Dirección y Organización n. 89 (2026)
https://doi.org/10.37610/89.758

An Intelligent and Scalable IoT-Based Architecture for Residential Energy Consumption Forecasting Using Deep Learning

Predicción del consumo energético en el hogar mediante aprendizaje profundo: una arquitectura inteligente y escalable basada en IoT

Javier R. Caparrós[1], Felipe Romero[2], Elvira Maeso-González[3]*, Luis F. Romero[4]

Received: 06/11/2025 | Accepted: 18/03/2026

Abstract

This paper presents a smart, scalable architecture that integrates Internet of Things (IoT) technologies and Deep Learning models to improve the accuracy and adaptability of energy consumption forecasting in residential environments. The proposed system is designed to support efficient data acquisition, storage, and analysis in dynamic home contexts, where consumption is influenced by multiple temporal, environmental, and behavioral variables. The system's foundation is a comprehensive IoT architecture developed for robust data collection. This infrastructure includes high-precision sensors to monitor power consumption across three phases, environmental sensors to capture weather variables like temperature and humidity, and occupancy detection mechanisms that infer human presence through smart device activity. Furthermore, a dedicated Android application facilitates the calibration of household appliance energy usage, enabling the identification of specific devices contributing to consumption fluctuations. Data is transmitted in real-time using the low-bandwidth MQTT (Message Queuing Telemetry Transport) protocol, managed via RESTful API services, and stored in JSON format within a highly scalable MongoDB NoSQL database, chosen for its big data capabilities. The predictive core of the system is a sophisticated neural network that combines Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), specifically employing LSTM/GRU blocks, to effectively extract spatiotemporal patterns and capture long-term dependencies in timeseries data. The model architecture consists of six hidden layers with 2048 fully connected neurons each and is trained using the Mean Absolute Error (MAE) as the loss function with an Adam optimizer. The model incorporates a wide range of contextual factors, such as time of day, day of the week, holidays, weather conditions, and occupancy. Critically, it also explores user specific behavioral indicators, such as the presence of specific individuals, to achieve a more granular and accurate prediction. Evaluation was conducted on a real-world dataset collected over a two-month period, split into 80% for training and 20% for testing, with five-fold cross validation to prevent overfitting. The model demonstrated remarkable accuracy, achieving a Mean Absolute Error below 5% under diverse conditions. The system is designed for scalability, making it adaptable for larger applications such as residential communities or smart grid energy management. Future work will focus on enhancing model generalization by incorporating larger datasets over extended time frames and exploring the conversion of energy consumption data into images to further leverage the pattern recognition capabilities of CNNs.

Keywords: electricity consumption forecasting, Internet of Things (IoT), deep learning, smart homes, digital twins.

* Correspondence author.

[1] School of Industrial Engineering, University of Malaga, 29071, Spain. Email: jr@uma.es ORCID: 0000-0003-1621-476X

[2] School of Industrial Engineering, University of Malaga, 29071, Spain. Email: fr@uma.es ORCID: 0000-0001-8475-2725

[3] School of Industrial Engineering, University of Malaga, 29071, Spain. Email: emaeso@uma.es ORCID: 0000-0003-4870-4923

[4] School of Industrial Engineering, University of Malaga, 29071, Spain. Email: felipe@uma.es ORCID: 0000-0003-2959-2030

Resumen

Este estudio presenta un sistema basado en inteligencia artificial para predecir el consumo energético en hogares. Utilizando una arquitectura de Internet de las Cosas (IoT) para la recolección de datos y técnicas de Aprendizaje Profundo (Deep Learning), el sistema logra pronósticos de alta precisión. El modelo integra múltiples factores como la hora, el día, festivos, condiciones meteorológicas y la ocupación del hogar, incluyendo la presencia de individuos específicos. Los datos se almacenan en bases de datos NoSQL y se transmiten mediante el protocolo MQTT (Message Queuing Telemetry Transport). A pesar de usar un conjunto de datos limitado, el modelo demostró una notable precisión, mostrando su potencial para optimizar la eficiencia energética residencial.

Palabras clave: predicción del consumo eléctrico, internet de las cosas (IoT), aprendizaje profundo (deep learning), hogares inteligentes, gemelos digitales.

