scarfó, Antonio 607
Scioti, Albina 560
Sena, Inês 210, 272
Severino, Ricardo 73
Sidumo, Aurélio 344
Silva, Bruno 327
Silveira, Ana 144
Simsek, Halil İbrahim 242
Siqueira, Lauryane Santos 163
Soares, Ana Jacinta 192
Sousa, Ana 308
Sousa, Vânia 327
Surova, Darya 402

T
Tartari, Maria 541
Tomaz, Graça 84

V
Verardi, Ferdinando 523
Vieira, António C. 327

Y
Yussubaliyeva, Madina 402, 419

Z
Zakimova, Allfiya 402