The Intersection of Finite Element Limit Analysis and Physics-Informed Machine Learning for Structural Vulnerability Assessment

Nima Dolatabadi[1], Renato Zona[2(✉)], and Vincenzo Minutolo[2]

[1] Department of Civil-Industrial Engineering and Architecture, University of Catania, Via Santa Sofia 64, 95123 Catania, Italy
nima.dolatabadi@phd.unict.it

[2] Department of Engineering, Università della Campania Luigi Vanvitelli, Via Roma 29, 81031 Aversa, CE, Italy
{renato.zona,Vincenzo.MINUTOLO}@unicampania.it

Abstract. Recent advances in computational mechanics and artificial intelligence have expanded the range of tools available for structural engineering. While conventional elastoplastic iterative algorithms can be computationally demanding, limit analysis offers an efficient way to estimate ultimate load capacity. This paper investigates the use of 3D FELA with a physics-informed machine learning workflow to predict structural vulnerability. The learning component is a multi-output multilayer perceptron surrogate, while the physics-informed aspect is introduced through the FELA-based generation of training data and targets derived from the static theorem of plasticity. The methodology is validated across five architectural morpho-types: Rectangular, L-type, C-type, Asymmetrical C-type, and H-atrium buildings. Within the investigated dataset, the proposed scheme removes repeated iterative analyses in parametric studies, achieves correlation coefficients above $R^2 = 0.995$, and reduces average analysis time by more than a factor of 1,200. These results indicate that the proposed PIML workflow can support rapid structural assessment and urban-scale seismic micro-zoning.

Keywords: Limit Analysis · Finite Element Method · Physics-Informed Machine Learning · Structural Vulnerability

1 Introduction

Artificial intelligence and machine learning are increasingly used in structural engineering for design support, response prediction, and data-driven assessment (Málaga-Chuquitaype, 2022). These methods complement traditional statistical tools, especially when relationships between variables are strongly non-linear or difficult to model explicitly (Suryawanshi et al., 2025). In parallel, direct limit

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16762, pp. 619–635, 2027.
https://doi.org/10.1007/978-3-032-30524-4_38

analysis has become an efficient alternative to incremental elastoplastic procedures for estimating structural capacity through mathematical programming (Spiliopoulos and Dais, 2012). To reduce the limitations of purely data-driven "black-box" models, recent studies have introduced physics-informed machine learning approaches that couple learning algorithms with mechanics-based constraints, simulations, or data-generation procedures (Imani et al., 2025). In the present study, the machine-learning model itself remains relatively standard, while the physics-informed character of the workflow is provided primarily by the FELA-based generation of training data. In the context of seismic vulnerability assessment, this combination is relevant because it can support rapid evaluation of large building stocks (Zona et al., 2025).

1.1 Evolution of Computational Mechanics

Over the last three decades, computational mechanics and artificial intelligence have substantially influenced structural engineering practice (Chen, 2021). Advances in high-performance computing have enabled sophisticated numerical models for the analysis of complex three-dimensional systems under severe loading conditions (Ferrari, 2017). However, detailed building-by-building modeling remains difficult to deploy quickly at the urban scale, particularly for seismic vulnerability assessment (Imani et al., 2025). Early computational approaches relied on step-by-step elastoplastic analyses, which required substantial time and computational resources. Limit analysis emerged as a more efficient alternative for estimating ultimate load-bearing capacity directly (Zona and Minutolo, 2025). Based on the theorems of plasticity, it avoids incremental elastic-plastic calculations and focuses instead on the collapse state of the structure (Bleyer,2013). The integration of the finite element method with limit analysis through linear and non-linear mathematical programming has therefore produced effective tools for structural assessment. (Makrodimopoulos, 2008). At the same time, machine learning has become a useful tool for identifying complex non-linear patterns in large structural datasets (Bishop, 2006). Deep neural networks and related architectures are increasingly used to predict structural response, detect damage, and support recovery planning after extreme events (Karniadakis et al., 2021; Sun et al., 2021). However, purely data-driven models may be difficult to interpret physically and can perform poorly when extrapolating beyond the domain represented in the training data. (Senders et al., 2018; Di Bonito et al., 2024). These concerns have motivated the development of explainable and physics-informed approaches, especially for safety-critical applications (Zheng et al., 2025; Gondaliya et al., 2025). In particular, embedding lower-bound finite element procedures into the learning framework helps preserve consistency with the governing mechanics. (Murashko and Tkachov,2025). For emergency planning, generating high-resolution numerical models for every building is often computationally prohibitive (Ferranti et al., 2024). Large-scale vulnerability studies therefore require methods that are both simplified and reliable. Traditional workflows also depend on extensive manual data collection, which limits their applicability in rapid assessments (Zona et al., 2025). Machine learning can help classify

buildings according to expected vulnerability by exploiting patterns in empirical and simulated datasets (Ahmad et al., 2026; Seo et al., 2012). Within this broader context, structural capacity information remains essential for risk management and loss reduction. Performance-based urban planning further benefits from grouping existing structures into representative morpho-types defined by geometric, mechanical, and historical attributes (Spence et al., 2005; Alipour et al., 2011). Such classifications can support seismic micro-zoning maps that identify urban sectors with similar structural vulnerabilities and can guide long-term regeneration strategies (Lagomarsino and Podestà, Cardellicchio et al., 2023; 2006).

