Recurrent Neural Networks for Predicting Inflows in Hydroelectric Reservoirs

E. Villar[1], Cecília Coelho[2], M. Fernanda P. Costa[1(✉)], and L. L. Ferrás[3,4]

[1] Centre of Mathematics (CMAT), University of Minho, 4710-57 Braga, Portugal
mfc@math.uminho.pt
[2] Institute for Artificial Intelligence, Helmut Schmidt University, 22043 Hamburg, Germany
cecilia.coelho@hsu-hh.de
[3] Centro de Estudos de Fenómenos de Transporte (CEFT), Faculty of Engineering, University of Porto, 4200-465 Porto, Portugal
lferras@fe.up.pt
[4] ALiCE, Faculty of Engineering, University of Porto, 4200-465 Porto, Portugal

Abstract. Hydropower plays a crucial role in many countries' electricity systems, but its operation often conflicts with ecological preservation and agricultural water needs. Reservoirs must simultaneously deliver reliable energy, maintain minimum ecological flows to protect downstream ecosystems, and guarantee sufficient water for irrigation, making inflow forecasting critical for sustainable operation. In this work, we study neural network-based prediction of inflows to the Miranda hydropower reservoir in Brazil, aiming to provide accurate, data-driven forecasts that can support environmentally conscious and socially responsible decision-making. Using publicly available hydrological data, we perform a systematic exploratory analysis and pre-processing of hourly and daily time-series, including treatment of missing values, negative measurements, and outliers, as well as normalization and feature engineering. We then compare several recurrent neural network architectures (RNN, LSTM, BiLSTM, and GRU) under different training and model configurations. Our results show that a single-layer GRU model trained on the daily, normalized inflow series with a one-day lag achieves the best predictive performance. The selected model is used to produce 7-day ahead forecasts, which closely follow short-term dynamics while exposing limitations in anticipating abrupt peaks, thus highlighting the remaining challenges in forecasting extreme inflow events for hydropower management.

Keywords: Hydropower · Inflow Prediction · Neural Networks · Time-series · Reservoir Management · Data Analysis · Climate AI · Sustainability

1 Introduction

The construction of hydroelectric power plants and the consequent damming of rivers cause profound changes in the landscape and in environmental dynamics

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16763, pp. 272–284, 2027.
https://doi.org/10.1007/978-3-032-30530-5_17

[14]. These impacts are not limited to biological diversity, they also affect social and economic structures [2,7].

As an illustrative example, consider the basin of the Araguari River in Brazil. The sequence of dams built for hydroelectric power generation has led to deep transformations in the region. The formation of reservoirs has had a significant impact on communities that depend on fishing, both as a source of food and as a source of income. Their activities were practically extinguished, since most large fish species that were previously caught in the river cannot reproduce under the environmental conditions imposed by the reservoirs [12].

Another consequence of the construction of reservoirs is the difficulty in guaranteeing the multiple use of water resources. Historically, priority has been given to electricity generation over other potential uses, such as irrigation, recreation, and fish farming [16].

Hydroelectric projects also alter the ichthyological biodiversity of rivers, leading to changes in species composition and a reduction in species richness, which can in many cases result in the decline or even extinction of certain species [6,12].

Hydropower is one of the oldest and most widely used sources of electricity generation in the world and is often considered a relatively clean and renewable alternative compared with fossil fuels [11]. However, the development and operation of hydroelectric projects have sparked intense debate due to their environmental [1] and social [3] impacts.

As mentioned above, the absence of a minimum ecological discharge in hydroelectric projects can have devastating consequences for ecosystems and for local communities. Without an adequate release of water, downstream aquatic ecosystems suffer from a drastic reduction in water levels, which leads to the degradation of habitats that are essential for aquatic fauna and flora. This can result in a decline in biodiversity, affecting fish species and other forms of life that depend on a continuous flow of water to survive [15].

Another often neglected issue is the scarcity of water available to farmers. When water is retained in reservoirs for power generation, the downstream flow frequently decreases significantly, leaving farmers without sufficient access to the water resources needed to sustain their crops. This may cause crop losses, reduced agricultural productivity, and increased costs for farmers, who may be forced to seek alternative, often more expensive and less efficient, water sources [5].

