Baltic J. Modern Computing, Vol. 13 (2025), No. 2, pp. 315–330
https://doi.org/10.22364/bjmc.2025.13.2.01

A Neural Network-Based Causal Model for Electricity Demand Estimation in Remote Areas: A Case Study in El Espino, Bolivia

Stefano SANFILIPPO[1] ⋆, José Juan HERNÁNDEZ-CABRERA[2] ⋆⋆, Christoph KÄNDLER[3] ⋆⋆⋆, José Juan HERNÁNDEZ-GÁLVEZ[2] †, José ÉVORA-GÓMEZ[2] ‡, Octavio RONCAL-ANDRÉS[2] §

[1] STAM S.r.l, Genoa, Italy
[2] Instituto Universitario de Sistemas Inteligentes y Aplicaciones Numéricas en Ingeniería - Universidad de Las Palmas de Gran Canaria, Las Palmas, Gran Canaria, Spain
[3] EIFER - Europäisches Institut für Energieforschung, Karlsruhe, Germany

Abstract. Designing microgrids in remote areas is challenging due to the lack of reliable and high-quality electricity demand data. These limitations arise from technological, economic, and logistical constraints, making it difficult to estimate demand—especially when daily demand curves cannot be fully constructed due to missing or incomplete data. Traditional methods, which rely on consistent and comprehensive datasets, often prove ineffective in such conditions.
To address this issue, this paper introduces a novel causal modelling approach, implemented using a neural network, to uncover the underlying relationships between key influencing factors—such as temperature, humidity, time of day, and seasonal variations—and electricity demand. Rather than requiring complete hourly demand curves as inputs, the proposed approach leverages available data to infer demand patterns more effectively.
We propose a neural network architecture that aims to capture causal dependencies in electricity demand by encoding input features into a high-dimensional latent space. Using an encoder-decoder structure, the encoder maps inputs to a latent space designed to preserve potential causal relationships, while the decoder generates the demand estimation. This approach hypothesizes that this configuration may help to get causal dependencies. To evaluate this, we compared our model against a simpler neural network architecture characterised by a triangular layer structure.

⋆ s.sanfilippo@stamtech.com ORCID: 0009-0001-0547-6222
⋆⋆ josejuan.hernandez@ulpgc.es ORCID: 0000-0003-2427-2441
⋆⋆⋆ christoph.kaendler@eifer.org ORCID: 0000-0002-0873-1137
† jose.galvez@ulpgc.es ORCID: 0009-0008-3626-7520
‡ jose.evora@ulpgc.es ORCID: 0000-0001-9348-7265
§ octavio.roncal@ulpgc.es ORCID: 0000-0003-3503-3833

Using real-world data from El Espino, Bolivia, our model achieved a Mean Squared Error (MSE) of 0.0511 with the Adam optimiser, representing a 61.8% improvement over the simpler neural network architecture.
A sensitivity analysis further confirmed the relevance of selected input variables, showing that excluding temporal-based features, such as the month of the year and weekend indicator, increased estimation error, with an 11.7% increase in MSE. These findings highlight the model's effectiveness in handling data limitations and its potential as a scalable solution for electricity demand estimation in remote areas.

Keywords: Microgrids, Remote areas, Electricity demand, Causal Model, Estimation, Neural Networks

1 Introduction

Microgrids are a key strategy for achieving electrification in remote areas (Nayanathara and Srilatha, 2018), as they integrate distributed energy resources to provide localised electricity generation, avoiding the challenges of transmission and distribution in inaccessible locations. Successfully designing such systems requires an accurate understanding of demand. Overestimating demand can inflate costs unnecessarily, while underestimating it may lead to undersized systems, which are unable to consistently meet energy needs (Sanfilippo and et al., 2023).

The term demand encompasses a range of needs, including electricity (Castillo et al., 2022), heat (Białek et al., 2022), and cooling (Abugabbara et al., 2022). In this study, the focus is specifically on electricity demand in remote areas, where reliable access to electricity not only improves local economic opportunities but also contributes to enhanced energy efficiency and integration of renewable energy (Stadler et al., 2016).

However, estimating electricity demand in remote areas is inherently challenging. In villages awaiting electrification, baseline demand data is often non-existent. Even in electrified areas, data can be sparse, unreliable, or of poor quality (Wassie and Ahlgren, 2023). Such data limitations hinder the ability to estimate demand accurately, making it difficult to design microgrids efficiently.

Addressing these constraints requires the development of robust, data-driven models tailored to conditions where high-quality, granular demand information is short. The case of El Espino, Bolivia, which offers an unusually large dataset compared to similar contexts, provides a good opportunity to apply artificial intelligence techniques for customised electricity demand estimation. The proposed approach aims to produce a model capable of delivering reliable estimations in remote areas facing persistent information gaps by leveraging this data-rich environment and integrating meta data (temperature and day time) to create a general representative model.

Given the challenges associated with electricity demand estimation in remote areas, various approaches have been explored to minimise the error and improve robustness. Existing methods generally fall into three categories: traditional statistical methods, computational intelligence methods, and hybrid approaches that combine both. This section provides an overview of these methodologies and their effectiveness in addressing the limitations of demand estimation in data-scarce environments.

2 Related Work

In recent years, the energy sector has encountered significant challenges in accurately estimating electricity demand. Various methodologies have been developed to address these challenges, which can generally be categorised into three main approaches: Traditional Statistical Methods, Computational Intelligence Methods, and Hybrid Methods. Each of these approaches offers distinct advantages and limitations, depending on the complexity of the problem, the availability of data, and the need for minimisation of error in forecasting.