1. Introduction

The increasing demand for energy-efficient solutions in residential buildings has driven significant interest in intelligent systems capable of accurately forecasting electricity consumption. Such systems are essential for optimizing energy usage, reducing operational costs, and promoting sustainability. Recent advancements in digital technologies, particularly those involving the Internet of Things (IoT) and Deep Learning, have enabled the development of sophisticated prediction models that consider various contextual factors beyond historical consumption data.

Traditional methods for energy consumption forecasting can be broadly classified into four categories: statistical, engineering-based, data-driven, and hybrid approaches (Meira de Oliveira and Cyrino, 2018). Statistical models rely on historical trends, engineering models simulate physical systems, data-driven models utilize machine learning algorithms, and hybrid approaches combine these methods for improved performance. At the residential level, machine learning approaches have also been successfully applied to capture household usage patterns and improve short-term demand forecasts (Chou & Tran, 2018). Among these, Deep Learning, specifically Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), has shown significant promise due to its capacity to model complex nonlinear relationships in high-dimensional datasets (Koprinska et al., 2018; Nakip et al., 2023). In addition to recurrent and convolutional architectures, temporal convolutional networks have also shown strong performance for energy-related time-series forecasting, as demonstrated by Lara Benítez et al. (2020).

Several studies have explored different combinations of features to enhance forecasting accuracy. Some focus solely on past consumption patterns (Azad et al., 2023; Emshagin et al., 2022), while others incorporate external variables such as weather conditions, calendar events, or occupancy levels (Ding et al., 2019). In line with this idea, the DECODE framework shows that combining historical meter readings with environmental and contextual features can significantly enhance building-level energy consumption prediction (Mishra et al., 2023). However, few approaches provide a comprehensive integration of heterogeneous data sources within a single predictive architecture. Our research addresses this gap by designing an intelligent system that incorporates not only standard contextual variables, such as temperature, humidity, and time-related features, but also fine-grained behavioral indicators, like the presence of specific individuals in the household.

This work builds upon the current body of literature and proposes an architecture that leverages IoT technologies for real-time data acquisition, efficient transmission using MQTT, and scalable storage via MongoDB NoSQL databases. The predictive core employs a hybrid neural model that combines the spatial capabilities of CNNs with the temporal memory of RNNs (specifically LSTM and GRU layers). The model has been trained and validated with real-world residential data, achieving high accuracy and demonstrating robust performance under varying domestic conditions. This approach aligns with recent trends in intelligent infrastructure design within industrial contexts, where the integration of IoT and AI has been shown to significantly enhance predictive capabilities and operational efficiency (Mateo-Casalí et al., 2025).

The following sections detail the factors considered in the prediction model (Section 2), the design of the IoT-based data acquisition and storage system (Section 3), the structure and training of the Deep Learning model (Section 4), and the experimental results and conclusions (Section 5).

2. External parameters included in the prediction model

A robust energy consumption forecasting model must account for a wide range of variables that influence household electricity demand. In this work, we adopt a multivariable approach that includes both intrinsic and extrinsic factors, which together provide a comprehensive view of the consumption behavior in residential settings. External parameters incorporated into predictive models include:

Seasonality: Time of day, day of week, and seasons significantly affect energy consumption. For example, heating and air conditioning use varies with the seasons, while the use of certain appliances may follow daily or weekly patterns.

Weather Conditions: Variables such as outdoor temperature, humidity, and solar radiation influence heating and cooling demand. Incorporating weather data into predictive models allows estimates to be adjusted based on expected weather conditions.

Occupancy and Inhabitant Behavior: The presence of people in the home, especially the members of the family, each specified individually, and their energy consumption habits, such as the use of appliances and lighting systems, directly impact consumption. Presence detection through mobile phone connections and modeling user behavior using IoT sensors are key components to improve the accuracy of predictions.

Special Events and Holidays: During holidays or special events, consumption patterns may deviate from the norm due to changes in occupancy and activities performed.

The selected input features in our model include commonly accepted factors such as temperature, special events, and occupancy levels, which have been validated in previous studies. These variables are known to significantly influence domestic energy consumption and are widely used in predictive modeling. Additionally, our framework explores user-specific behavioral indicators, such as the presence of specific individuals in the household, which, while not commonly addressed in the literature, are included based on empirical intuition and domain-specific reasoning. We hypothesize that these fine-grained behavioral patterns may improve model performance under certain conditions, and we include them as a basis for future exploration. Their integration into predictive models is facilitated by the use of IoT architectures, which allow for the real-time collection of relevant data, providing a constant stream of information to feed Deep Learning models (Koprinska et al., 2018).