Hydroelectric energy thus plays a crucial role in the provision of sustainable electricity worldwide. Nevertheless, the efficient and responsible operation of these plants poses substantial challenges, especially when it is necessary to balance the maximization of electricity production with ecosystem preservation, through compliance with prescribed ecological releases. In addition, hydroelectric plants must meet water allocation requirements for agricultural use, which further increases the complexity of water resources management [18].

The use of artificial intelligence (AI) in hydroelectric plant management has increased in recent years. In general, however, existing approaches focus primarily on maximizing energy production, without explicitly enforcing minimum

ecological discharges for ecosystem preservation and water allocation for agricultural irrigation [5,6].

Building on [5] and [6], this work focuses on the prediction component of a hybrid model that combines neural networks with constrained optimisation for the management of hydroelectric plants. The goal is to develop and analyse neural network models that can accurately predict reservoir inflows, providing reliable inputs to downstream optimisation procedures and contributing to more transparent and ethically grounded AI-based decision support systems. The goal of the predictive component is to forecast the volume of water that enters the hydroelectric reservoir at each time step. A neural network is implemented to predict these inflows, both on an hourly and on a daily scale. Accurate inflow prediction is crucial for any subsequent optimisation model, since it provides information about future water availability and thus enables more efficient management of current and future water resources.

This paper is organised as follows. Section 2 introduces the case study, the Miranda hydroelectric plant, describing its main characteristics and presenting an exploratory analysis and pre-processing of the hydrological dataset of reservoir inflows. Section 3 presents the neural network-based inflow prediction methodology and reports performance results along with a comparison discussion. Finally, Sect. 4 concludes the work and discusses future research directions.

2 Case Study: Miranda Hydroelectric Plant

In this work, the Miranda hydroelectric plant is considered as a case study since its hydrological data are publicly available in [13] (a database provided in Portuguese), with hourly and daily measurements. The plant is located on the Araguari River, in the municipality of Indianópolis, Minas Gerais, Brazil (see Fig. 1). The reservoir covers an area of 50.61 km^2 and has a useful storage volume of $1.12 \times 10^9 m^3$. The dam has a maximum height of 79 m and a length of 2050 m, with an effective head of 66.4 m for power generation. The installed capacity is 408 MW, distributed across three generating units equipped with vertical Francis turbines of 136 MW each. The firm energy for commercialisation is, on average, 198.2 MW [8].

In this study, the variables of interest are "val_vazaoafluente" and "din_instante". This dataset will be referred to simply as the *dataset*. As mentioned before, measurements are available at an hourly resolution. Time stamps specify year, month, day, and hour, for example: "2021-01-01 01:00:00", "2021-01-01 02:00:00". Figure 2 shows the inflows over time (variable "val_vazaoafluente") for the *dataset* (Table 1).

2.1 Exploratory Data Analysis

Data quality analysis. We start by analysing the variable "val_vazaoafluente". This variable contains one missing value (NaN) and 1434 negative values. Since the physical meaning of this variable is not compatible with negative values, as



Fig. 1. Location of the Miranda hydroelectric plant.

Table 1. Description of variables in the hydrological dataset

Attribute	Type	Description/Values
id_subsistema	Qualitative (nominal)	Subsystem identifier: {SE, S, N, NE}
nom_subsistema	Qualitative (nominal)	Basin regions: {Sudeste/Centro-Oeste, Sul, Norte, Nordeste}
tip_reservatório	Qualitative (nominal)	Reservoir type: Fio d'água; Reservatório com centrais; Reservatório sem centrais; Reservatório com centrais de bombeamento
nom_bacia	Qualitative (nominal)	Basin name (e.g., Amazonas, Uruguai)
id_reservatorio	Qualitative (nominal)	Reservoir identifier (e.g., AMUHBB)
nom_reservatorio	Qualitative (nominal)	Reservoir name (e.g., Balbina, São Roque)
cod_central	Quantitative (discrete)	Power plant code (e.g., 88, 277)
din_instante	Qualitative (ordinal)	Timestamp (year, month, day, hour)
val_vazaoafluente	Quantitative (continuous)	Inflow discharge to the reservoir
val_nivelmontante	Quantitative (continuous)	Upstream water level
val_niveljusante	Quantitative (continuous)	Downstream water level
val_vazaodefluente	Quantitative (continuous)	Outflow discharge from the reservoir
val_vazaoturbinada	Quantitative (continuous)	Turbined flow used for energy generation
val_vazaovertida	Quantitative (continuous)	Spilled flow not used for energy generation
val_vazaotransferida	Quantitative (continuous)	Water transferred to another reservoir/location
val_vazaovertidanaoturbinavel	Quantitative (continuous)	Spilled flow not passing through turbines and not usable for generation