Traditional statistical methods are grounded in mathematical principles and rely on historical data patterns, probability distributions, and regression techniques to predict electricity demand. Among these, time series analysis (Velasquez et al., 2022; Dilaver and Hunt, 2011) represents a fundamental approach, leveraging past consumption patterns to identify trends and seasonal variations that inform future demand projections. This method has been widely applied due to its interpretability and strong theoretical foundation. Additionally, econometric models (Dieudonné et al., 2022; Nasr et al., 2000; Gómez and Rodríguez, 2019; Zamanipour et al., 2023) extend the traditional statistical approach by incorporating relationships between electricity demand and macroeconomic indicators such as gross domestic product, economic growth, urbanization, and financial development. These models provide valuable insights into long-term dependencies and external factors affecting energy consumption patterns.

Further expanding on statistical modelling, probability-based approaches are employed to handle uncertainties in electricity demand estimation. Stochastic methods (Lombardi et al., 2019), for instance, integrate randomness into forecasting models, allowing for flexible predictions that account for variations in external influences such as weather fluctuations and human behaviour. Similarly, structural models (Michalik et al., 1997) adopt an engineering-based perspective, focusing on the physical and technical characteristics of the electricity system to determine demand based on infrastructure constraints and efficiency measures.

As the complexity of electricity demand forecasting has increased, computational intelligence methods have gained prominence. These methods, often associated with artificial intelligence and machine learning, are designed to handle non-linear relationships and large-scale datasets, surpassing the predictive capabilities of traditional statistical techniques. Neural networks, for example, have been widely used to model electricity demand by learning intricate consumption patterns through adaptive training mechanisms (Kandananond, 2011; Foldvik Eikeland et al., 2021). Similarly, Random Forest (Shin and Woo, 2022) has been applied to electricity consumption forecasting, offering enhanced accuracy and robustness in handling diverse input variables.

Beyond traditional statistical and computational intelligence approaches, hybrid methods (Shiraki et al., 2016) have emerged as a powerful alternative by integrating multiple forecasting techniques to improve accuracy and adaptability. By combining statistical models with AI-driven approaches, hybrid techniques mitigate the limitations of individual methods and leverage their strengths. For instance, scenario-based methods (Xia et al., 2022) often incorporate both econometric and machine learning components to assess how external variables such as climate change, economic fluctuations, and policy decisions influence electricity demand. Similarly, Geographic Information Systems methods (Torabi Moghadam et al., 2018) benefit from the integration of traditional spatial analysis with computational intelligence techniques to improve region-specific demand estimation, considering factors such as population density, land use, and climatic conditions.

Furthermore, agent-based models highlight another key area where hybridization of techniques proves beneficial. These models simulate the behaviour of individual consumers or groups, allowing for a more dynamic representation of energy demand influenced by social and behavioural patterns (Tian and Chang, 2020). When combined with stochastic and machine learning techniques, agent-based models become highly effective in capturing demand variability and providing more refined insights into consumption trends.

Overall, the landscape of electricity demand forecasting has evolved through the interplay of traditional statistical methods, computational intelligence approaches, and hybrid methodologies. While traditional statistical models offer well-established theoretical foundations and interpretability, computational intelligence techniques provide superior predictive power in handling complex, high-dimensional data. Hybrid methods, by integrating these diverse approaches, present a promising direction for enhancing accuracy and adaptability in demand estimation. The selection of an appropriate method depends on factors such as data availability, forecasting horizon, and the specific characteristics of the electricity market under analysis.

While these approaches have contributed significantly to demand estimation, they often rely on pattern recognition rather than explicitly modelling the causal factors driving electricity consumption. Traditional statistical methods assume stable demand patterns, while computational intelligence methods, such as neural networks, excel at pattern recognition. Hybrid approaches attempt to bridge this gap, yet they remain constrained by data limitations. To overcome these challenges, we propose a model that explicitly encodes causal relationships, enabling a more interpretable and robust framework for electricity demand estimation in remote areas.

3 Proposed Solution

Modelling electricity demand poses significant challenges due to its complex and nonlinear nature, as well as the intricate interdependencies that arise, particularly because user behaviour plays a central role (Lazzari et al., 2022). Traditional demand estimation methods rely on predefined demand curves, which assume stable and well-defined consumption patterns. However, these approaches fail to capture the evolving and dynamic nature of electricity demand, particularly in environments with incomplete or unreliable data. Moreover, they are based on correlations rather than identifying the causal mechanisms that drive demand variations.

Electricity demand is not merely the sum of independent factors but rather the result of dynamic interactions between environmental conditions, socio-economic factors, and technological adoption. These elements influence each other, creating causal dependencies that cannot be fully understood through conventional demand curve-based models alone. As a result, there is a need for an approach that moves beyond static demand profiles to one that models the underlying factors driving electricity consumption.

To address this, the proposed model is designed to describe the causal relationships between these factors, rather than relying on predefined demand curves or purely correlational patterns. The final proposal integrates variables—such as temperature, humidity, time of day, month, and whether it is a weekday or weekend—which were selected based on expert domain knowledge, ensuring that they are known to influence electricity demand. By incorporating expert, driven insights, the model is designed to capture the actual causal factors driving consumption, rather than relying solely on statistical correlations.

By structuring the model around cause-effect relationships, it aims to represent the way external conditions and user behaviour interact to shape electricity consumption. This approach allows for a more interpretable and robust demand estimation framework, capable of adapting to scenarios with incomplete or noisy data while maintaining a meaningful representation of the underlying processes driving electricity demand. The core innovation of this work lies in designing a neural network architecture specifically to capture these causal relationships, rather than merely identifying patterns, a fundamental shift from how neural networks are typically used.