2 Method and Materials

The procedure adopted in this research combines a mechanics-based structural analysis framework with a data-driven machine learning process. A core computational system performs elastic analysis and plastic limit analysis of three-dimensional frames to estimate the global collapse factor and the local collapse modes. To avoid the cost of repeated analyses over a large parameter space, a regression-based neural network is trained to predict the target quantities from a normalized dataset of geometric and physical input parameters.

2.1 Structural Analysis Core and the Static Theorem Framework

Instead of finite element analyses with elemental stress limits against yielding, this analysis package determines the state of the whole structure failure using the static theorem of limit analysis. The whole structure domain is conceptualized as a system of spatial joints where the elements are the structure members idealized as 3D Euler-Bernoulli beams. These members are connected at their ends providing to each joint 12 degrees of freedom, considering their full 3D kinematics along with the 6 degrees of freedom at both the start and the end node of each beam element. The internal forces included in the analysis comprise axial force, weak and strong axis bending moments and shears, and torque. These structures are constrained to their respective spatial domains by setting to zero selected nodal degrees of freedom such as restraints for translations and rotations of nodal points that represent rigid foundations and supports. Finding the maximum load capacity is expressed as an extremely large linear programming optimization problem, for which the primal optimization problem aims at maximizing globally collapsing load multiplier λ, which refers to the safety factor regarding total plastic collapse subject to equality constraints of global equilibrium, which is stated as:

$$\mathbf{Hx} = \lambda \mathbf{F}_{var} + \mathbf{F}_{perm} \tag{1}$$

In which $\mathbf{H}$ is the overall equilibrium matrix, $\mathbf{x}$ is the overall vector of internal forces, $\mathbf{F}var$ is the vector of variable loads for which the external actions are scaled to provide plastic collapse and $\mathbf{F}perm$ is the vector of loads for permanent

load (due to the gravity effect on dead load). Superposition is handled automatically by the solver. It calculates an elastic response to permanent load, and adds this to the variable load case using a load combination factor before trying to solve for the plastic collapse. The overall stiffness and equilibrium matrices are constructed in an efficiently vectorized form and do not involve an iterative algorithmic loop for computational reasons.

2.2 Piecewise Linear Yield Surface and Interaction Domains

In order to guarantee that internal forces vector x does not violate the structural framework capacity. The inequality constraint are set strictly on the boundaries of the optimization problem. The structure yield point formation or plastic hinge is mathematically formulated by eight-faced linear piecewise yield surface. This domain clearly includes the cross-section interaction of axial force with multi-axial moments bending. The physical constraints of orthogonal frontiers for the individual linear facet are based on the geometric cross-section properties as well as material properties given to the member components. For the building morphotypes analyzed here, the constraints that form boundaries of the optimization have been selected as a symmetrical $275.0MPa$ as max compression and max axial tension limits. In addition, algorithm succeeded to sum-in the dead loads by a combination factor of 0.3 prior to the determination of the lateral plastic collapse states, the plasticity yield constraint in the LP bounds is defined as follows:

$$\mathbf{A}_{yield}\mathbf{x} \leq \mathbf{b}_{capacity} \tag{2}$$

Here $\mathbf{A}yield$ is a matrix that transforms the directed normal vectors from the eight faceted polyhedron and $\mathbf{b}capacity$ contains the axial and moments limits. This mathematical surface is then used in the post-processing where it becomes possible to establish exactly which facet of the surface has been violated and where in the mechanism collapse occurs, in addition with each distinct yield mode separated based upon the interaction that controls the failure.

2.3 Data Preprocessing and Matrix Normalization

The structural data fed into the neural network is derived from six distinct input files that define the three-dimensional frame. A summary of these required input files, their mathematical formats, and structural engineering significance is presented in Table 1 and detailed below.

Table 1. Summary of required input files, their mathematical formats, and structural engineering significance.

File Type	Format	Description/Engineering Concept
Node Coordinates	Matrix $[N_{nodes} \times 3]$	Defines the global X, Y, and Z coordinates for all structural joints in the 3D wireframe.
Element Connectivity	Matrix $[N_{elements} \times 4]$	Defines the structural topology by linking start/end nodes, and assigning section and load IDs.
Boundary Conditions	Vector $[(N_{nodes} \times 6) \times 1]$	Locks specific nodes to the foundation by defining 6 degrees of freedom as free (1) or fixed (0).
Nodal Loads	Vector $[(N_{nodes} \times 6 \times N_{cases}) \times 1]$	Contains the vertically stacked external forces, separated into permanent dead loads and variable loads.
Beam Orientation	Matrix $[N_{elements} \times 3]$	Provides the perpendicular normal vector to determine the strong and weak bending axes.
Section Properties	Matrix $[N_{sections} \times 11]$	Defines material stiffness and physical strength limits dictating when and where plastic hinges form.

- **Node Coordinates:** This file translates the physical structure into a numerical wireframe. It records the exact locations in 3D space where columns and beams intersect or touch the foundation.
- **Element Connectivity:** This file establishes the structural members. It maps out the connections between specific starting and ending nodes, while also assigning the designated cross-sectional shapes and external load identifiers to each physical beam or column.
- **Boundary Conditions:** This file serves as the restraints for the computational model, preventing the structure from undergoing infinite rigid body translation or rotation under applied forces. It establishes constraints for all 6 spatial degrees of freedom at every node, identifying which joints are free to move and which are glued to an immovable foundation.
- **Nodal Loads:** This file defines the external mechanical actions exerted on the building's joints. The data is stacked to apply permanent loads, which represent the constant gravity effect of the structure's own weight, followed by the variable loads representing wind or seismic forces that are scaled to determine the safety factor.
- **Beam Orientation:** This file defines the spatial rotation of each element. By providing a vector perpendicular to the main axis of the beam, it ensures the solver correctly identifies the strong axis and weak axis for bending, which fundamentally controls the element's structural stiffness.
- **Section Properties:** This file assigns the physical and material thresholds for the cross-sections used in the elements. It provides the elastic stiffness parameters alongside the critical plastic limits for axial compression, axial tension, and multiaxial bending moments. When the internal forces reach these designated strength limits, the elements yield and form rotating plastic hinges instead of snapping.