Dataset composed of 26,274 records and 18 attributes, constructed from hourly hydrological data files covering the period 2021–2023 [13].

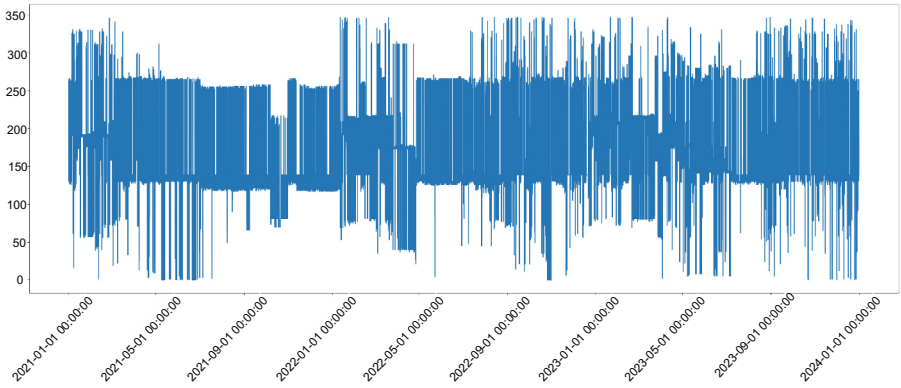

Fig. 2. Visualization of the inflows over time.

it represents the volume of water entering the reservoir, these negative entries were treated as measurement errors and thus removed. Furthermore, Fig. 3 shows a highly right-skewed distribution for this variable in which most observations are concentrated in a relatively low range, with the median located well below the upper whisker. The interquartile range (IQR) measure is fairly compact, indicating that typical inflow values do not vary dramatically. However, there is a large number of outliers (with 3473 identified) extending far above the upper whisker, reaching values above 1200. This suggests the presence of extreme inflow events which can be associated with heavy rainfall periods or seasonal hydrological peaks. The lower whisker is close to zero, showing that very low inflow conditions also happen, but without significant lower-end outliers.

Data Cleaning. Guided by the data analysis, a copy of the dataset was created in which outliers (the NaN value, and the negative values) were replaced by the median. These replacements affect approximately 19% of the data.

Other strategies were also tested: replacing negative values by zero; replacing outliers with the mean of their adjacent neighbours; and using the *SimpleImputer* class from *sklearn.impute* with the strategies *mean*, *median*, *most_frequent*, and *constant*. It was observed that replacing these values by the median leads to better performance when modelling inflows, measured in terms of the mean squared error (MSE).

Table 2 reports the four best MSE values obtained with these strategies, denoted E1 to E4.

Thus, replacing outliers and negative values by the median yields the lowest MSE. The other imputation strategies led to higher MSE values than those reported in Table 2. Since strategy E4 achieves the lowest MSE, this cleaned version of the data is used in the subsequent analysis and will be referred to as Dataset-H.

