Development of a System for Prediction and Optimization of Electricity Consumption in Smart Homes, Based on Artificial Intelligence

Dejana Zorić[1], Goran Đukanović[2]
[1]*Banja Luka, Bosnia and Herzegovina, dejana.zoricc@gmail.com, ORCID ID: 0009-0006-7829-2981*
[2]*Banja Luka, Bosnia and Herzegovina, goran.z.djukanovic@apeiron-edu.eu, ORCID ID: 0009-0000-6828-5725*

Original scientific paper
https://doi.org/10.7251/JIT2601005Z UDC: 007.52:621.317.38

Abstract: This paper presents a machine-learning-based approach for short-term forecasting of household electricity consumption. The study aims to model temporal consumption patterns and support intelligent energy management in residential environments. Historical power consumption data were collected, cleaned, normalized and transformed into supervised learning sequences using sliding window techniques. A Long Short-Term Memory (LSTM) neural network was developed to capture time-dependent characteristics of electricity usage. The model was trained using the Adam optimization algorithm and evaluated using standard regression metrics, including Mean Absolute Error (MAE), which indicated high prediction accuracy and robustness. To ensure practical applicability, the proposed system integrates edge computing principles. Experimental results demonstrate that deep learning-based time-series forecasting can effectively predict short-term energy consumption. The proposed approach contributes to smart home energy monitoring by providing a scalable, efficient and reliable solution, and supports sustainable electricity usage through data-driven decision-making. The findings highlight the importance of integrating predictive analytics into future intelligent energy systems.

Keywords: LSTM, UCI, AI, Smart Homes

Introduction

Electricity consumption has been steadily increasing worldwide because of industrialization, urbanization and the growing reliance on electrical devices in everyday life. [1] This trend poses significant challenges for energy systems, particularly in the context of climate change mitigation and the transition towards sustainable energy use. As a result, energy efficiency has been recognized as one of the most effective instruments for reducing greenhouse gas emissions while maintaining economic and social development. In this framework, residential buildings represent a critical sector, as households account for a substantial share of total electricity consumption. [2]

Smart home technologies have emerged as a promising approach to addressing these challenges by enabling advanced monitoring, control and optimization of household energy usage. A central component of smart home systems is energy management, which aims to provide end users with insights into their consumption behavior and tools for improving efficiency. However, most existing energy management systems operate in a reactive manner, responding to current or past conditions without the capability to anticipate future demand. Such approaches limit the potential for proactive decision-making and optimal resource utilization.

Recent advances in artificial intelligence, particularly in deep learning, have introduced new possibilities for electricity consumption forecasting. [3] Deep learning models are capable of learning complex temporal dependencies from historical data, making them well suited for time-series prediction tasks. Among these models, Long Short-Term Memory (LSTM) neural networks have demonstrated strong performance in capturing long-term dependencies in sequential data and have been successfully applied in various energy-related forecasting problems. [4] Despite these advances, many existing studies remain focused on theoretical performance evaluation or cloud-based solutions, with limited attention given to practical deployment in real-world smart home environments.

A notable gap exists between demonstrated potential of artificial intelligence techniques in academic research and their integration into affordable, user-oriented energy management systems. Challenges related to system complexity, computational requirements and deployment constraints often hidden into translation of research prototypes into practical solutions. Furthermore, there is a lack of studies addressing the application of intelligent energy management systems in developing of under-researched regions, such as Bosnia and Herzegovina, where smart home adoption and energy efficiency research are still at an early stage.

This paper addresses these limitations by proposing a complete and practical system for short-term household electricity consumption forecasting based on deep learning techniques. The primary objective of the research is to design and implement an end-to-end solution that combines accurate prediction capabilities with practical deployment considerations for smart home environments. The proposed approach utilizes an LSTM-based neural network trained on historical consumption data to generate short-term forecasts that can support proactive energy management decisions.

The contribution of this work is threefold. First, it demonstrates the effectiveness of deep learning-based time-series forecasting for residential electricity consumption. Second, it presents a practical system architecture that supports deployment in resource-constrained smart home environments. Third, it provides insights into the applicability of intelligent energy management solutions in the local context of Bosnia and Herzegovina, contributing to the broader understanding of smart home technologies in emerging markets. The results of this research highlight the potential of integrating artificial intelligence into future energy management systems to support sustainable and efficient electricity consumption.