While the exact mechanisms of collapse may be completely identified by means of direct linear programming solvers, a full exploration of the parametric design space requires too much computational expense. Therefore, a deep learning regression model is developed to estimate, rapidly, the six values that define the lateral variable loads along both principle planar directions and the elastic

and plastic multipliers associated with the fixed loads. Training and validation sets for both the features and the target matrix are separated using a ratio of eighty to twenty. To prevent large-scaled input features with high numerical values from unduly dominating the gradient of optimization, statistical scaling is applied to the features. Calculation of mean μ_j and variance σ_j and normalization of the input features occurs via computation of the empirical mean and standard deviation of the features.

$$\mu_j = \frac{1}{N}\sum_{i=1}^{N} X_{i,j} \tag{3}$$

$$\sigma_j = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(X_{i,j} - \mu_j)^2} \tag{4}$$

Every input tensor is subsequently normalized via a standard transformation:

$$X_{i,j}^{(scaled)} = \frac{X_{i,j} - \mu_j}{\sigma_j} \tag{5}$$

These scaling coefficients are stored and serialized for use by deployment applications to ensure the mathematical consistency of the models when they are inferred in the future.

The predictive surrogate model is in the form of a fully connected multilayer perceptron intended to capture the extremely non-linear relations that exists between frame geometry and its plastic capacity. A completely connected fully connected multilayer perceptron network comprised of an input layer matching the shape of the feature set, three contracting fully connected dense hidden layers, each containing 64, 32, and 16 neurons, is used in the prediction:

$$f(z) = \max(0, z) \tag{6}$$

The structural data is then fed through the parameters in the network in a fixed sequence of matrix transformations, as outlined below:

First Hidden Layer:

$$\mathbf{H}_1 = \max(0, \mathbf{X}_{scaled}\mathbf{W}_1 + \mathbf{b}_1) \tag{7}$$

Second Hidden Layer:

$$\mathbf{H}_2 = \max(0, \mathbf{H}_1\mathbf{W}_2 + \mathbf{b}_2) \tag{8}$$

Third Hidden Layer:

$$\mathbf{H}_3 = \max(0, \mathbf{H}_2\mathbf{W}_3 + \mathbf{b}_3) \tag{9}$$

The output layer: since the structural load multipliers are continuous real numbers over the entire elastic and collapse behavior regions of operation, the output layer has six neurons giving the linear combinations directly without any limitations on activation:

$$\hat{\mathbf{Y}} = \mathbf{H}_3\mathbf{W}_4 + \mathbf{b}_4 \tag{10}$$

In which the matrix $\hat{\mathbf{Y}}$ readily gives the prediction vectors for the structural response, respectively.

The parameters that can be learned are the internal matrices of weight and bias for this neural network, which are optimized through minimization of a specific empirical risk. Specifically, the objective function which is minimized is the mean squared error that severely punishes major discrepancies between the predictive output and the true structural capacities. To reflect the specific physical parameters being predicted—the elastic (λ_e) and collapse (λ_c) load multipliers under permanent, lateral X, and lateral Y conditions—the loss function is explicitly defined as:

$$\mathcal{L}_{MSE} = \frac{1}{N_{batch}} \sum_{i=1}^{N_{batch}} \sum_{j\in\{perm,sx,sy\}} \left[\left(\lambda_{e,j}^{(i)} - \hat{\lambda}_{e,j}^{(i)}\right)^2 + \left(\lambda_{c,j}^{(i)} - \hat{\lambda}_{c,j}^{(i)}\right)^2 \right] \tag{11}$$

Meanwhile, the model also calculates the mean absolute error as its main tracking criterion. This value provides a direct interpretation of the prediction deviation on the true physical scale of the structural load multipliers:

$$\mathcal{M}_{MAE} = \frac{1}{N_{batch}} \sum_{i=1}^{N_{batch}} \sum_{j\in\{perm,sx,sy\}} \left[\left|\lambda_{e,j}^{(i)} - \hat{\lambda}_{e,j}^{(i)}\right| + \left|\lambda_{c,j}^{(i)} - \hat{\lambda}_{c,j}^{(i)}\right| \right] \tag{12}$$

The optimizer is driven by an adaptive moment estimation algorithm with initial learning rate at 0.001. The training is carried out for a defined maximum number of steps, during which training and validation loss are both recorded to verify smooth mathematical convergence.

Three concurrent protocols are used to constrain the gradient descent loop in such a way that the neural network generalizes well and correctly maps un-seen structures rather than memorizing.