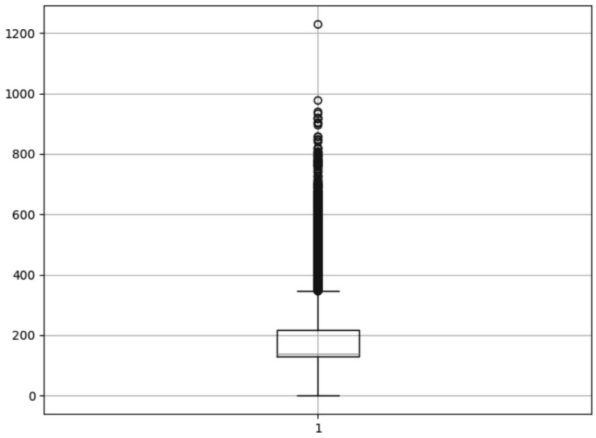

Fig. 3. Boxplot of the inflow discharge variable "val_vazaoafluente". The distribution is strongly right-skewed, with a compact interquartile range and numerous high-value outliers.

Table 2. Best MSE values obtained for different strategies for replacing outliers and negative values.

Strategy	Negative values	Outliers	MSE
E1	Neighbour mean	Neighbour mean	5337
E2	Neighbour mean	Not replaced	10545
E3	0	Median	3354
E4	Median	Median	2628

The attribute "din_instante" was converted to *datetime* format, ensuring that the data follow a correct chronological order.

For the purposes of this work, a daily dataset was also created to reduce the temporal resolution. To this end, hourly inflows were aggregated by day using a sum, yielding a daily dataset, Dataset-D, derived from Dataset-H.

Data Analysis. We first investigate how many past lags (days or hours) are most appropriate for predicting the inflow for the next day (or hour).

As an initial step, we study the autocorrelation function, which measures the correlation of a time-series with itself at different lags. Figure 4 shows the autocorrelation function (ACF) for Dataset-D at different time lags, helping to identify repeating patterns (seasonality). The results indicate that autocorrelation is strongest at with a time lag of 1, suggesting that inflows are most strongly correlated with the previous day.

The partial autocorrelation function (PACF) measures the correlation between a time-series and its own lags, while controlling for the effect of intermediate lags. While the ACF captures both direct and indirect correlations between

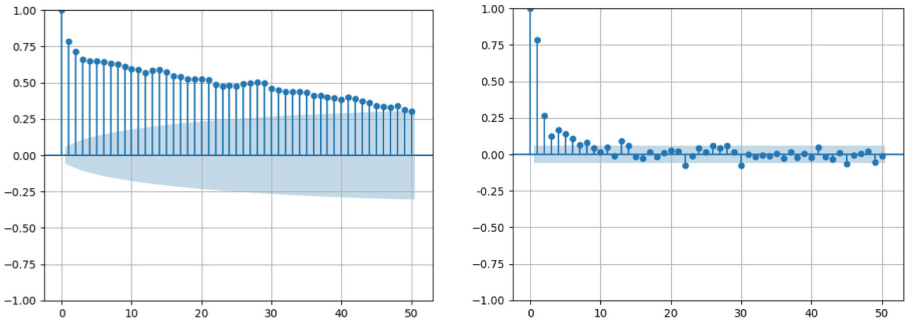

Fig. 4. Autocorrelation (left) and partial autocorrelation (right) function over different time lags for Dataset-D.

current and past values, the PACF removes the effect of preceding lags and focuses on the direct correlation. Figure 4 shows the PACF for different lags. The PACF is useful for identifying the order of autoregressive models, where significant peaks indicate lags that have a direct relationship with the current value after removing the influence of earlier lags. As with the ACF, the largest partial autocorrelation occurs at a time lag of 1.

Note that, the results for Dataset-H are not displayed due to the similarity with the ones obtained for Dataset-D.

From this analysis, only the previous day (or hour) is used in the pre-processing stage to predict the inflow at the current day (or hour).

Since Dataset-D has a coarser temporal resolution, an additional study was performed to decompose the data into trend, seasonality, and residual components. For this purpose, the `seasonal_decompose` function from the *statsmodels* Python package was used. This decomposition is useful for understanding the dynamics of inflows over time, separating long-term trends, seasonal behaviour, and unexplained fluctuations. In the pre-processing stage we will explore whether including trend, seasonal, and residual components as additional features benefits the predictive performance.

Note that the attributes used as inputs correspond to the previous day, and the inflow to be predicted for the current day is not fed back as an input.