In addition, storing this data in NoSQL databases allows for efficient handling of large volumes of information, while GPU processing speeds up the training and execution of predictive models. Summarizing, the combination of IoT and Deep Learning in predicting residential energy consumption represents a significant step towards efficient energy management in homes. By considering a wide range of factors that influence consumption, these systems can provide more accurate estimates and offer personalized recommendations to optimize energy use, thus contributing to sustainability and cost reduction for users.

3. IoT-based data acquisition and storage system

The success of any energy consumption prediction model largely depends on the quality, granularity, and timeliness of the input data. To address these requirements, we developed a custom IoT-based infrastructure tailored to the needs of a residential environment, capable of capturing both electrical and contextual data in real time. Figure 1 illustrates the overall data-feeding architecture of the proposed neural network, which can be conceptually divided into four main functional blocks: (i) data acquisition components, (ii) communication and integration mechanisms, (iii) data storage and preprocessing infrastructure, and (iv) the data flow sequence feeding the predictive model. The following subsections describe these elements in detail.

Figure 1. An outline of the data feeding system to the NN.

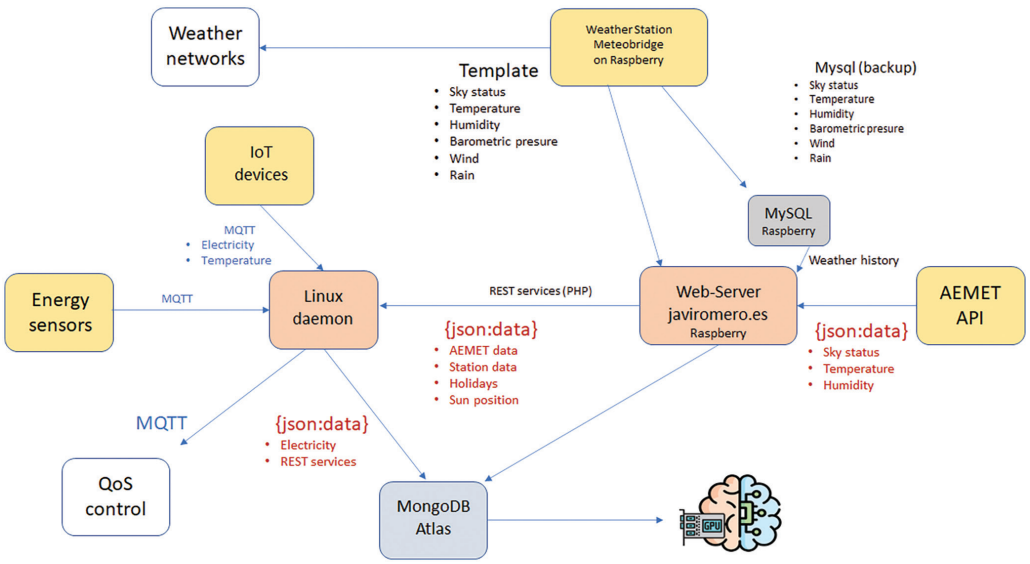

The main components include:

- Power and Energy Sensors: High-precision sensors are installed in the household electrical system to measure power consumption in three phase installations. to measure power consumption in three-phase installations, using specific hardware such as the Shelly™ 3EM sensor. These sensors continuously monitor voltage, current, and power factor to provide a detailed breakdown of energy usage per phase.
- Environmental Sensors: The system integrates meteorological sensors that capture temperature, humidity, and atmospheric pressure. These factors are crucial for understanding variations in energy consumption due to weather conditions.
- Occupancy Detection: Human presence in the household is inferred through indirect indicators such as smart device activity and network connectivity monitoring. These data points help estimate occupancy patterns and their influence on energy consumption.
- Smart Appliance Monitoring: An Android application is developed to calibrate the power consumption of household appliances. This calibration allows the identification of specific devices contributing to energy fluctuations, enabling better predictive modeling and potential user feedback mechanisms. The application publishes calibration data via MQTT, which is then captured by a central service to label power spikes in the database accordingly.
- Electricity Price Forecasting Subsystem: This subsystem predicts electricity price fluctuations using a deep learning model based on RNNs. The architecture includes Long Short-Term Memory (LSTM) layers to capture temporal dependencies and fully connected layers to refine predictions, ensuring optimized energy consumption strategies.