Traditional neural networks excel at recognizing statistical dependencies in data but do not inherently distinguish between correlation and causation. A neural network is a computational model inspired by the structure and function of neurological systems, designed to represent complex, non-linear functions (Neervannan, 2018). It consists of interconnected layers of neurons, with each layer's arrangement determining its specific purpose or function. Optimised using methods such as backpropagation, neural networks effectively map input variables to output estimations, making them ideal for modelling systems like electricity demand, where relationships are non-linear and evolve over time.

Unlike traditional models, neural networks automatically identifying patterns and relations within multi variable datasets and utilize these insights of usually unseen couplings (Scarborough and Somers, 2006). Their ability to capture deep and complex dynamics allows them to represent emergent behaviours, surpassing the limitations of linear or polynomial models (Somers and Casal, 2009).

However, causality remains a challenge in neural networks, as they are traditionally designed to identify patterns rather than capture causal relationships. Our hypothesis is that by encoding inputs into a higher-dimensional space, it becomes possible to reveal underlying causal relationships that may not be apparent in the original input space. By doing so, the network moves beyond conventional pattern recognition, instead aiming to represent how different factors interact and influence electricity consumption in a structured manner.

In the context of electricity demand estimation, this means that the model could begin to identify meaningful dependencies that drive consumption, rather than merely recognizing statistical correlations. We posit that by leveraging the expressive power of high-dimensional representations, the model can uncover the true interactions that govern demand variations.

4 Data Collection, Preparation, and Analysis

During the data collection and analysis phase, available data were gathered, explored, and examined. Previous studies have aimed to establish common data platforms (Fioriti et al., 2023) that streamline this process. Such platforms would significantly reduce the effort required to obtain authentic electricity demand data, currently a task involving extensive literature reviews, contacting authors, interviewing stakeholders, and identifying specialised websites.

Although real measured data were identified for locations in Bolivia, Namibia, Mexico and Tanzania, many datasets contained missing measurements. After removing incomplete curves, the final dataset included a total of 869 days of measured electricity demand. Among these, El Espino in Bolivia provided the largest number of complete days, 578, which surpassed any other region under consideration. Consequently, El Espino was chosen for this study due to its comparatively abundant and reliable data.

The El Espino dataset, sourced from a GitHub repository (Balderrama Subieta, 2022), spans from 1 January 2016 to 31 July 2017, covering 578 days of recorded measurements. It encompasses data from 128 households, a hospital, a school, and street lighting systems and the wattage consumed at each timestamp. First, the data were preprocessed to address inconsistencies, remove outliers, and format it into a structured, consolidated, and normalised dataset.

Visual examinations were conducted to illustrate inherent variability in energy demand patterns. Figures 1 and 2 show the distribution of power usage by hour and by month, respectively, with outliers indicated by dots. Outliers were identified using the three-standard-deviation rule, where any data point x that deviates more than three times the standard deviation (σ) from the mean (μ) is considered an outlier. Mathematically, this is expressed as in Equation 1.

$$x < \mu - 3\sigma \quad \text{or} \quad x > \mu + 3\sigma. \tag{1}$$

By applying this approach, we mitigate the risk of extreme values disproportionately affecting the learning process. This helps preventing the exploding gradient problem, ensuring stable and efficient model training.

This method assumes a normal distribution of energy demand data and effectively detects extreme variations while maintaining robustness in identifying significant deviations. Figure 1 reveals a minimum at 8 a.m. and a peak at 8 p.m., while Figure 2 indicates that October experiences the highest demand.

Figure 3 illustrates the energy demand on weekends, highlighting significant patterns influenced by the hour of the day, the month of the year, and the effects seen during weekend. The heatmap reveals a distinct hourly pattern, with higher demand during specific hours, such as evenings. Additionally, a seasonal effect is evident, with increased energy usage during colder months, likely due to heating needs, and during warmer months, possibly due to cooling systems. The effects seen during weekend is also noticeable, as the patterns differ from those typically observed on weekdays, reflecting variations in social and economic activities. These insights are crucial for modelling energy demand and considering time-of-use factors in energy estimation.

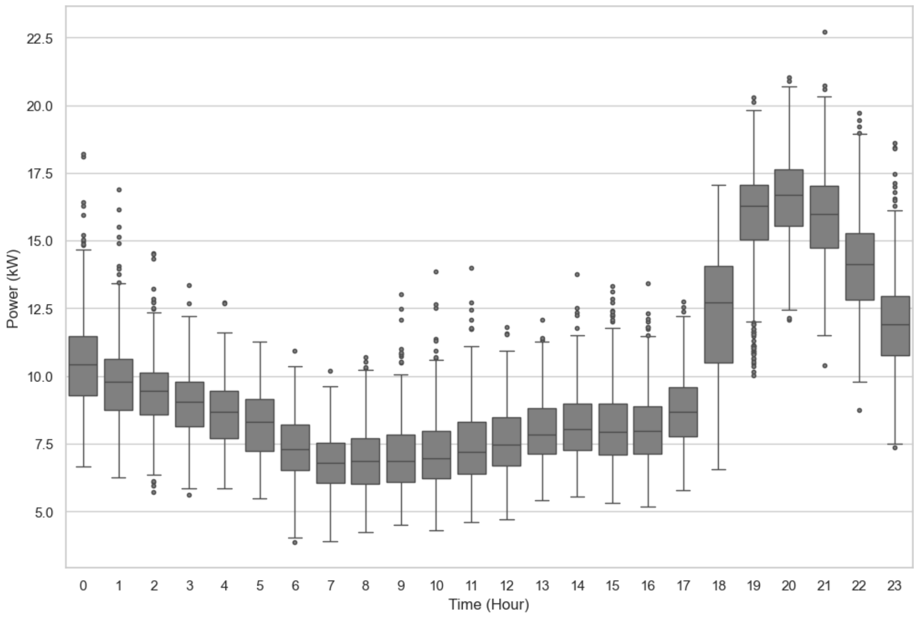

Fig. 1. Hourly Distribution of Electricity Demand Represented as a Boxplot with Outliers