Figure 5(a) shows the daily inflow time-series. The series exhibits high variability, with significant fluctuations. Some seasonal oscillations can be observed, with inflow peaks in certain periods followed by sharp drops. There appears to be an overall stabilization or slight decline in the most recent data, after a period of increase at the beginning of 2022.

In Fig. 5(b), the trend reflects the long-term direction of the data, smoothing short-term fluctuations. It suggests a gradual increase in inflows from mid-2021 until the end of 2022. From 2023 onwards, there is a slight deceleration or decline in the trend, indicating stabilization or a potential decrease in reservoir inflows.

Figure 5(c) shows the seasonal component, which represents repetitive and predictable patterns occurring at regular intervals. It reveals marked seasonal

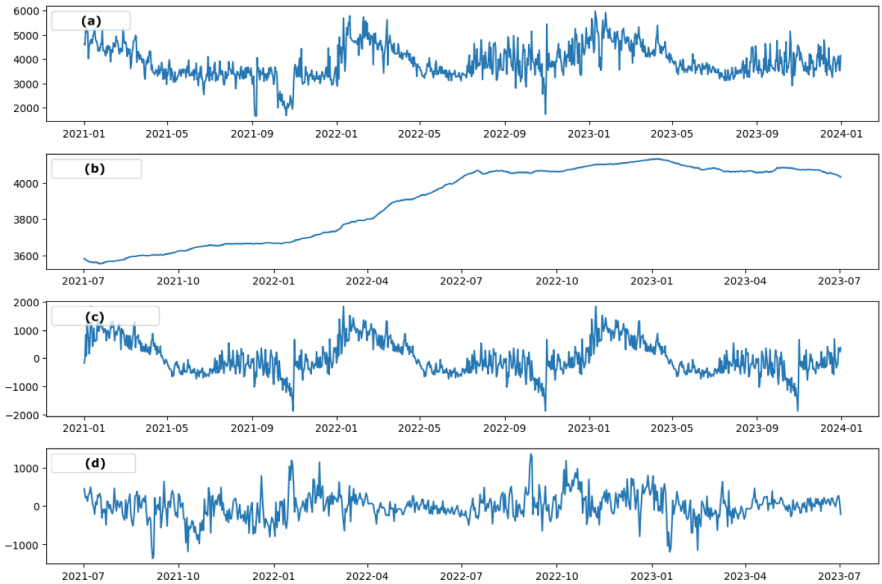

Fig. 5. Data decomposition analysis for Dataset-D: (a) original daily inflow series; (b) trend; (c) seasonal component; (d) residuals.

variations, with peaks and troughs that likely correspond to climatic cycles such as rainy and dry seasons.

Figure 5(d) displays the residual component, which contains the remaining variability after removing trend and seasonality from the original series. The residuals appear largely random, with no clear pattern, indicating that most of the structure in the series has been captured by the trend and seasonal components. A few pronounced peaks remain, which may correspond to isolated events or anomalies not explained by the other components.

Pre-processing. As discussed above, we use a time lag of 1, *i.e.* one day for Dataset-D and one hour for Dataset-H, and the target variable $\{y_t\}_{t=1}^{T}$ is the inflow volume to the reservoir at a given day (or hour). Several pre-processing techniques were investigated to improve the performance of the neural network models, including log-transformation, differencing, standardisation, noise injection, normalisation, and the creation of additional features:

- Log transformation: Applying a logarithmic transformation to the target using *numpy*, i.e., $y \leftarrow \ln(y_t)$, where ln denotes the natural logarithm;
- Differencing: A standard technique in time-series analysis to obtain stationarity by removing long-term trends or seasonal patterns. The first difference is computed as $d_t = y_t - y_{t-1}$;
- Standardization: Transforming the target to have zero mean and unit variance, $y_t \leftarrow \dfrac{y_t - \mu}{\sigma}$, where μ and σ denote the mean and standard deviation