METHODS AND MATERIALS

For the development of the predictive model, the publicly available UCI Individual Household Electric Power Consumption was used. [5] The dataset contains detailed electricity consumption measurements from a single household located in a suburb of Paris, recorded over the period from December 2006 to November 2010. It comprises 2,075,259 individual observations collected at one-minute intervals, making it one of the most comprehensive publicly available datasets for residential energy consumption analysis.

Each observation includes nine attributes that provide a detailed representation of the household's electrical load profile. These attributes include the date and time of measurement, total active and reactive power, voltage, current intensity and energy consumption recorded by three sub-metering channels corresponding to the kitchen, laundry room and air-conditioning systems. Active power represents the actual electricity consumption of the household, while reactive power provides additional insight into power quality characteristics.

The dataset is provided in a text-based format with a semicolon as the field separator. Missing values are indicated by a question mark and were handled during the data preprocessing stage. The high temporal resolution and detailed sub-metering information enable precise analysis of consumption patterns as both aggregate and appliance-group levels, making the dataset particularly suitable for time-series forecasting and smart home energy management research.

The primary development environment used in this study was Visual Studio Code, which enabled integrated development of both the Python-based backend and the React Native frontend application.

The backend component was implemented using Python 3.11, with the Flask framework employed for the development of RESTful application programming interfaces. TensorFlow was used to implement the LSTM model for electricity consumption prediction, while NumPy, Pandas and Scikit-learn were utilized for data processing, preprocessing and statistical analysis.

The mobile frontend application was developed using React Native version 0.73, providing a cross-platform solution for user interaction. Communication with the backend services was handled using the Axios library for HTTP requests, while data visualization was implemented using react-native-chart-kit, enabling intuitive graphical representation of electricity consumption data and prediction results.

The complete data-processing pipeline consists of several sequential stages: raw data acquisition, preprocessing and cleaning, normalization, sliding-window sequence generation, LSTM-based prediction, threshold-based optimization logic and mobile application alert generation. This workflow enables transformation of raw household electricity measurements into actionable user recommendations for energy optimization.

The dataset loading process was implemented in the Python programming language using the Pandas library, which provides efficient tools for manipulating structured data. The loading routine was designed to automatically detect the presence of the UCI dataset on the local file system. In cases where the dataset was unavailable, the system generated simulated data that preserved the statistical and temporal characteristics of real household electricity consumption measurements. This approach ensured system robustness and independence from external data availability.

The dataset loading function first verified the existence of the required data files. If the files were present, a parsing process was initiated that involved several steps. The raw data were read from a text-based file with explicit specification of the field separator and the symbols representing missing values. Although Pandas supports automatic data type inference, additional parameters were defined to correctly handle special characters and non-standard numeric representations present in the dataset.

Following data import, the date and time attributes were combined into a single datetime object. This transformation enabled temporal aggregation and facilitated the analysis of seasonal and cyclical consumption patterns. Python's datetime objects provide extensive functionality for time-series manipulation, which was essential for subsequent preprocessing and modelling stages.

Subsequently, relevant columns were selected for further analysis. The focus of this study was placed on Global Active Power and the three sub-metering variables representing different appliance groups. Other variables, such as voltage and reactive power, while informative for advanced power quality analysis, were excluded as they were not directly required for the baseline predictive model.

Data cleaning represents a critical phase in preparing the dataset for machine learning applications. Despite its high quality, the UCI dataset contains a certain proportion of missing values resulting from measurement errors or communication interruptions in the monitoring equipment.

In the original dataset, missing values were denoted by a question mark. During the loading process, these values were identified and converted into NumPy NaN representations. A dropout-based strategy was then applied, removing all records containing at least one missing value. Although this approach resulted in the removal of a limited number of observations, it ensured data integrity and prevented the propagation of invalid values through the modeling pipeline.

An additional challenge involved the conversion of string-based values into numerical formats. While Pandas attempts automatic type conversion, the presence of special characters and varying decimal formats required explicit conversion using the pd.to_numeric function. The parameter ***errors='coerce'*** was applied to ensure that any non-convertible values were replaced with *NaN*, allowing consistent handling of conversion errors.