Learning rate scheduling: The convergence of the gradient descent path to the minimum value is ensured using dynamic learning rate. Validation loss is observed, and whenever the minimum has been reached and stagnant for certain steps, the learning rate is divided by two geometrically. The process continues until an absolute minimum learning rate is reached to control the optimization weights accurately.

$$\alpha_{new} = \begin{cases} 0.5 \times \alpha_{old} & \text{if } \Delta\mathcal{L}_{val} \approx 0 \text{ for threshold epochs} \\ \alpha_{old} & \text{otherwise} \end{cases} \tag{13}$$

Early stopping: We prevent over-fitting of the model-where the algorithm still decreases training error but not validation error-by early stopping the optimization. This mechanism is implemented based on monitoring of the generalization gap - the discrepancy between the empirical risk on the training data $Jtrain(\theta)$

and risk on an independent data set, $Jval(\theta)$. Early stopping depends on patience parameter τ - a discrete interval of epochs. Stopping criteria are reached at epoch t when the validation error is not decreased during τ consecutive iterations:

$$\min\{J_{val}(\theta_{t-\tau}), \ldots, J_{val}(\theta_{t-1})\} \leq J_{val}(\theta_t) \tag{14}$$

In a more mathematical formulation, suppose that θ^* are the parameters that obtain the minimum on the whole past of training up to time t.

$$\theta^* = \arg \min_{\theta \in \{\theta_1, \ldots, \theta_t\}} J_{val}(\theta) \tag{15}$$

When the termination condition is met, the parameters matrices of that specific iteration is manually reset to the one that returned the global minimal error in validation and thus eliminates the effect of further parameter updates on the degraded generalization. Therefore, by resetting the network to its most predictive states for structural limit analysis, one ensures that the returned θ^* parameters is that which is closest to the physical relation between structural geometry and collapse multiplier:

Optimal State Preservation: This protocol complements the training loop. The states are examined and archived at the end of each training step. Archiving events to serialize the final mathematical graph is initiated only when the current validation error surpasses a new absolute minimum in the validation loss terrain (Fig. 1).

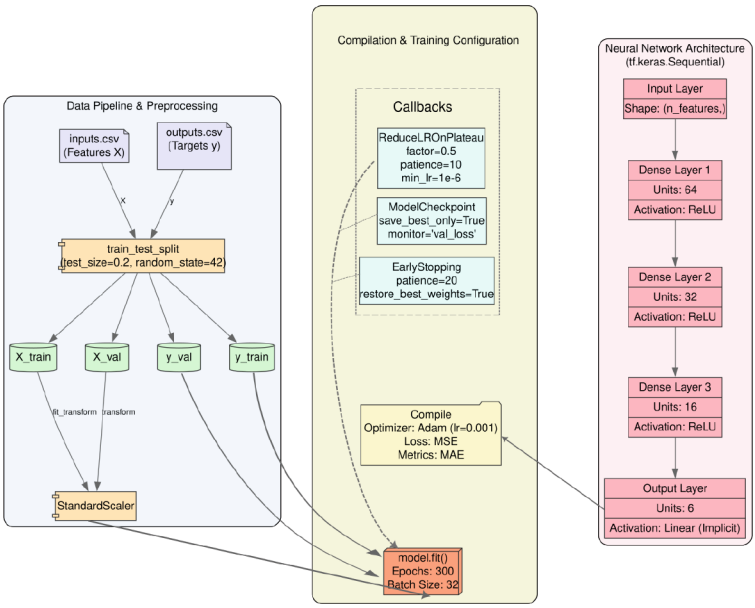

Fig. 1. Schematic representation of the Neural Network architecture, illustrating the integration of limit analysis constraints within the loss function for structural vulnerability assessment.

2.4 Surrogate Model Deployment and Inference

This deployment stage provides a dramatic departure from the economics of computationally evaluating response for intricate 3D frames. The linear programming simplex algorithm would otherwise have to run with increasingly indeterminate stiffness matrices and solve for many structures with each new geometry, but the neural network performs the forward pass approximation instantly. In essence, we have transformed the problem from constraint-based optimization to a deterministic functional map, which allows for "real-time" structure evaluations. This inference pipeline takes as an input an unobserved vector $\mathbf{x}_{new} \in \mathbb{R}^n$ which corresponds to a structure with a particular geometry and material response. It is important to provide the input vector as a transformation of an input in a space compatible with the model-meaning the input vector is subjected to an identical affine transformation using the distribution's parameters collected during the training phase:

$$\mathbf{x}_{scaled} = (\mathbf{x}_{new} - \boldsymbol{\mu}) \div \boldsymbol{\sigma} \tag{16}$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the vector of means and standard deviation for each feature respectively. The main part of the deployment is the application of the trained neural operator F_θ. The predicted value $\hat{\mathbf{y}}$ is computed via a composite sequence of nonlinear functions:

$$\hat{\mathbf{y}} = \mathbf{W}_L \cdot \sigma_{L-1}(\mathbf{W}_{L-1} \ldots \sigma_1(\mathbf{W}_1 \mathbf{x}_{scaled} + \mathbf{b}_1) \cdots + \mathbf{b}_{L-1}) + \mathbf{b}_L \tag{17}$$

Given the specific architecture of the surrogate model, where σ represents the Rectified Linear Unit (ReLU) activation function, the internal state of each hidden layer l is computed as:

$$\mathbf{h}_l = \max(0, \mathbf{W}_l \mathbf{h}_{l-1} + \mathbf{b}_l) \tag{18}$$

The inference here is computationally less complex from the $\mathcal{O}(n^3)$ or $\mathcal{O}(2^n)$ needed in iterative limit analysis solvers, and can be expressed as a series of matrix-vector multiplications that require complexity $\mathcal{O}(\sum l = 1^L nlnl - 1)$, where nl is the number of neurons in layer l. The recovered neural structure evaluates the entire array of elastic and plastic load multipliers at once, and the resulting vector $\hat{\mathbf{y}}$ is transformed back into physical space:

$$\hat{\mathbf{y}} = [\lambda_{e,perm}, \lambda_{e,sx}, \lambda_{e,sy}, \lambda_{c,perm}, \lambda_{c,sx}, \lambda_{c,sy}]^T \tag{19}$$

This set of analytic data is organized and extracted as needed. Structural response and failure locations can be immediately predicted, using only the initial geometry and material information, and these can be rapidly communicated for immediate use in the fight against natural hazards.