of y, respectively. This places all variables on a common scale so that they contribute proportionally during model training;
- Noise injection: Adding noise to the training data, in which the noise level controls the amount of random perturbation added to the data and is often defined as a fraction of the standard deviation, since σ reflects the natural scale of the data. Noise injection can increase model robustness and help prevent overfitting by encouraging the model to tolerate small variations: $y_t \leftarrow y_t + \varepsilon$, where ε is a random perturbation with prescribed variance;
- Normalisation: Scaling the data to a fixed interval. Here, values are mapped to $[0, 1]$ via $y_t \leftarrow \dfrac{y_t - y_{\min}}{y_{\max} - y_{\min}}$, where $y_{\min}$ and $y_{\max}$ are the minimum and maximum observed values of $(y_0, y_1, \ldots, y_N)$. This is particularly useful since neural networks are sensitive to the magnitude of input variables;
- Feature engineering: Using the `seasonal_decompose` function from *statsmodels* to decompose the target into trend, seasonal, and residual components and including these as additional features. For Dataset-H, additional calendar features were also created from "din_instante": "day", "month", and "hour". Using these features and the elbow method, a clustering step with 11 clusters was performed, resulting in an additional categorical feature "cluster".

Among all tested techniques, normalisation yielded the best predictive performance in terms of the adopted evaluation metrics. Therefore, normalisation was chosen as the main pre-processing step for the development of this work.

3 Neural Networks for Inflow Prediction

Different recurrent neural network architectures were tested: Long Short-Term Memory (LSTM) [9], Simple Recurrent Neural Network (RNN) [10], Bidirectional LSTM (BiLSTM) [17], and Gated Recurrent Unit (GRU) [4]. In addition, several network design choices were explored, including increasing or decreasing the number of neurons per layer, changing the number of layers, adding regularization and/or dropout layers, modifying the learning rate and/or the optimiser, changing the activation functions, and varying the weight initialisation schemes.

Both datasets were split into 80% training, 10% test, and 10% validation sets, while preserving the chronological order.

Note that, the neural network architectures were tested on both hourly (Dataset-H) and daily (Dataset-D) datasets. However, while the finer temporal granularity of Dataset-H captures short-term fluctuations, it also amplifies noise from measurement errors and irregular hourly variations, leading to unstable training dynamics and inferior generalisation performance across all architectures. For these reasons, we focus subsequent analysis and modelling exclusively on Dataset-D.

Table 3 reports the root mean squared error (RMSE) and mean absolute percentage error (MAPE) for the different strategies and approaches, using Dataset-D. The acronym NSOVN indicates there was no replacement of outliers and negative values.

Table 3. Results for different strategies and networks for the daily data, Dataset-D.

Scenario	Strategy	Method	Network	RMSE	MAPE
1	E4	Normalisation	RNN	508.6	11
2	NSOVN	Normalisation	RNN	606.6	16
3	E4	Normalisation	LSTM	611.8	14
4	E4	Normalisation	GRU	364.9	7

All networks were trained with the Adam optimiser for 100 epochs, using a batch size of 64 and an early-stopping callback with patience equal to 10. The models are configured to generate forecasts over a 7-day horizon in a recursive manner, *i.e.*, the predicted value for a given day is fed back as input to predict the inflow for the next day. All implementations were done using *Pytorch*, while data analysis and preprocessing was done using *Pandas* and *Scikit-learn*.

The analysis of Table 3 shows that Scenario 4 achieves the best performance. The selected model consists of a single GRU layer with 100 neurons, taking as input the inflow value of the previous day, followed by a dropout layer with dropout rate 0.1. The training process of the chosen network was analysed by monitoring the evolution of the training loss and validation loss (val_loss), both computed as MSE, Fig. 6. The rapid decrease of the loss during the first epochs indicates that the model is able to learn effectively from the data, with only minor improvements after approximately four epochs. Moreover, the training and validation loss curves remain close and do not diverge significantly, suggesting that the model is well fitted and does not exhibit clear signs of overfitting.