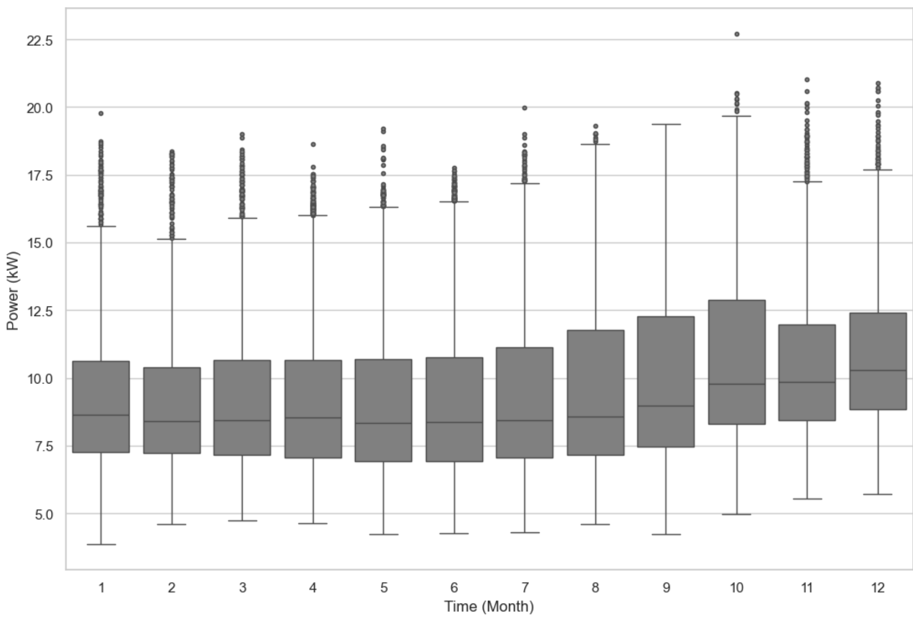

Fig. 2. Monthly Distribution of Electricity Demand Represented as a Boxplot with Outliers

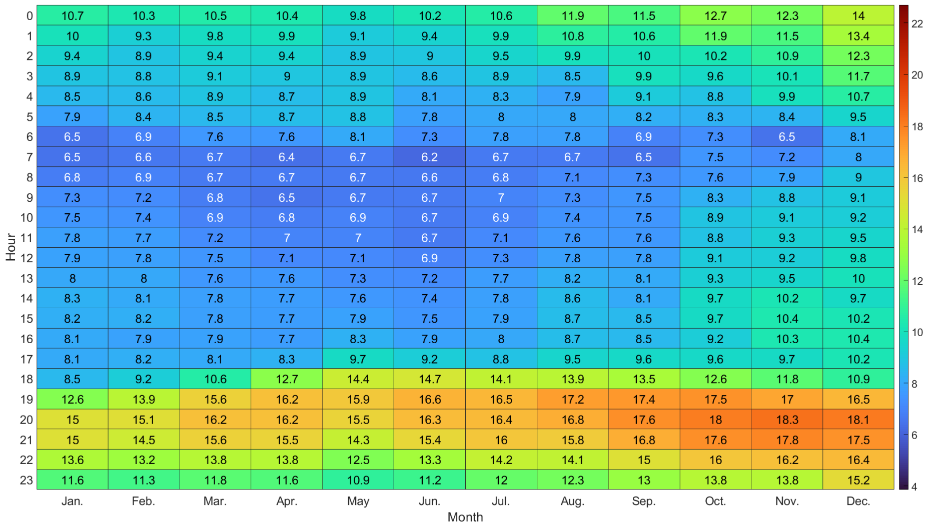

Fig. 3. Heatmap of the mean energy demand on weekends, showing the variation across different hours of the day and months of the year.

Figure 4 presents the double standard deviation of Figure 3, which is significant lower than the single standard deviation plotted in Figures 1 and 2. The conclusion of this comparison is that the training of the neural network needs to be trained by including as meta data the hours, the days and if it is a weekend day or not. A decrease in the standard deviation of the metadata-sensitive demand profiles means less uncertainty due to unknown, influencing variables and minimising the error of the generated demand profiles. In principle, applying a two-input-based heat map might be an interesting approach to evaluate whether metadata affects the error and to assess the relative impact of the two inputs on each other.

The original measurements were recorded every 5 minutes. Since the objective is to achieve hourly estimations suitable for pre-design analysis, all 5-minute measurements within each hour were averaged. Additionally, two other components were integrated alongside hourly energy demand: a variable was included to denote *weekends (1)* versus *weekdays (0)* to account for potential differences in demand patterns arising from social or economic activities, and the *month* was included to capture potential seasonal variations in energy usage.