Multi-output Regression Architecture. The surrogate model is a feedforward Multi-Layer Perceptron (MLP) that has been trained to predict multiple structural responses at the same time. Instead of determining the damage levels separately, the network has been trained as a multi-output regression model to directly approximate the six key response quantities that need to be predicted: the elastic multipliers associated with dead loads, X- direction earthquake loads, Y- direction earthquake loads and their associated plastic collapse multipliers. The mapping can be expressed mathematically as follows:

$$\hat{\mathbf{y}} = f(\mathbf{X}; \boldsymbol{\theta}) \tag{20}$$

In which $\mathbf{X}$ is the normalized matrix of structural input features, $\hat{\mathbf{y}} \in \mathbb{R}^6$ is the target vector predicted by the network, and $\boldsymbol{\theta}$ is composed of the network's weights and biases.

3 Results

This section presents the quantitative validation of the proposed framework using high-fidelity ground truth generated by the 3D Frame Limit Analysis Suite. The results are organized into four parts: structural response, surrogate-model convergence, predictive accuracy, and computational efficiency.

3.1 Structural Response and Collapse Mechanisms

Before training the machine learning model, the physical consistency of the ground-truth dataset was checked by simulating the response of several three-dimensional structural topologies under different load conditions. Although the framework can be applied to broader classes of buildings, this study focuses on five morpho-types: Rectangular, L-type, C-type, Asymmetrical C-type, and H-atrium. Their corresponding collapse mechanisms are shown in Fig. 2.

Table 2 summarizes the elastic and collapse multipliers obtained from the lower-bound formulation. During the solution of the large equilibrium systems, the SVD-based reduction stabilized the constraint matrices. The matrix rank increases with structural complexity, from 231 for the L-type case to 1248 for the H-atrium configuration.

Comparing the C-type and Asymmetrical C-type frames under lateral Load Combinations 2 and 3 shows a marked reduction in collapse multiplier for the asymmetric configuration ($\lambda_{collapse} = 3.9636 \times 10^{-2}$ versus 8.6473×10^{-2} for LC2). This result indicates higher sensitivity to lateral loading in the asymmetric case. The plastic reserve ratios reported in Table 2 also suggest lower redundancy in the Asymmetrical C-type frame, where torsional effects and earlier mechanism formation reduce the post-yield reserve. The idealized capacity curves support the same interpretation. Although limit analysis does not resolve the full incremental degradation path, the initial elastic stiffness $K = \lambda/\Delta_e$ still provides a useful reference for relating elastic response to collapse capacity.

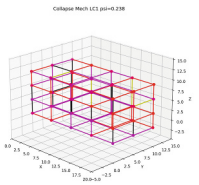

(a) Rectangular (LC1, $\psi = 0.238$)

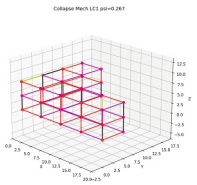

(b) L-type (LC1, $\psi = 0.267$)

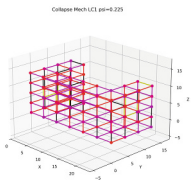

(c) C-type (LC1, $\psi = 0.225$)

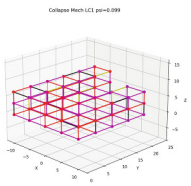

(d) Asym. C-type (LC1, $\psi = 0.099$)

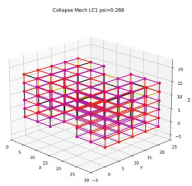

(e) H-atrium (LC1, $\psi = 0.288$)

Fig. 2. 3D frame limit-analysis results for five structural morpho-types under Load Combination 1 (LC1). The figure reports the collapse mechanism and load multiplier (ψ) for each geometry: (a) Rectangular, (b) L-type, (c) C-type, (d) Asymmetrical C-type, and (e) H-atrium.

Table 2. Summary of structural load multipliers and equilibrium matrix constraint reduction across the five distinct building morpho-types.