Following type conversion, an additional consistency check was performed to identify records containing negative consumption values. As negative electricity consumption is physically implausible, such records were interpreted as measurement errors and removed from the dataset. The outcome of this process was a clean and consistent dataset containing valid numerical values suitable for further preprocessing and predictive modelling.

To enable system testing in scenarios where the original UCI dataset was unavailable, an algorithm for generating synthetic electricity consumption data was developed. This approach ensured that the system could operate autonomously without external dependencies while preserving realistic consumption behaviour.

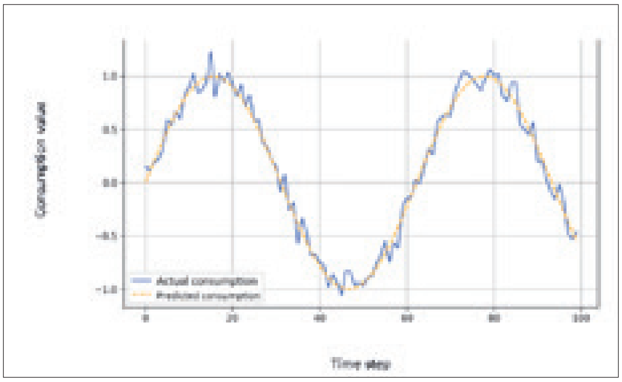

Figure 1. *Synthetic electricity consumption signal generated using sinusoidal daily patterns and Gaussian noise, compared with LSTM-predicted values to validate realism of generated data.*

The algorithm employed a combination of sinusoidal functions to model daily consumption cycles, with added Gaussian noise to simulate real-world variability. Kitchen energy consumption was modelled using a Gaussian curve centered around midday, reflecting typical appliance usage patterns. The amplitude and width of the curve were calibrated to match the average values observed in the UCI dataset (Figure 1).

Laundry room consumption was generated using a stochastic approach, in which discrete washing events were created at random intervals. These events were assigned fixed durations and amplitudes, simulating the intermittent and discrete nature of washing machine usage.

The third consumption group, encompassing water heating and air-conditioning systems, was generated using a combination of constant baseline consumption and a variable component influenced by daily temperature patterns. An evening peak was introduced to reflect increased hot water usage during typical household routines.

Global active power was calculated as the sum of all three sub-metering components, along with an additional baseline load representing other household appliances. The final signal was processed through a filtering mechanism to ensure the absence of negative values and to constrain all measurements within realistic operational ranges. This methodology produced synthetic data that statistically resembled the original UCI dataset, enabling reliable system testing under diverse operational scenarios.

An LSTM-based neural network was developed for short-term electricity consumption forecasting. The model consists of two stacked LSTM layers followed by fully connected layers. The first LSTM layer contains 64 units and processes input sequences of 24 hourly time steps, corresponding to one day of historical consumption. This layer is configured to return full sequences in order to preserve temporal dependencies for subsequent processing.

A dropout layer with a rate of 0.2 is applied to reduce overfitting. The second LSTM layer contains 32 units and outputs only the final hidden state, which represents a condensed temporal representation of the input sequence. This is followed by a dense layer with 16 neurons using ReLu activation. The final output layer consists of a single neuron with linear activation, producing a continuous-valued consumption prediction for the next hour.

The original time series was transformed into overlapping input-output pairs using a sliding window approach. Each input sample consists of 24 consecutive hourly values, while the target corresponds to the following hour. Prior to sequence construction, the data were normalised using Min-Max scaling to the range [0,1] to ensure stable training and faster convergence. The resulting input tensor has the shape (*samples, 24, 1*), which is the standard format for LSTM networks.

The model was trained using the Adam optimizer with a learning rate of 0.001. Mean Squared Error (MSE) was used as the loss function, while Mean Absolute Error (MAE) was monitored as an additional performance metric. The dataset was divided into training and validation subsets using an 80/20 split. Training was performed for 50 epochs with a batch size of 32.

The model converged rapidly, with a significant reduction in training loss observed during the first 20 epochs. After 50 epochs, the final training MSE reached 0.0234, while the validation MSE was 0.0267, indicating good generalisation and no significant overfitting (Figure 2).