3.1. Communication protocols and data integration

The system utilizes efficient communication protocols to ensure seamless data transmission between sensors, servers, and applications:

- MQTT (Message Queuing Telemetry Transport): A lightweight, publish-subscribe messaging protocol is used for real-time data transmission from IoT devices to a centralized server. MQTT ensures minimal bandwidth consumption and efficient data delivery.
- RESTful API Services: The system integrates RESTful APIs for structured and scalable communication between different components. These APIs allow external applications, such as mobile apps and analytics platforms, to interact with stored data.
- JSON: Data is structured in JSON format for storage and integration with bigdata frameworks.

Data is transmitted in real-time using the low-bandwidth MQTT protocol. This data is captured by a custom C++ background service (known as *MQTT'daemon*), which subscribes to sensor topics and persists the information into the database. Concurrently, environmental and contextual data (like weather, holidays, or occupancy) is generated and served by a custom RESTful API built in PHP, which is also consumed by the main daemon.

3.2. Data storage, preprocessing and management

A robust storage infrastructure is implemented to manage the vast amount of collected data efficiently, including:

MongoDB (NoSQL Database): A highly scalable NoSQL database is chosen because of its ability to handle large, unstructured datasets efficiently and its horizontal scalability with flexible schema design.

Data Preprocessing: The system applies preprocessing techniques such as normalization, outlier detection, and segmentation. These steps ensure data integrity and improve the accuracy of predictive models.

Big Data Scalability: The architecture is designed to support future expansion beyond single households, enabling potential deployment in residential communities or energy distribution networks.

The integration of high-performance IoT sensors, efficient communication protocols, and a scalable data storage solution ensures a comprehensive and robust system for energy consumption analysis. This infrastructure provides a solid foundation for the deep learning models used in forecasting, ultimately improving energy efficiency and resource management in smart homes.

3.3. Data flow and sequence description

As depicted in Figure 1, the data-feeding process follows a sequential pipeline. Electrical and environmental data are first acquired through IoT sensors and smart devices, then transmitted in real time via MQTT and RESTful services to the central server. Once stored in the NoSQL database, the data are preprocessed and structured as time-series inputs, which are subsequently fed into the hybrid CNN–RNN model for training and inference. This pipeline ensures consistency and traceability throughout the data lifecycle.

4. Structure and training of the deep learning model

The predictive core of the proposed architecture relies on a hybrid Deep Learning model that combines Convolutional Neural Networks (CNNs) with Recurrent Neural Networks

(RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) layers. This hybrid configuration leverages CNNs for extracting local patterns and correlations from input sequences and RNNs for modeling the temporal dependencies across time steps. Recent studies have also demonstrated the effectiveness of image-based approaches using CNNs for time series forecasting, offering a complementary perspective on sequence modeling (Semenoglou et al., 2023).

Figure 2 illustrates the hybrid CNN-RNN architecture adopted in this study, highlighting the main processing blocks and the flow of information from input features to energy consumption prediction.

As shown in Figure 2, the convolutional blocks are responsible for extracting local patterns from the input sequences, while the recurrent layers model temporal dependencies across time steps.

The architecture is based on a combination of RNNs and CNNs, optimized for handling time-series data efficiently, including i) an input layer with a number of neurons equal to the number of features; ii) six hidden Layers with 2048 fully connected neurons per layer (the activation function is ReLU and the architecture combines CNN and RNN (LSTM/GRU) blocks for extracting spatiotemporal patterns and capturing long-term dependencies); and iii) an output Layer with 3 fully connected neurons and no activation function.

The input to the model includes the set of variables described in Section 2, such as temporal features, environmental conditions, occupancy status, and manually logged behavioral events. These inputs are normalized using min-max normalization to the [0,1] range and formatted as time-series sequences with a fixed-length sliding window to ensure consistency across training samples.

The dataset was divided into training and testing subsets using an 80/20 split. A five-fold cross-validation technique was applied to evaluate model robustness and prevent overfitting. The model was trained using the Adam optimizer, with Mean Absolute Error (MAE) as the primary loss function. Early stopping and dropout layers were incorporated to enhance generalization performance.