Morpho-Type	Load Combination	Eq. Matrix Rank	$\lambda_{elastic}$	$\lambda_{collapse}$	$\lambda_{collapse}/\lambda_{elastic}$
Rectangular	LC 2 (Dir X)	360	1.8139×10^{-2}	9.6167×10^{-2}	5.302
	LC 3 (Dir Y)	360	1.5975×10^{-2}	8.3392×10^{-2}	5.220
L-type	LC 2 (Dir X)	231	9.8123×10^{-3}	1.0694×10^{-1}	10.900
	LC 3 (Dir Y)	231	9.7168×10^{-3}	9.8255×10^{-2}	10.110
C-type	LC 2 (Dir X)	625	1.3109×10^{-4}	8.6473×10^{-2}	5.906
	LC 3 (Dir Y)	625	1.2161×10^{-4}	7.9018×10^{-2}	6.105
Asym C-type	LC 2 (Dir X)	435	2.7002×10^{-2}	3.9636×10^{-2}	1.468
	LC 3 (Dir Y)	435	2.5470×10^{-2}	3.7022×10^{-2}	1.454
H-atrium	LC 2 (Dir X)	1248	3.5236×10^{-3}	1.1458×10^{-1}	32.520
	LC 3 (Dir Y)	1248	3.5326×10^{-3}	1.0678×10^{-1}	30.230

The numerical results define the response regime of each structure. For clarity, the tabulated quantities are also presented graphically in Fig. 3, which highlights the relation between topology, load-carrying capacity, and structural redundancy across the considered morpho-types.

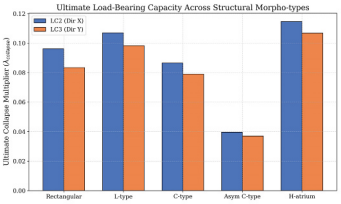

(a) Ultimate Load-Bearing Capacity

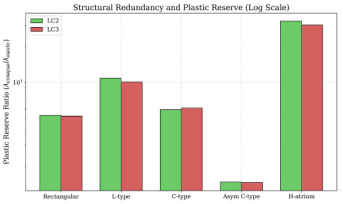

(b) Plastic Reserve Ratio (Log Scale)

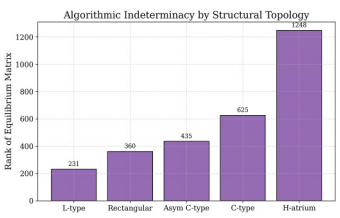

(c) Equilibrium Matrix Rank

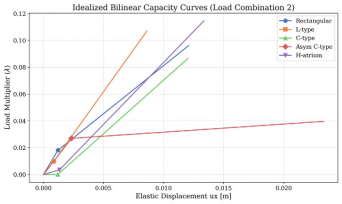

(d) Smooth Idealized Capacity Curves

Fig. 3. Quantitative comparison of the five architectural morpho-types: (a) ultimate load-bearing capacity, (b) plastic reserve ratio, (c) equilibrium matrix rank, and (d) idealized capacity curves.

3.2 Surrogate Model Training and Convergence

The training history of the multi-output multilayer perceptron surrogate is shown in Fig. 4. To capture the non-linear relation between structural geometry and limit capacity while reducing overfitting, regularization was introduced during optimization. The Mean Squared Error ($\mathcal{L}_{MSE}$) is reported on a log scale and shows smooth convergence for both the training and validation sets.

Dynamic learning-rate scheduling helps stabilize optimization by refining the parameter updates as training progresses. Early stopping identifies the onset of overfitting and retains the parameter set associated with the minimum validation loss, which in this case occurs at epoch 202 ($\min \mathcal{L}_{val} = 1.3801 \times 10^{-2}$). The corresponding optimal weights ($\boldsymbol{\theta}^*$) are then used for inference.

The predictive capability of the framework was assessed by comparing surrogate-model outputs with the ground truth generated by finite element limit analysis (FELA). Figure 5 presents parity plots for the six target variables, with FELA values on one axis and neural-network predictions on the other.

The predicted values lie close to the identity line ($y = x$), indicating limited systematic bias across both the elastic and collapse regimes. Quantitatively, the coefficients of determination are above 0.995 for all target variables and above 0.9999 for the elastic multipliers. The corresponding MAEs are on the order of 10^{-4} and 10^{-6} for the load multipliers. Within the analyzed parameter space, these results show that the multi-output multilayer perceptron can reproduce the

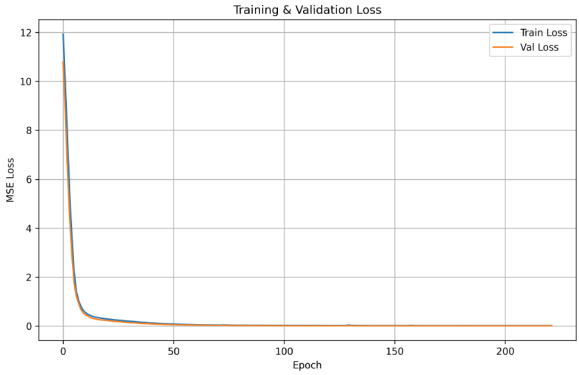

Fig. 4. Training and validation loss curves for the surrogate model. Early stopping selects the best-performing weights at epoch 202.

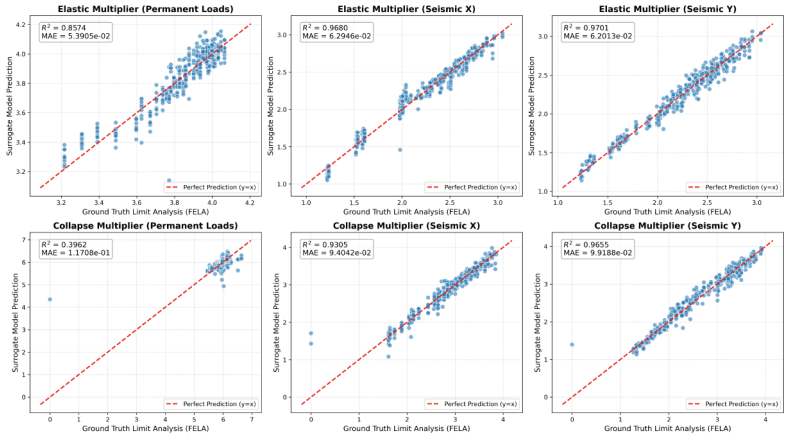

Fig. 5. Parity plots comparing neural-network predictions with finite element limit-analysis ground truth for six target variables. Each panel reports R^2 and MAE.