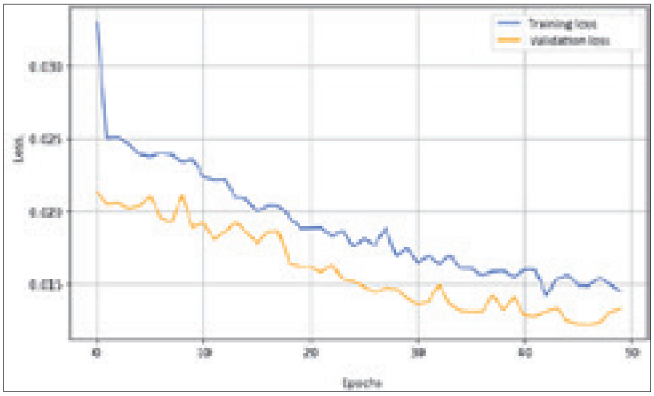

Figure 2. *Training and validation loss curves during LSTM model training, demonstrating stable convergence and absence of significant overfitting after 50 epochs.*

The final validation MAE of 0.094 kW corresponds to an average relative prediction error of approximately 4%, demonstrating high forecasting accuracy. The model performed best during stable consumption periods, while larger deviations were observed during rapid load changes, particularly during morning and evening peak hours. These results confirm suitability of LSTM networks for short-term household energy consumption prediction.

The system consists of a Flask REST API backend connected to an LSTM prediction model, and a React Native mobile application. This architecture enables real-time electricity consumption prediction, device management and visualization for smart homes.

The backend is implemented using Flask, selected for its simplicity, flexibility and seamless integration with Python's data science ecosystem. The server exposes multiple RESTful endpoints:

- /api/predict generates 24-hour electricity consumption forecasts using an autoregressive LSTM model.
- /api/devices supports CRUD operations for smart home devices and tracks their real-time status.
- /api/current-consumption provides real-time total household consumption.
- /api/statistics computes daily, weekly and monthly consumption statistics, including peak and low hours.
- /api/health returns server and model status for monitoring and diagnostics.

In-memory instances of the LSTM model, scaler objects, datasets and device logs are maintained to ensure fast response times. Robust error handling ensures the server can manage unexpected inputs without crashing.

The mobile application is developed using React Native, offering cross-platform support for Android and iOS. It communicates with Flask API to fetch predictions, device data, real-time consumption and statistics. Interactive charts visualize energy forecasts, while users can control smart home devices directly from the app.

Figure 3. *Main interface of the Smart Home Energy Saver application displaying real-time electricity consumption, 24-hour LSTM-based forecasts and peak consumption statistics for proactive energy monitoring.*

Figure 3 represents the Smart Home Energy Saver mobile application interface, which serves as the primary user interaction point for the system. The application displays real-time consumption data (0.935 kW), LSTM-based 24-hour consumption predictions with hourly granularity, and comprehensive statistics including peak usage (1.16 kW) and average consumption (1.15 kW). The predictive model successfully captures daily consumption patterns, enabling users to anticipate and optimize their energy usage.

The interface has been localized in Serbian for deployment in the target market.

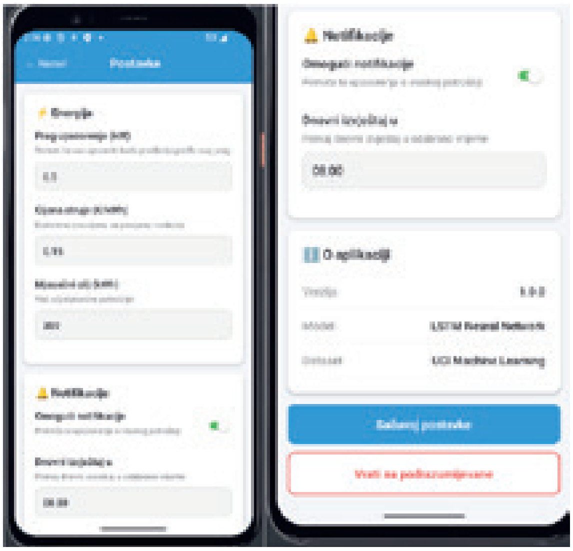

Figure 4 *Configuration interface allowing users to customize energy pricing, alert thresholds, reporting schedules and notification preferences within the smart home energy management system.*

The application includes a settings button that provides access to the configuration menu (Figure 4), where users can define key parameters such as electricity specifications, notification preferences and application information. Users can customize energy pricing (€/kWh), consumption thresholds (kWh), enable push notifications for excessive usage alerts and set daily reporting schedules. The interface also displays technical information including the LSTM Neural Network model and UCI Machine Learning dataset. This configuration capability enables personalization and adaptation to individual user requirements.