The dataset was collected from a single residential dwelling equipped with a three-phase electrical installation and a comprehensive set of IoT sensors. The residence corresponds to a typical private household, including standard domestic appliances and variable occupancy patterns. Data were collected continuously over a period of approximately two years. 80% of randomly selected observations were used for training, while the remaining 20% were reserved for testing and validation.

In addition to MAE, other metrics such as Root Mean Squared Error (RMSE) and R-squared (R^2) were calculated to assess predictive accuracy. The hybrid model consistently achieved MAE values below 5% across different validation sets, confirming its effectiveness in capturing the complex and nonlinear patterns in residential energy consumption.

The entire model was implemented in Python using TensorFlow and Keras libraries. The training process was executed on a standard workstation equipped with a GPU, ensuring reasonable training times and model scalability.

This hybrid modeling strategy, combining spatial and temporal learning, offers a flexible and robust solution for energy forecasting, particularly suited to environments with dynamic behavioral patterns and heterogeneous data inputs.

5. Results

The performance evaluation of the proposed hybrid Deep Learning model was carried out using real-world data collected over two years from a residential environment. The model was tested under various operational conditions, including regular weekdays, weekends, and special events such as holidays or user absences.

As stated above, the dataset was split into 80% for training and 20% for testing and cross validation was used to prevent overfitting. The results highlight the effectiveness of the deep learning approach employed.

The training process (shown in Figure 3) converged stably, with the training and validation curves remaining close to each other, indicating good generalization and the absence of significant overfitting.

In Figure 3, the horizontal axis represents the training epochs, while the vertical axis corresponds to the Huber loss value. The decreasing trend reflects model convergence,

Figure 2. Hybrid CNN-RNN architecture used for residential energy consumption forecasting.

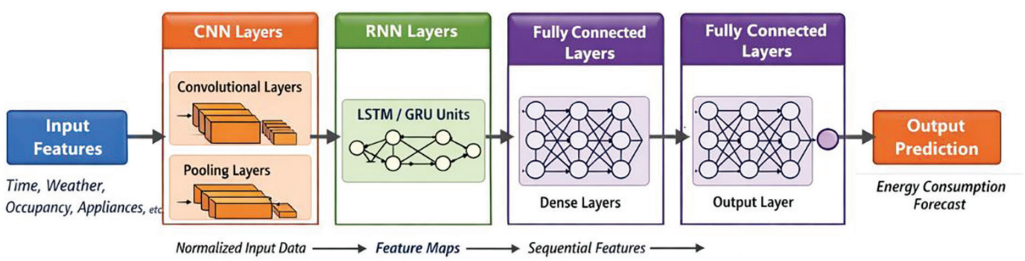

Figure 3. Evolution of the Huber loss during training and validation epochs.

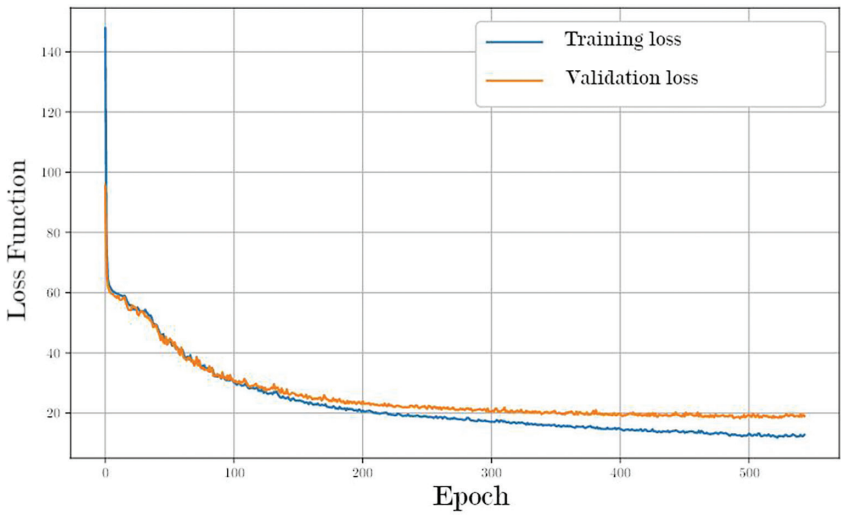

whereas the oscillations are associated with batch-based optimization and validation dynamics.