FELA-based target responses with high accuracy from geometric and material inputs alone.

3.3 Computational Efficiency and Application

The motivation for combining physics-informed machine learning with standard limit analysis is to reduce the computational cost associated with large-scale territorial risk assessment. Table 3 compares the conventional solver with the deployed surrogate model.

With the original solution strategy, estimating ultimate load-bearing capacity requires assembling highly indeterminate stiffness matrices, generating piecewise-linear yield surfaces, and solving a linear-programming problem. The

Table 3. Comparative analysis of algorithmic execution efficiency and underlying mathematical complexity.

Methodology	Mathematical Complexity	Avg. Time per Analysis (s)	Acceleration Factor
Finite Element Limit Analysis	$\mathcal{O}(n^3)$ or $\mathcal{O}(2^n)$	2.450	Baseline
Multi-Output Surrogate Model	$\mathcal{O}(\sum_{l=1}^{L} n_l n_{l-1})$	0.002	1,225x

resulting computational burden, of the order of $O(n^3)$ or higher depending on the problem structure, limits exhaustive scans of large parameter spaces in emergency-planning applications.

Using the pre-trained multilayer perceptron changes the computational workflow to a deterministic sequence of matrix-vector operations through the network layers, with cost scaling as $O(\sum_{l=1}^{L} n_l n_{l-1})$. Based on the values reported in Table 3, the surrogate-model inference is more than 1200 times faster than the original solver. This reduction in runtime supports near-real-time structural capacity assessment and makes the workflow more suitable for seismic micro-zoning studies at the city scale.

4 Discussion

The proposed framework integrates physical mechanics with machine learning in a consistent PIML workflow. In this formulation, the machine-learning component is a standard multi-output multilayer perceptron, whereas the physics-informed aspect arises mainly from the FELA-based generation of the training data and target responses. The surrogate model achieves $R^2 > 0.995$ with small MAE values, indicating that it captures the non-linear relation between input parameters and the FELA-derived responses with high accuracy within the analyzed domain. As shown in Fig. 5, the model response remains close to the identity line over both elastic and collapse states.

The framework also reduces computational cost substantially. The observed speedup of more than 1,200 times helps address one of the main limitations of conventional iterative analyses when applied to large building inventories. By moving most of the computational effort to an offline training stage, the pre-trained surrogate model enables near-real-time structural capacity evaluation during deployment. This property is relevant for rapid parametric scans, building-stock categorization, and the preparation of seismic micro-zoning maps.

4.1 Limitations and Future Directions

Despite the strong predictive performance, the presented PIML model remains constrained by the range and distribution of its training data. As an empirical surrogate, it is expected to be most reliable within the parameter space spanned by the structural geometries, load conditions, and material properties used for

training. Extrapolation to unseen topologies, materially different systems, or substantially different loading scenarios may reduce accuracy. In addition, the surrogate inherits the assumptions embedded in the FELA ground truth, including the adopted boundary conditions, idealizations in material behavior, and the use of the static theorem as the reference description of collapse. A further limitation is that the present evaluation focuses on agreement with the FELA-generated targets and computational efficiency, but it does not include systematic comparisons with alternative surrogate models or other regression architectures. Such comparisons would help clarify whether the observed performance is specific to the chosen multilayer perceptron or can also be achieved by simpler or different models. These limits should be considered when extending the workflow to other structural classes or operating conditions.

Subsequent work can address these limitations by broadening the training space to include a wider variety of morpho-types, such as anisotropic historical masonry and damaged composite structures. It would also be useful to benchmark the present multilayer perceptron against alternative surrogate models, including simpler regression baselines and other non-linear learning architectures, to better position the method within the broader model-selection landscape. Further extensions may incorporate dynamic time-history targets and structural health monitoring data to update the surrogate model over time within a digital-twin framework.

5 Conclusion

This paper presents an integrated computational framework that combines Finite Element Limit Analysis (FELA) with a data-driven surrogate model for large-scale seismic vulnerability assessment. The objective is to reduce the computational cost that limits repeated analyses across wide parametric spaces. By formulating collapse through the static theorem of plasticity, adopting piecewise-linear yield surfaces, and stabilizing the equilibrium system with Singular Value Decomposition, the workflow generates consistent databases of ultimate load-carrying capacity for complex structural topologies.

A multi-output multilayer perceptron surrogate model was then trained to reproduce the non-linear relation between structural geometry, loading, and limit multipliers. In this sense, the learning architecture is relatively standard, while the distinguishing feature of the workflow lies in its coupling with physics-based FELA data generation. The framework was evaluated on five morpho-types: Rectangular, L-type, C-type, Asymmetrical C-type, and H-atrium. The quantitative results show high predictive accuracy, with all reported R^2 values above 0.995 and very small absolute errors across the analyzed targets.