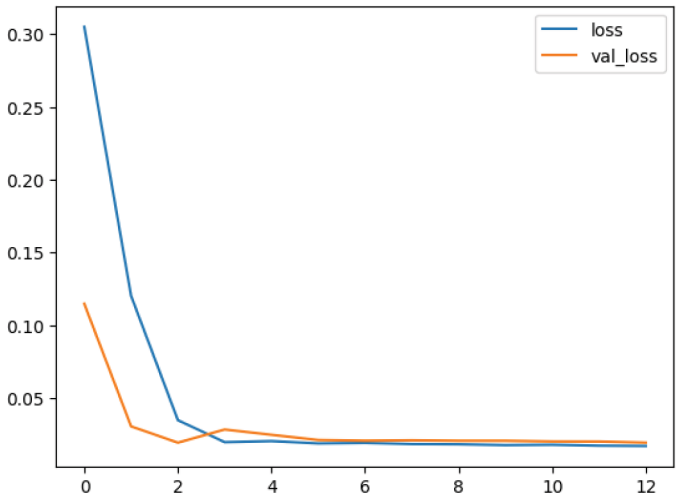

Fig. 6. Evolution of the training and validation loss for the best performant model.

The early-stopping callback interrupted training after 12 epochs since the validation performance did not improve for 10 consecutive epochs. Other configurations, such as different weight initialisations, additional dropout layers, regularisation, alternative activation functions, or modified learning rates did not lead to further performance improvements.

Table 4 reports the predicted inflows for the 7-day horizon, together with the corresponding observed values.

Table 4. Predicted versus observed inflow values for a 7-day horizon.

Day	Observed value	Prediction
1	3581	3596
2	3598	3513
3	3522	3471
4	3516	3439
5	4074	3414
6	3932	3396
7	3828	3382

From Table 4, it can be seen that, for the first three days, the model performance is relatively good as the forecasts follow the observed trend, and the error remains within an acceptable range. On day 4, however, a significant event (a pronounced peak) occurs in the observed series, which the model fails to predict. Due to the recursive forecasting scheme, any prediction error is fed back as if it were a true input, and thus tends to accumulate and amplify over time. Another contributing factor may be the limited ability of the network to capture abrupt changes as the model appears to have learnt a smoothed pattern, and consequently fails to reproduce this anomalous event. Despite these challenges, the overall results confirm the effectiveness of neural networks, in particular GRU, for reliable short-term inflow prediction in hydropower applications.

4 Conclusion

This work has demonstrated the feasibility of neural network-based inflow prediction for the Miranda hydroelectric reservoir using three years of hydrological data. Through exploratory data analysis, cleaning strategies, and feature engineering, we constructed a high-quality hourly (Dataset-H) and daily (Dataset-D) time-series, with the daily dataset proving superior for predictive modelling.

Among the tested recurrent architectures (Recurrent Neural Network, Long-Short Term Memory, Bidirectional Long-Short Term Memory, and Gated Recurrent Unit), a single-layer Gated-Recurrent Unit model with 100 neurons, trained on normalised daily inflows using a one-day lag, achieved the best performance, with root mean square error of 364.9 and mean absolute percentage error of

7%. The 7-day ahead forecasts accurately captured short-term dynamics but struggled with abrupt peaks, demonstrating the challenges of modelling extreme inflow events in recursive schemes where errors accumulate over time.

These findings provide a reliable predictive component for sustainable reservoir management, where accurate inflow forecasts are essential to balance energy production with ecological minimum flows and agricultural water demands.

Future work will focus on improving peak prediction through hybrid architectures (e.g., attention mechanisms or CNN-LSTM ensembles), incorporating meteorological covariates, and extending the forecast horizon while maintaining low error rates. Integrating this predictor into a constrained optimisation framework will enable sustainable operating policies for the Miranda plant.

Acknowledgments. We thank the support of the High-Performance Computing Center at the University of Évora funded by FCT I.P. under the project "OptXAI: Constrained optimisation in NNs for Explainable, Ethical and Greener AI", reference 2024.00191.CPCA.A1, platform Vision. This work was financially supported by Fundação para a Ciência e Tecnologia (FCT) through the FCT/MCTES (PIDDAC) under the project 2022.06672.PTDC-iMAD (Improving the Modelling of Anomalous Diffusion and Viscoelasticity: solutions to industrial problems), DOI 10.54499/2022.06672.PTDC (https://doi.org/10.54499/2022.06672.PTDC); and by the projects LA/P/0045/2020 (ALiCE), UIDB/00532/2020, and UIDP/00532/2020 (CEFT), and UID/00013/2025 (https://doi.org/10.54499/UID/00013/2025) (CMAT). It was also financially supported by Fundação "la Caixa"|BPI and FCT through project PL24-00057: "Inteligência Artificial na Otimização da Rega para Olivais Resilientes às Alterações Climáticas". C. Coelho would like to thank the KIBIDZ project funded by dtec.bw-Digitalization and Technology Research Center of the Bundeswehr; dtec.bw is funded by the European Union-NextGenerationEU.