The interface uses a modern, card-based layout for clarity, displaying current consumption, predictive graphs, device controls and alerts.

The system continuously monitors predicted consumption against predefined thresholds. When a threshold is exceeded, the app generates an alert and recommends actions to optimize energy usage, allowing users to take immediate corrective measures.

RESULTS

The performance of the proposed LSTM-based prediction model was evaluated using standard regression metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE) and the coefficient of determination (R^2).

The model achieved an MAE of 0.0943 kW (94.3W). Considering the average household consumption of 2.38 kW, this corresponds to a relative error of approximately 3.96%, indicating high prediction accuracy in absolute terms.

The RMSE value of 0.1634 kW reflects the presence of occasional larger deviations, as expected due to the quadratic penalization of errors. The moderate difference between RMSE and MAE suggests that large prediction errors are not dominant and that the model remains stable over time.

The MAPE of 4.2% confirms strong performance across varying consumption levels. According to commonly accepted industry thresholds, MAPE values below 5% are considered excellent, [6] placing the proposed model in the highest accuracy category.

The coefficient of determination ($R^2 = 0.9534$) indicates that over 95% of the variance in energy consumption is explained by the model, demonstrating strong predictive capability.

To provide comparative context for evaluating the proposed LSTM architecture, its expected performance can be related to simpler forecasting approaches commonly used in time-series prediction tasks, such as ARIMA models and Feedforward Neural Networks (FNN).

Traditional ARIMA models are effective for modeling linear temporal dependencies but often struggle with highly nonlinear and dynamic household consumption patterns. Similarly, FNN architectures can model nonlinear relationships but lack internal temporal memory mechanisms required for sequential time-series learning.

In contrast, LSTM networks are specifically designed to capture long-term temporal dependencies and complex sequential patterns, which makes them more suitable for short-term electricity consumption forecasting in smart home environments.

The Flask-based backend was evaluated under local development conditions to access responsiveness and resource usage.

The health-check endpoint exhibited the lowest latency, responding within 3-5ms, as it performs minimal processing. The device status endpoint responded within 8-12ms on average.

The prediction endpoint, which includes data preprocessing, normalization, 24-step LSTM inference and result formatting required 35-50ms per request. Approximately 30ms of this time is attributed to LSTM inference. These response times are suitable for real-time interaction in a mobile application context.

The initial server startup time, including model loading into memory, was approximately 2.3 seconds. Once loaded, the model instance is reused, significantly reducing latency for subsequent requests.

Memory profiling showed that the LSTM model occupies approximately 8 MB of RAM, while the total Python process uses around 156 MB, which is acceptable for small-scale deployment.

The mobile application was tested using an Android emulator (Pixel 4, API 36). The average startup time was 1.8 seconds, providing a responsive user experience.

UI rendering performance remained stable at 60 FPS, ensuring smooth navigation and interaction. Memory consumption ranged between 85-95 MB, including application logic, data and graphical components.

Network usage was minimal. Initial data loading required approximately 8 KB, while periodic updates of real-time consumption (every 5 seconds) transferred 150-200 bytes per request. The 24-hour consumption graph required approximately 2 KB. Overall, daily network usage remained below 1 MB.

The system employs a proactive energy optimization strategy based on 24-hour consumption forecasts. A dynamic detection threshold is computed as the 90^{th} percentile of historical consumption, allowing the system to adapt to household-specific usage patterns.

When predicted consumption exceeds the threshold, the system generates a warning prior to the actual occurrence of high load. The maximum predicted value and its expected time are identified, enabling contextual recommendations (e.g., anticipated peak hours).

Device selection for optimization is based on three criteria: current power consumption, operational status and device type. Devices with the highest active consumption are prioritized, while heuristics ensure that critical appliances are deprioritized. This selection process is implemented via descending sorting based on real-time power usage.