The proposed system demonstrates excellent computational efficiency for practical deployment. Inference testing on the trained model shows an average prediction time of 0.054 milliseconds per sample, with a throughput of approximately 469,000 samples per second. For a typical hourly forecast (60 samples), the total prediction time is only 0.128 milliseconds, enabling near real-time operational decisions at the household level. The model architecture comprises 11,267 trainable parameters with a memory footprint of 0.04 MB, making it suitable for deployment on edge computing devices. The training process, conducted on an NVIDIA GeForce RTX 3080 GPU (10GB VRAM), achieved convergence in approximately 2-10 minutes for 20 epochs with a batch size of 128, depending on the dataset size. This computational efficiency supports the system's applicability in real-world residential energy management scenarios, where both offline training and online inference must be performed reliably and with minimal latency.

The model was trained using the Mean Absolute Error (MAE) as both the loss function and evaluation metric. Optimization, performed with the Adam algorithm, and early stopping with a patience of 10 epochs, is applied to prevent overfitting. Figure 4 shows the "everyone at home" and "vacation" scenarios, as they represent negative extremes relevant to assessing model robustness. The model demonstrated strong forecasting capabilities, achieving a Mean Absolute Error (MAE) below 5%, even under varying conditions like weather and occupancy. Including temporal features such as weekdays and holidays significantly boosted accuracy by capturing daily and seasonal trends. Environmental inputs like temperature and humidity further improved adaptability to climate-driven changes. Designed for scalability, the system can be extended to larger residential settings and smart grid applications, maintaining consistent performance across different household types.

The model's robustness and adaptability are further demonstrated in Figure 5, which compares the actual and predicted power on August 5, 2025. This result is particularly significant as solar panels were added to the installation in June 2025. As observed in the "Total power" plot, the system learned to correctly model the photovoltaic generation (indicated by negative power values) during central daylight hours in just two months, a pattern drastically different from the original training data. This validates the model's high capacity to adapt to fundamental changes in the domestic environment. This research aligns with broader digital transformation trends in both residential and industrial contexts. As highlighted by Galaz de Miguel et al. (2025), the integration of AI and IoT technologies is becoming increasingly prevalent in organizational innovation strategies. The proposed system exemplifies how such integration can be effectively implemented to enhance energy efficiency, user awareness, and predictive capabilities.

6. Conclusions and future research lines

This work successfully developed a deep learning-based system for forecasting energy consumption in residential settings. The integration of IoT sensors and an advanced neural network model resulted in precise and adaptive energy predictions. The proposed architecture allows users to anticipate their energy consumption under different personalized scenarios, facilitating planning and savings. Unlike static approaches, our model adapts its predictions to real and changing conditions in the home environment.

Figure 4. Power prediction under two assumptions, and the actual power consumed.

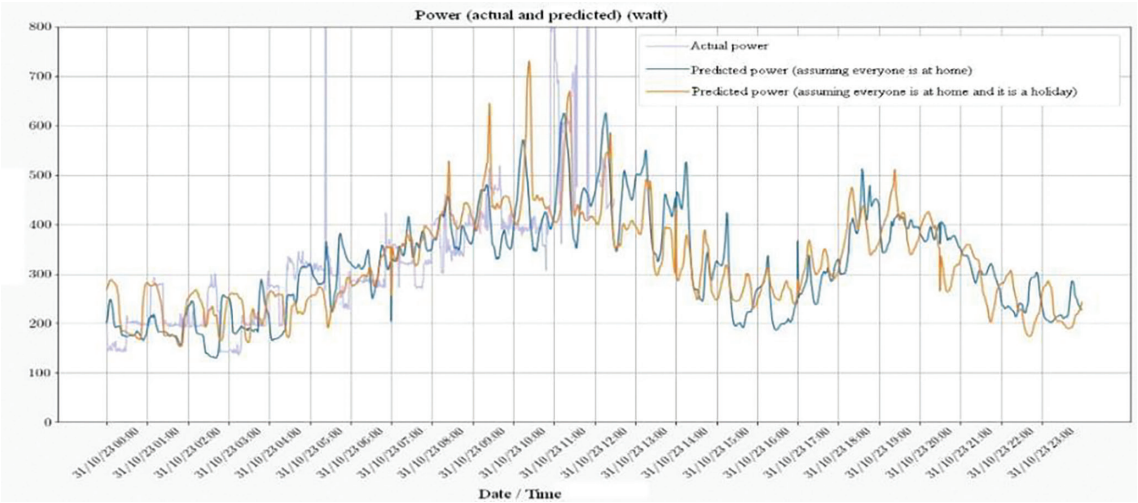

Figure 5. Comparison of actual and predicted power (total and per-phase) for 08/05/2025, showing RMSE values after the installation of solar panels.