The inclusion of *temperature* and *humidity* variables was essential to complement the energy demand records due to the nature of the used appliances, such as refrigerators and other temperature-sensitive devices. These appliances exhibit energy demand patterns that are heavily influenced by ambient temperature and humidity. For instance, higher temperatures increase the cooling demand of refrigerators, while humidity levels can affect their efficiency and operation cycles. Incorporating these meteorological variables allows the model to account for external factors that significantly impact en-

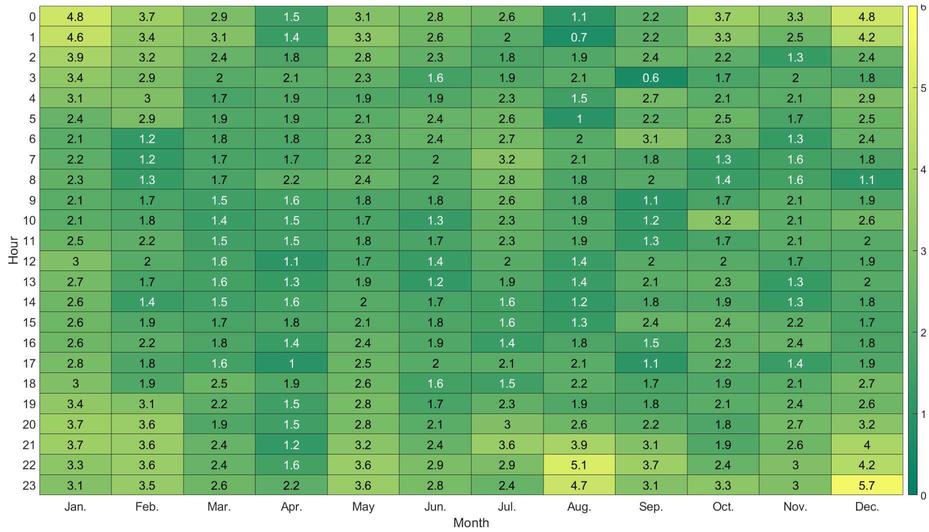

Fig. 4. Heatmap of double standard deviation energy demand on weekends, showing the variation across different hours of the day and months of the year.

ergy demand (Raza and Khosravi, 2015), thereby minimising the error and enhancing the reliability of the estimations. *Temperature* and *humidity* data were retrieved from (Weather Forecast API, 2023). These data were aligned by timestamps to ensure precise temporal synchronisation of all variables, producing a unified dataset for the El Espino case.

The resulting dataset provides the foundation for the modelling process described in the following section.

5 Architecture Definition

This work employs a neural network model designed to estimate hourly electricity demand using a set of metadata inputs: temperature (in degrees Celsius), humidity (in percent), hour of the day, month of the year, and a binary indicator for type of day (weekday or weekend). These variables were selected based on their broad availability in open-source datasets and their known influence on electricity demand patterns. The output node corresponds to the estimated hourly power in kW, meaning that for a given input (e.g., a specific hour and month along with the corresponding temperature and humidity), the model produces a single kilowatt value.

A key feature of the proposed approach is the use of a neural network architecture that transforms the input variables into a higher-dimensional space before making demand estimations. By expanding the input space, the model can capture complex dependencies and uncover latent causal structures that may not be evident in the original feature set. This higher-dimensional representation allows the network to move beyond simple correlations, enabling it to better model the intricate relationships that govern electricity demand.

The flexibility of neural network models in handling incomplete or irregular data has been highlighted in several studies (Owda et al., 2014; Hooshmand and Sharma, 2019). Unlike other methods that require complete daily data to generate estimations, this approach leverages any available measurement, allowing the construction of a dataset from incomplete or irregular records. It is possible to construct an entire demand curve for a selected time period by iterating this process across various hours.

The proposed architecture is represented in Figure 5. It incorporates a min-max normalization at the input stage, as shown in Equation 2, ensuring that all features are scaled to the $[0, 1]$ range. The minimum ($\min(x)$) and maximum ($\max(x)$) values for each feature are computed from the training dataset, ensuring consistency during inference. This normalization step is critical, as it prevents disproportionate influence from features with naturally larger magnitudes and generally improves model convergence and stability.

$$x_{\text{scaled}} = \frac{x - \min(x_{\text{train}})}{\max(x_{\text{train}}) - \min(x_{\text{train}})}. \tag{2}$$

The network consists of seven layers, structured to progressively increase the dimensionality of the feature space before refining the output. It begins with an input layer of 7 nodes, followed by hidden layers with 50, 250, and 750 nodes, capturing increasingly complex representations. The network then transitions through 300 and 150-node layers before reaching the single-node output layer. This progressive expansion in dimensionality enables the model to disentangle intricate dependencies, facilitating the capture of underlying causal relationships in the data.

Rectified Linear Unit (ReLU) activation functions are applied to each layer (unless otherwise specified) due to their effectiveness in modelling complex, non-linear relationships (Xu, 2015). Additionally, a batch normalization layer (indicated in orange) is incorporated after one of the hidden layers to stabilize and accelerate training. Finally, the red layer at the end represents a min-max descaling step, transforming the output from the normalised scale back to kilowatts, ensuring that estimations remain interpretable in domain-relevant units.

This hierarchical expansion and contraction of the feature space serves as a fundamental component of the model's ability to capture causal dependencies, as the higher-dimensional layers allow for richer representations before refining the output to a single predicted demand value.

The chosen architecture, normalisation strategies, and hyperparameters were guided by established best practices in neural network modelling and iterative empirical testing. Nonetheless, additional sensitivity analyses, alternative architectures (e.g., recurrent or attention-based networks), and more systematic hyperparameter optimisation could further enhance the model's estimation performance and transferability to different contexts.