Recommendations are presented through a clearly visible warning card containing: a quantitative problem description, peak consumption details and a concrete actionable recommendation.

User actions result in immediate feedback, including updated device status and real-time consumption changes, while preserving full user control over decision-making.

DISCUSSION

The achieved predictive performance confirms that the proposed LSTM model is well suited for household energy forecasting. With a Mean Absolute Error of 94.3 W and a Mean Absolute Percentage Error of 4.2%, the model demonstrates a high level of accuracy relative to the average household load. Such error levels are considered excellent in energy forecasting applications and indicate reliable short-term predictions.

The high coefficient of determination (R^2 = 0.9534) shows that the model captures the dominant temporal patterns in the data, explaining over 95% of the observed variance. The relationship between MAE and RMSE suggests that large prediction errors are infrequent and not systematic. From a practical standpoint, these results indicate that the model provides stable and consistent forecasts rather than occasional extreme deviations.

Importantly, the model's performance should be interpreted in relation to its intended purpose. The primary objective is not exact point-wise prediction, but timely identification of high-consumption period. In this context, the achieved accuracy is sufficient to support proactive energy management decisions.

The predictive model serves as an enabling component for higher-level optimization logic rather than an isolated forecasting tool. Since the system relies on dynamic, percentile-based thresholds, minor numerical prediction errors do not significantly affect the detection of high-load intervals. As long as relative consumption trends are preserved, the system can reliably anticipate critical periods.

The proactive nature of the approach is a key advantage. By generating warnings based on predicted values, the system allows users to act before peak consumption occurs. Even moderate forecasting inaccuracies do not undermine this functionality, as early notification remains beneficial compared to reactive responses.

Moreover, the optimization strategy prioritizes devices based on their relative contribution to total consumption. This ranking-based approach is inherently robust to small prediction errors and ensures that recommended actions consistently target the most impactful devices. As a result, the system translates predictive insights into practical, actionable recommendations with minimal cognitive effort required from the user.

From a system-level viewpoint, the integration of the predictive modeling, backend services and a mobile user interface demonstrates the feasibility of deploying machine learning-based energy optimization in an interactive application. Low response latency and modest computational requirements enable real-time interaction without compromising user experience.

The design intentionally abstracts technical complexity from the user. Instead of exposing forecasts or raw data, the system delivers concise warnings and concrete recommendations. This user-centric approach increases the likelihood of engagement and highlights the importance of coupling predictive accuracy with effective presentation and usability.

Compared to many commercial smart home solutions that focus primarily on automation or reactive monitoring, the proposed system emphasizes proactive decision support through forecasting. While existing platforms often provide historical insights or rule-based automation, they rarely integrate short-term prediction directly into user-facing optimization recommendations.

In contrast to many academic studies that remain limited to offline model evaluation, this work demonstrates an end-to-end system that bridges the gap between prediction, optimization and user interaction. The contribution lies not in outperforming all existing methods in terms of accuracy, but in demonstrating a practical and deployable integration of these components.

Despite encouraging results, several limitations should be acknowledged. The use of a single-household dataset restricts generalizability, and simulated device control does not fully capture the complexities of real-world IoT deployments. Additionally, the recursive prediction strategy may introduce cumulative errors over longer horizons.

These limitations suggest that while the results are promising, further validation on diverse datasets and real hardware integrations is necessary before large-scale deployment. Nevertheless, they do not invalidate the conclusions regarding system feasibility and design effectiveness.

The presented system provides a foundation for several future extensions, including adaptive thresholding, personalized recommendation strategies based on user behavior and integration with variable electricity tariffs. Incorporating real IoT devices and exploring multi-output forecasting models could further enhance robustness and practical impact.

Overall, this work demonstrates that predictive modelling can be effectively translated into actionable decision support for smart energy management, highlighting the importance of system-level design alongside model accuracy.

Conclusion

This study presented the development and evaluation of a functional system for short-term electricity consumption prediction and optimization in smart homes using LSTM neural networks. The central objective was to investigate whether an artificial intelligence-based system could provide accurate energy forecasts and translate them into personalized, actionable recommendations through an accessible mobile application. The results demonstrate that this objective was successfully achieved.