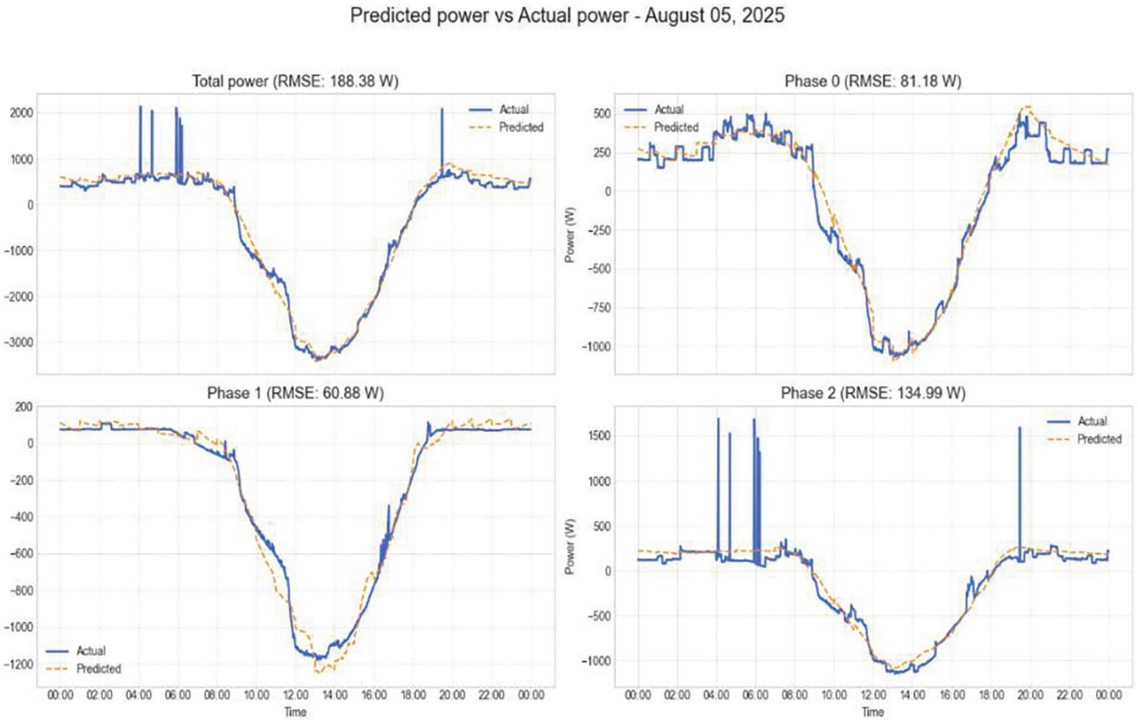

The system's scalability makes it suitable for broader energy management applications, including demand response programs and smart grid optimization.

7. Managerial implications

The proposed forecasting architecture supports decision-making in residential energy management, particularly in situations characterized by short-term variability and changing household behavior. At the household level, energy consumption forecasts can be used to support operational decisions such as scheduling high-consumption activities, anticipating demand peaks, or assessing the impact of changes in occupancy patterns (e. g., vacations or teleworking days). The system allows users to compare alternative personalized scenarios before decisions are implemented.

For energy service providers and aggregators, household-level scenario-based forecasts provide a more detailed representation of demand dynamics than aggregate consumption profiles. This information can support decisions

related to the design of demand-response programs, the identification of relevant user segments, or the evaluation of how contextual and behavioral factors affect short-term demand patterns.

In residential contexts with coordinated energy management, such as multi-dwelling buildings or housing complexes, the forecasting results can inform decisions related to energy efficiency measures, technology upgrades, or the integration of distributed energy resources. Scenarios can be defined by modifying contextual variables already included in the model, such as weather conditions, occupancy levels, or the presence of photovoltaic generation, enabling an ex-ante assessment of alternative operating strategies.

Overall, the managerial contribution of the proposed system lies in its ability to support scenario-based planning in residential energy contexts, providing decision-makers with forecasts that explicitly account for contextual and behavioral variability rather than relying on static consumption assumptions.