References

1. Botelho, A., Ferreira, P., Lima, F., Pinto, L.M.C., Sousa, S.: Assessment of the environmental impacts associated with hydropower. Renew. Sustain. Energy Rev. **70**, 896–904 (2017)
2. Catolico, A., Maestrini, M., Strauch, J., Giusti, F., Hunt, J.: Socioeconomic impacts of large hydroelectric power plants in Brazil: a synthetic control assessment of Estreito hydropower plant. Renew. Sustain. Energy Rev. **151**, 111508 (2021)
3. Cernea, M.M., et al.: Social impacts and social risks in hydropower programs: preemptive planning and counter-risk measures. In: Keynote address: Session on social aspects of hydropower development. United Nations Symposium on Hydropower and Sustainable Development Beijing, China (2004)
4. Cho, K.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014)
5. Coelho, C., Costa, M.F.P., Ferrás, L.L.: The influence of neural networks on hydropower plant management in agriculture: addressing challenges and exploring untapped opportunities. In: Mercier-Laurent, E., Kayakutlu, G., Owoc, M.L., Wahid, A., Mason, K. (eds.) IFIP International Workshop on Artificial Intelligence for Knowledge Management, pp. 150–160. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-61069-1_11

6. Coelho, C., Jing, M., Costa, M.F.P., Ferrás, L.L.: An adaptive hydropower management approach for downstream ecosystem preservation. arXiv preprint arXiv:2403.02821 (2024)
7. De Faria, F.A., Davis, A., Severnini, E., Jaramillo, P.: The local socio-economic impacts of large hydropower plant development in a developing country. Energy Econ. **67**, 533–544 (2017)
8. ENGIE: Miranda hydroelectric power plant (2026). https://www.engie.com.br/en/usinas/miranda-hydroelectric-power-plant/. Accessed 26 Feb 2026
9. Hochreiter, S.: Long short-term memory. Neural Comput. (1997)
10. Keras: SimpleRNN. https://keras.io/api/layers/recurrent_layers/simple_rnn/. Accessed 26 Feb 2026
11. Killingtveit, Å.: Hydropower. In: Managing Global Warming, pp. 265–315. Elsevier (2019)
12. Loures, R.C., Pompeu, P.S.: Temporal changes in fish diversity in lotic and lentic environments along a reservoir cascade. Freshw. Biol. **64**(10), 1806–1820 (2019)
13. ONS: Dados hidrologicos horario 2024-01 (2024). https://dados.ons.org.br/dataset/dados_hidrologicos_ho/resource/ca306dc0-abbe-47d2-af9d-da601fc2c09f?inner_span=True. Accessed 26 Feb 2026
14. Pandit, M.K., Grumbine, R.E.: Potential effects of ongoing and proposed hydropower development on terrestrial biological diversity in the Indian Himalaya. Conserv. Biol. **26**(6), 1061–1071 (2012)
15. Renöfält, B.M., Jansson, R., Nilsson, C.: Effects of hydropower generation and opportunities for environmental flow management in Swedish riverine ecosystems. Freshw. Biol. **55**(1), 49–67 (2010)
16. Ribas, J.R.: An assessment of conflicting intentions in the use of multipurpose water reservoirs. Water Resour. Manage. **28**(12), 3989–4000 (2014)
17. Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. IEEE Trans. Sig. Process. **45**(11), 2673–2681 (1997)
18. Zeng, R., Cai, X., Ringler, C., Zhu, T.: Hydropower versus irrigation—an analysis of global patterns. Environ. Res. Lett. **12**(3), 034006 (2017)