By replacing repeated constraint-oriented linear-programming analyses with a sequence of matrix-vector operations, the proposed framework reduces average evaluation time from 2.450 s to 0.002 s per analysis, corresponding to a speedup of more than 1200 times. Within the investigated domain, this computational gain makes rapid structural vulnerability assessment feasible for geometrically diverse

building configurations. At the same time, broader benchmarking against alternative models would strengthen the evaluation and clarify the relative advantages of the proposed surrogate choice. The workflow may therefore support structural engineers, city planners, and emergency managers in large parametric studies, seismic micro-zoning, and related decision-making tasks.

References

Ahmad, O., Sadeghi, K., Nouban, F.: Optimizing seismic performance assessment: a web-based enhanced visual screening method integrated with machine learning for reinforced concrete structures. Appl. Sci. **16**, 1271 (2026)

Alipour, A., Shafei, B., Shinozuka, M.: Performance evaluation of deteriorating highway bridges located in high seismic areas. J. Bridge Eng. **16**, 597–611 (2011)

Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)

Bleyer, J.: Yield surface approximation for lower and upper bound yield design of 3D composite frame structures. Comput. Struct. **129**, 86–98 (2013)

Cardellicchio, A., et al.: Vulma project: vulnerability analysis using machine learning. Procedia Struct. Integr. **44**, 1714–1721 (2023)

Chen, W.F.: Plasticity, Limit Analysis, Stability And Structural Design: An Academic Life Journey From Theory To Practice. World Scientific (2021)

Di Bonito, L.P., Campanile, L., Di Natale, F., Mastroianni, M., Iacono, M.: Explainable artificial intelligence in process engineering: promises, facts, and current limitations. Appl. Syst. Innov. **7**, 121 (2024) (2024).explainable artificial intelligence in process engineering: Promises, facts, and current limitations.*Applied System Innovation*, 7:121

Ferranti, L., et al.: Seismic vulnerability assessment of urban environments through machine learning algorithms. Buildings **14**, 309 (2024) Seismic vulnerability assessment of urban environments through machine learning algorithms.*Buildings*, 14:309

Ferrari, R.: Computational elastoplastic limit analysis of the paderno d'adda bridge (Italy, 1889). Arch. Civil Mech. Eng. **17**, 374–384 (2017)

Gondaliya, K., et al.: Machine learning-based approach for assessing the seismic vulnerability of reinforced concrete frame buildings. Structures **71**, 105821 (2025) Machine learning-based approach for assessing the seismic vulnerability of reinforced concrete frame buildings. *Structures*, 71:105821

Imani, H., Zona, R., Arcieri, A., Di Bonito, L.P., Palladino, S., Minutolo, V.: Explainable artificial intelligence framework for structure's limit load extimation. In: IoTBDS 2025, pp. 469–480 (2025) Minutolo, V. (2025).explainable artificial intelligence framework for structure's limit load extimation.In *IoTBDS 2025*, pages 469–480

Karniadakis, G.E., et al.: Physics-informed machine learning. Nat. Rev. Phys. **3**, 422–440 (2021)

Lagomarsino, S., Podestà, S.: Macroseismic and mechanical models for the vulnerability and damage assessment of current buildings. Bull. Earthq. Eng. **2**, 415–443 (2006)

Makrodimopoulos, A.: Finite-element limit analysis of mohr-coulomb materials in 3D using semidefinite programming. J. Eng. Mech. **134**, 339–347 (2008)

Murashko, O., Tkachov, Y.: Physics-informed neural networks for solving physical partial differential equations in aerospace engineering. Zenodo (2025)

Málaga-Chuquitaype, C.: Machine learning in structural design: an opinionated review. Front. Built Environ. **8**, 815717 (2022)

Senders, J.T., et al.: Machine learning and neurosurgical outcome prediction: a systematic review. World Neurosurg. **109**, 476–486 (2018)

Seo, J., et al.: Metamodel-based regional vulnerability estimate of irregular steel moment-frame structures subjected to earthquake events. Eng. Struct. **45**, 181–195 (2012)

Spence, R., et al.: Residential building and occupant vulnerability to tephra fall. Bull. Volcanol. **67**, 72–85 (2005) Residential building and occupant vulnerability to tephra fall. *Bulletin of Volcanology*, 67:72–85

Spiliopoulos, K.V., Dais, N.G.: A powerful force-based approach for the limit analysis of three-dimensional frames. Arch. Appl. Mech. **82**, 1111–1127 (2012)

Sun, H., Burton, H.V., Huang, H.: Machine learning applications for building structural design and performance assessment. J. Build. Eng. **33**, 101816 (2021)

Suryawanshi, G.A., Mahajan, L.S., Bhagat, S.R.: Applications of artificial intelligence in structural engineering: a review. IJIES **10**, 25 (2025)

Sweeney, E.M., et al.: A comparison of supervised machine learning algorithms and feature vectors. PLoS ONE **9**, e95753 (2014)

Zheng, H., et al.: A comparative analysis of explainable artificial intelligence (XAI) models for predicting concrete elastic dynamic modulus. Smart Constr. **4**, 100015 (2025) A comparative analysis of explainable artificial intelligence (xai) models for predicting concrete elastic dynamic modulus.*Smart Construction*, 4:100015

Zona, R., Imani, H., Losco, S., Modano, M., Minutolo, V.: Fast evaluation of seismic vulnerability of RC building: a proposal of micro-zoning procedure for safety assessment. Nat. Hazards **121**, 23329–23347 (2025)

Zona, R., Minutolo, V.: A direct finite element limit analysis for assessing the structural safety of masonry under anisotropic behavior. Int. J. Numer. Methods Eng. **126**, e70221 (2025)