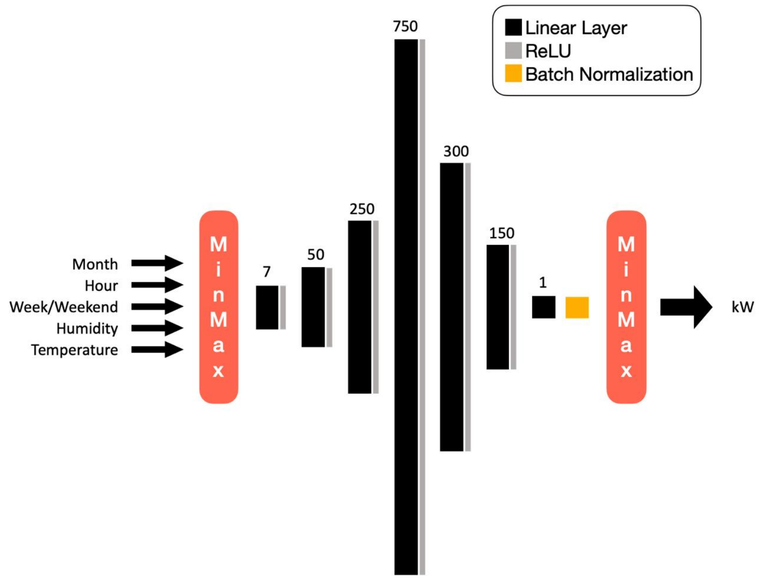

Fig. 5. Proposed neural network architecture, including input normalisation, hidden layers with ReLU activation, batch normalisation, and a final output descaling layer.

6 Model training

The dataset of 13,872 hourly measurements was randomly split into training (8,878 hours), validation (2,220 hours), and testing (2,774 hours) sets before training. This approximate 64%, 16%, 20% division is a standard practice aimed at ensuring robust model evaluation and preventing overfitting (Goodfellow et al., 2016). The training set was exclusively used to update model parameters, while the validation set was employed to monitor performance during training and prevent overfitting using early stopping techniques. The testing set was reserved for final evaluation, simulating real-world performance.

All weight parameters of the neural network were randomly initialised, a standard procedure in deep learning workflows. ReLU activation was uniformly applied to all nodes, as it avoids the vanishing gradient problem common in traditional sigmoid or tanh activations. MSE and MAE were used as the primary performance metric, given its widespread acceptance and straightforward interpretation (Hyndman and Koehler, 2006). Specifically, MAE quantifies the average magnitude of errors without considering their direction. In contrast, MSE assigns greater weight to larger errors due to its squared term, offering additional insight into the frequency and impact of significant discrepancies.

Four different optimisers were tested: Adaptive Moment Estimation (Adam), Adaptive Gradient (Adagrad), Adaptive Delta (Adadelta), and Stochastic Gradient Descent (SGD) (Tian et al., 2023), which are well-established and have demonstrated their applicability across a wide range of problems in neural network training. Each optimiser was applied with an initial learning rate of 0.01. Training was capped at a maximum of 20 epochs, and an early stopping mechanism was implemented to halt training once no further improvement on the validation set was observed, thus preventing overfitting.

7 Results

The primary objective of this study was to validate whether the selected metadata and input variables lead to an reliable approximation of electricity demand by capturing underlying causal relationships. By integrating key external factors—temperature, humidity, hour, month, and weekend indicator—the model aims to identify the causal influences that shape electricity demand patterns. Rather than relying on predefined demand curves or purely statistical correlations, this approach enhances interpretability and robustness in demand estimation.

To evaluate the model's performance, we analyse the impact of different optimisers on the error metrics. Table 1 presents the Mean Absolute Error (MAE) and Mean Squared Error (MSE) (Bhuyan et al., 2016) for each optimiser tested. The results indicate that the Adam optimiser achieves the lowest error values, with an MAE of 0.0536 and an MSE of 0.0511, making it the most effective optimiser for minimising errors.

To further validate the effectiveness of the proposed architecture, we compare its performance with a simpler model, the triangle-shaped architecture, as shown in Table 2. Across all tested optimisers, our approach consistently outperforms the simpler architecture in both MAE and MSE. Specifically, for the Adam optimiser, our architecture achieves an MAE of 0.0536 and an MSE of 0.0511, whereas the triangle-shaped architecture exhibits significantly higher errors, with an MAE of 0.1345 and an MSE of 0.1339. This results in a 60.2% reduction in MAE and a 61.8% reduction in MSE in our approach.

The performance gap is particularly evident across different optimisers. For Adagrad, our model's MAE is 0.0691, compared to 0.1818 in the triangle-shaped architecture, leading to a 62.0% improvement in MAE. Similarly, the MSE is reduced from 0.1744 in the simpler model to 0.0538 in our approach, which is a 69.2% reduction. Even for Adadelta, which produces the highest errors, our model achieves a lower MAE (0.1018 vs. 0.1364), resulting in a 25.3% improvement and MSE (0.1124 vs. 0.1962), which shows a 42.6% reduction, reinforcing the limitations of the simpler structure in capturing complex demand patterns. A similar trend is observed for the SGD optimiser.

These results confirm that our architecture enhances model precision and robustness, effectively capturing complex dependencies while minimising errors across different optimisation techniques. The substantial reduction in error compared to the triangle-shaped architecture supports the hypothesis that increasing model complexity, particularly through higher-dimensional encoding and carefully designed layers, leads to improved electricity demand forecasting.