8. Limitations and future research lines

This study presents several limitations that should be considered when interpreting the results. First, the empirical analysis is based on data collected from a single residential environment, which may limit the generalizability of the findings to other household types, dwelling configurations, or geographic contexts. Second, the dataset covers a relatively short observation period, which constrains the analysis of long-term consumption patterns and seasonal effects. Third, although the proposed model incorporates a broad set of contextual and behavioral variables, other relevant factors, such as dynamic electricity pricing schemes or additional socio-demographic characteristics, are not explicitly considered.

Building on these limitations, future research will focus on enhancing model generalization by incorporating additional real-world datasets collected from multiple residential units and over extended time periods. This will allow a more robust assessment of model performance across different household profiles and seasonal conditions. Further work will also explore training strategies aimed at improving scalability and reducing computational costs, including the use of transfer learning techniques and alternative data representations. In addition, the integration of external signals such as photovoltaic generation profiles and electricity price information will be investigated to extend the applicability of the proposed system to more advanced energy planning and optimization scenarios.

Finally, the robust data architecture developed in this study provides the basis for a future digital twin of the residential environment, in line with recent approaches proposed in the literature (Semenoglou et al., 2023). Such a digital representation could be used to test hypothetical scenarios by modifying selected model inputs, such as occupancy patterns, appliance usage, weather conditions, or photovoltaic generation, and by observing their expected effects on energy consumption. This would allow decision-makers to compare alternative operating strategies before implementation. In addition, continuous updates with sensor data and short-term forecasts could support real-time decision-making, for example by enabling load rescheduling, better alignment between consumption and on-site generation, or responses to anticipated demand peaks. The system's scalable design supports its potential extension to larger residential communities and smart grid applications.

References

AZAD, M. I., RAJABI, R., & ESTEBSARI, A. (2023). *Non-intrusive load monitoring (NILM) using deep neural networks: A review*. *arXiv*. https://arxiv.org/abs/2306.05017

CHOU, J. S., & TRAN, D. S. (2018). Forecasting energy consumption time series using machine learning techniques based on usage patterns of residential householders. *Energy, 165*, 709–726.

DING, Z., TURKCAN, M. K., & BOULANGER, A. (2019). *Using an ancillary neural network to capture weekends and holidays in an adjoint neural network architecture for intelligent building management*. *arXiv*. https://arxiv.org/abs/1902.06778

EMSHAGIN, S., HALIM, W. K., & KASHEF, R. (2022). *Short-term prediction of household electricity consumption using customized LSTM and GRU models*. *arXiv*. https://arxiv.org/abs/2212.08757

GALAZ DE MIGUEL, B., GUTIÉRREZ, J. A., & GARCÍA, R. (2025). Transformación digital en las empresas españolas: Análisis empírico del uso de tecnologías inteligentes. *Dirección y Organización, 85*, 55–65.

KOPRINSKA, I., WU, D., & WANG, Z. (2018). Convolutional neural networks for energy time series forecasting. En *2018 International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE.

LARA BENÍTEZ, P., CARRANZA GARCÍA, M., LUNA ROMERA, J. M., & RIQUELME, J. C. (2020). Temporal convolutional networks applied to energy-related time series forecasting. *Applied Sciences, 10*(7), 2322.

MATEO-CASALÍ, J., LABURU, C., BARANDIARAN, I., & RODRÍGUEZ, I. (2025). Smart maintenance for the circular economy: A digital twin proposal with AI-driven models for critical equipment in the food industry. *Dirección y Organización, 86*, 64–74.

MEIRA DE OLIVEIRA, E., & CYRINO, F. L. (2018). Forecasting mid-long term electric energy consumption through bagging ARIMA and exponential smoothing methods. *Energy, 144*, 776–788.

MISHRA, A., LONE, H. R., & MISHRA, A. (2023). *DECODE: Data driven energy consumption prediction leveraging historical data and environmental factors in buildings*. *arXiv*. https://arxiv.org/abs/2309.02908

NAKIP, M., ÇOPUR, O., BIYIK, E., & GÜZELİŞ, C. (2023). *Renewable energy management in smart home environment via forecast embedded scheduling based on recurrent trend predictive neural network*. *arXiv*. https://arxiv.org/abs/2307.01622

SEMENOGLOU, A., SPILIOTIS, E., & ASSIMAKOPOULOS, V. (2023). Image-based time series forecasting: A deep convolutional neural network approach. *Neural Networks, 157*, 39–53.