The proposed system integrates an LSTM forecasting model with a Flask-based backend and a React Native mobile application into a cohesive end-to-end solution. The predictive model achieved high accuracy at an hourly resolution, confirming the suitability of LSTM architectures for household energy forecasting. More importantly, the integration of prediction with optimization logic enabled a transition from reactive energy monitoring to proactive energy management.

Beyond predictive performance, the main contribution of this work lies in demonstrating the feasibility of deploying machine learning-based energy optimization without reliance on expensive hardware infrastructure. The system operates at a proof of concept showing that advanced artificial intelligence techniques can be embedded into lightweight, user-friendly applications suitable for real-world use. From a technical perspective, the achieved backend response times and stable mobile application performance indicate that such systems can support interactive and responsive user experiences.

Several limitations should be acknowledged. The use of a single-household dataset limits generalizability, and device control was simulated rather than implemented on real IoT hardware. Additionally, the recursive prediction strategy may lead to cumulative errors over longer forecasting horizons. These limitations suggest that further validation is required before large-scale deployment. It should be emphasized that the presented results are based on a single-household dataset and therefore primarily demonstrate proof-of-concept feasibility rather than universal generalization. Additional validation using heterogeneous multi-household datasets and real IoT deployments will be necessary to confirm scalability and applicability in broader smart home scenarios.

Future work will focus on integrating real smart devices using standardized IoT protocols, exploring alternative multi-output forecasting approaches and conducting long-term user studies to evaluate sustained effectiveness and user engagement. Overall, this work establishes a solid foundation for further research and practical applications in smart home energy management, demonstrating that predictive modeling combined with thoughtful system design can meaningfully contribute to improved energy efficiency.

Acknowledgements
The Faculty is acknowledged for its contribution to the academic framework within which this work was conducted. The author gratefully acknowledges the UCI Machine Learning Repository for making the dataset used in this study publicly available.

References

[1] I. E. A. (IEA), “”World Energy Outlook 2023”,” IEA Publications, Paris, 2023.
[2] I. E. A. (IEA), “”Energy Efficiency 2023”,” IEA Publications, Paris, 2023.
[3] Y. LeCun, Y. Bengio and G. Hinton, “”Deep learning”,” *Nature,* vol. 521, no. 7553, pp. 436-444, 2015.
[4] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation, 9(8),* pp. 1735-1780, 1997.
[5] G. Hebrail and A. Berard, “Individual household electric power consumption Data Set,” UCI Machine Learning Repository, 2012.
[6] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, Melbourne: OTexts, 2021.

Received: February 16, 2026
Accepted: April 30, 2026

About the authors

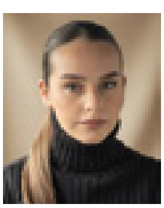

Dejana Zorić holds a Bachelor’s degree in Computer Science and Infromatics from the Faculty of Information Technologies, Pan-European Univerity Apeiron, Banja Luka. Her primary research interests lie in the fields of Artificial Intelligence, Internet of Things (IoT) and programming languages. She is especially interested in applying machine learning and deep learning methods within IoT systems and in designing efficient, scallable software solutions.

Goran Đukanović received his Ph.D from the Faculty of Electrical Engineering in Banja Luka. He has published more then fifty scientific papers and three university textbooks, and has gained additional scientific experience by making numerous reviews. He is full professor at the Pan-European University Apeiron (computer sciences) and member of the editorial board of the journal JITA. He is a member of the IEEE, with participation in Computer Society and the Internet of Things Community. His research interest includes mobile computing in IoT, with focus on resource management algorithms and event-driven mobile applications design.

For citation

Dejana Zorić, Goran Đukanović, Development of a System for Prediction and Optimization of Electricity Consumption in Smart Homes, Based on Artificial Intelligence, *JITA – Journal of Information Technology and Applications, Banja Luka*, Pan-Europien University APEIRON, Banja Luka, Republika Srpska, Bosna i Hercegovina, JITA 16(2026)1:5-13, (UDC: 007.52:621.317.38), (DOI: 10.7251/JIT2601005Z), Volume 16, Number 1, Banja Luka, June (1-76), ISSN 2232-9625 (print), ISSN 2233-0194 (online), UDC 004