Next, we conducted a sensitivity analysis using the Adam optimiser, as it performed the best with all the features, to determine the impact of excluding temporal features, such as the month of the year and weekend indicator, on the model's error. The results in Table 3 confirm that removing these variables leads to an increase in both MAE and

Optimiser	MAE	MSE
Adam	0.0536	0.0511
Adagrad	0.0691	0.0538
Adadelta	0.1018	0.1124
SGD	0.0904	0.1165

Table 1. Performance metrics of the proposed Neural Network Architecture.

Optimiser	MAE	MSE
Adam	0.1345	0.1339
Adagrad	0.1818	0.1744
Adadelta	0.1364	0.1962
SGD	0.1327	0.1403

Table 2. Performance metrics of the triangle-shaped architecture.

MSE, with a 10.7% increase in MAE and an 11.7% increase in MSE, further validating their relevance in electricity demand estimation.

8 Conclusions

Electricity demand estimation is crucial for effective microgrid design, particularly in remote areas where data availability is limited (Sanfilippo and et al., 2023). Unlike traditional methods that rely on predefined demand curves, the proposed approach introduces a causal model, implemented using a neural network, to uncover the underlying relationships driving electricity consumption. This study leveraged a dataset of 578 days from El Espino, Bolivia, integrating measured electricity demand with readily available metadata—such as temperature, humidity, hour, month, and whether the day is a weekday or weekend—to develop a novel estimation method. A key contribution of this work is the design of a neural network architecture specifically oriented toward capturing causality. Unlike conventional neural networks that primarily focus on pattern recognition, this architecture is structured to model causal relationships between input variables and electricity demand, ensuring that the learned representations reflect meaningful dependencies rather than surface-level correlations.

The study shows the superior performance of the causal-oriented neural network architecture over the triangle-shaped architecture in estimating hourly electricity demand. Our model consistently achieves lower errors, with a MAE of 0.0536 and an MSE of 0.0511, outperforming the triangle-shaped architecture by significant margins—up to 62% for MAE and 69% for MSE across various optimisers. The Adam optimiser, which outperforms Adagrad, Adadelta, and SGD in this context, enhances the model's ability to handle incomplete and noisy data. This approach not only offers a more reliable and

Feature Set	MAE	MSE
Full Feature Set	0.0536	0.0511
No Temporal Features	0.0593	0.0571

Table 3. Impact of Removing Temporal Features on Model Performance.

robust estimation but also provides greater interpretability, making it a valuable tool for complex demand pattern predictions. The results indicate that the neural network effectively captures causality between external factors and electricity demand. Additionally, the identified factors—such as temperature, humidity, time of day, month of the year, and weekend indicator-prove to be highly relevant in shaping electricity demand patterns. The sensitivity analysis, conducted using the Adam optimiser, further confirms that excluding these temporal-based features leads to an increase in error, demonstrating their importance in demand estimation. Specifically, removing the temporal variables results in a 10.7% increase in MAE and an 11.7% increase in MSE, underscoring their relevance. These results highlight that the selected variables not only contribute to minimising the error but also play a crucial role in capturing the causal relationships underlying electricity consumption, reinforcing the effectiveness of the proposed approach.

However, as datasets, methodologies, and computational tools continue to evolve, so too will the capacity to develop more scalable and widely applicable models. The integration of causality-driven neural networks into demand estimation holds significant potential for enhancing the reliability and efficiency of microgrid systems, particularly in underserved regions where precise energy planning is essential for sustainability and resilience.

While the chosen architecture—including layer sizes and the number of layers—was informed by preliminary experiments and computational constraints, future research could explore alternative neural network architectures to further improve the model's ability to capture causal relationships. Additionally, further work could investigate the impact of alternative input variables beyond those considered in this study. Incorporating additional environmental, socio-economic, or behavioural factors could enhance the model's ability to disentangle causal dependencies and improve demand estimation. Sensitivity analyses on a broader set of variables would help determine the most relevant causal factors for different microgrid contexts, ensuring the model remains adaptable and robust across diverse energy systems.

Although this approach was tested on a single region, its implications extend beyond El Espino. As data collection efforts improve—through expanded measurement campaigns, data-sharing platforms, or innovative sensing technologies—a similar methodology could be adapted for other remote areas, microgrids, renewable energy communities, and even individual buildings. Achieving broader applicability requires access to more extensive and diverse datasets, alongside ongoing methodological refinements. Future research should focus on validating the causal model across multiple sites and contexts to ensure robustness and transferability.

Acknowledgment

We acknowledge the European Institute for Energy Research (EIFER) for financing this internship, which contributed significantly to the development of this research.

References

Abugabbara, M., Javed, S., and Johansson, D. (2022). A simulation model for the design and analysis of district systems with simultaneous heating and cooling demands. *Energy*, 261:125245.

Balderrama Subieta, S. L. (2022). *Optimal design and deployment of isolated energy systems: The Bolivian pathway to 100 % rural electrification*. PhD thesis, ULiège - Université de Liège, Liège, Belgium.

Bhuyan, M. K., Mohapatra, D. P., and Sethi, S. (2016). Software reliability assessment using neural networks of computational intelligence based on software failure data. *Baltic J. Modern Computing*, 4(4):1016–1037.

Białek, J., Bujalski, W., Wojdan, K., Guzek, M., and Kurek, T. (2022). Dataset level explanation of heat demand forecasting ann with shap. *Energy*, 261:125075.

Castillo, V. Z., de Boer, H.-S., Muñoz, R. M., Gernaat, D. E., Benders, R., and van Vuuren, D. (2022). Future global electricity demand load curves. *Energy*, 258:124741.

Dieudonné, N. T., Armel, T. K. F., Vidal, A. K. C., and René, T. (2022). Prediction of electrical energy consumption in cameroon through econometric models. *Electric Power Systems Research*, 210:108102.

Dilaver, Z. and Hunt, L. C. (2011). Industrial electricity demand for turkey: A structural time series analysis. *Energy Economics*, 33(3):426–436.

Fioriti, D., Stevanato, N., Ducange, P., Marcelloni, F., Colombo, E., and Poli, D. (2023). Data platform guidelines and prototype for microgrids and energy access: Matching demand profiles and socio-economic data to foster project development. *IEEE Access*, 11:73218–73234.

Foldvik Eikeland, O., Bianchi, F. M., Apostoleris, H., Hansen, M., Chiou, Y.-C., and Chiesa, M. (2021). Predicting energy demand in semi-remote arctic locations. *Energies*, 14(4).

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

Gómez, M. and Rodríguez, J. C. (2019). Energy consumption and financial development in nafta countries, 1971–2015. *Applied Sciences*, 9(2).

Hooshmand, A. and Sharma, R. (2019). Energy predictive models with limited data using transfer learning. In *Proceedings of the tenth ACM international conference on future energy systems*, pages 12–16.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22(4):679–688.

Kandananond, K. (2011). Forecasting electricity demand in thailand with an artificial neural network approach. *Energies*, 4(8):1246–1257.

Lazzari, F., Mor, G., Cipriano, J., Gabaldon, E., Grillone, B., Chemisana, D., and Solsona, F. (2022). User behaviour models to forecast electricity consumption of residential customers based on smart metering data. *Energy Reports*, 8:3680–3691.

Lombardi, F., Balderrama, S., Quoilin, S., and Colombo, E. (2019). Generating high-resolution multi-energy load profiles for remote areas with an open-source stochastic model. *Energy*, 177:433–444.

Michalik, G., Khan, M., Bonwick, W., and Mielczarski, W. (1997). Structural modelling of energy demand in the residential sector: 1. development of structural models. *Energy*, 22(10):937–947.

Nasr, G., Badr, E., and Dibeh, G. (2000). Econometric modeling of electricity consumption in post-war lebanon. *Energy Economics*, 22(6):627–640.

Nayanathara, C. and Srilatha, R. (2018). Electrifying villages using microgrids. In *2018 International Conference on Power, Energy, Control and Transmission Systems (ICPECTS)*, pages 300–305.

Neervannan, A. (2018). Evaluating the effectiveness of deep reinforcement learning algorithms in a walking environment. *Baltic J. Modern Computing*, 6(4):335–348.

Owda, H. M., Omoniwa, B., Shahid, A. R., and Ziauddin, S. (2014). Using artificial neural network techniques for prediction of electric energy consumption. *arXiv preprint arXiv:1412.2186*.

Raza, M. Q. and Khosravi, A. (2015). A review on artificial intelligence based load demand forecasting techniques for smart grid and buildings. *Renewable and Sustainable Energy Reviews*, 50:1352–1372.

Sanfilippo, S. and et al. (2023). Microgrid design optimization in benin within the leopard project: Evaluating the impact of inaccurate load profile estimation. In *2023 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*.

Scarborough, D. and Somers, M. J. (2006). *Neural Networks in Organizational Research: Applying Pattern Recognition to the Analysis of Organizational Behavior*. American Psychological Association, Washington, DC.

Shin, S.-Y. and Woo, H.-G. (2022). Energy consumption forecasting in korea using machine learning algorithms. *Energies*, 15(13).

Shiraki, H., Nakamura, S., Ashina, S., and Honjo, K. (2016). Estimating the hourly electricity profile of japanese households – coupling of engineering and statistical methods. *Energy*, 114:478–491.

Somers, M. and Casal, J. (2009). Using artificial neural networks to model nonlinearity: The case of the job satisfaction–job performance relationship. *Organizational Research Methods - ORGAN RES METHODS*, 12:403–417.

Stadler, M., Cardoso, G., Mashayekh, S., Forget, T., DeForest, N., Agarwal, A., and Schönbein, A. (2016). Value streams in microgrids: A literature review. *Applied Energy*, 162:980–989.

Tian, S. and Chang, S. (2020). An agent-based model of household energy consumption. *Journal of Cleaner Production*, 242:118378.

Tian, Y., Zhang, Y., and Zhang, H. (2023). Recent advances in stochastic gradient descent in deep learning. *Mathematics*, 11(3).

Torabi Moghadam, S., Toniolo, J., Mutani, G., and Lombardi, P. (2018). A gis-statistical approach for assessing built environment energy use at urban scale. *Sustainable Cities and Society*, 37:70–84.

Velasquez, C. E., Zocatelli, M., Estanislau, F. B., and Castro, V. F. (2022). Analysis of time series models for brazilian electricity demand forecasting. *Energy*, 247:123483.

Wassie, Y. T. and Ahlgren, E. O. (2023). Determinants of electricity consumption from decentralized solar pv mini-grids in rural east africa: An econometric analysis. *Energy*, 274:127351.

Weather Forecast API (2023). Open-source weather api. https://open-meteo.com/.

Xia, Z., Ma, H., Saha, T. K., and Zhang, R. (2022). Consumption scenario-based probabilistic load forecasting of single household. *IEEE Transactions on Smart Grid*, 13(2):1075–1087.

Xu, B. (2015). Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*.

Zamanipour, B., Ghadaksaz, H., Keppo, I., and Saboohi, Y. (2023). Electricity supply and demand dynamics in iran considering climate change-induced stresses. *Energy*, 263:126118.

Received December 11, 2024 , revised February 21. 2025, accepted March 3, 2025