

A Comparative Analysis of Univariate Time Series Models for Industrial Electricity Consumption Forecasting

Melike BİLDİRİCİ

Turkish Aerospace, Ankara, Türkiye,
Industrial Engineering Department,
Gazi University, Ankara, Türkiye
melikesarikaya98@gmail.com

Prof. Mehmet KABAK and Assoc. Prof. Ebru YÜKSEL HALİLOĞLU

Industrial Engineering Department
Gazi University, Ankara, Türkiye
mkabak@gazi.edu.tr, ebruyuksel@gazi.edu.tr

Time Series Analysis, Data Analytics, Energy Management

Abstract

Sustainability goals and compliance with ISO 50001 standards in industrial facilities require the development of high-accuracy energy forecasting models. However, operational dynamics on-site make it difficult to integrate external variables (e.g. production volume, weather conditions) into forecasting models in real time and without data loss. This study presents a practical, univariate forecasting framework that relies solely on historical electricity consumption data, without depending on complex external data. Daily electricity consumption data for 28 manufacturing and assembly buildings at a large-scale facility over a 6-year period (2020–2025) were dynamically aggregated based on their commissioning dates. During the data preprocessing stage, cumulative data gaps caused by communication interruptions were filled using the uniform distribution method, and outliers were corrected using IQR-based clipping. The predictive performance of ARIMA, Random Forest, XGBoost, and Long Short-Term Memory (LSTM) models was comparatively analyzed on the prepared dataset. The findings demonstrated that tree-based algorithms are superior in modeling sudden operational disruptions (such as shifts and holidays) in the industrial consumption profile. While ARIMA underperformed on nonlinear dynamics (83.04% test accuracy), LSTM—despite its stable performance (93.13%)—lagged behind other machine learning models on the tabulated dataset. Random Forest, supported by temporal feature engineering, emerged as the most successful predictor in the study, achieving a test accuracy rate of 94.10% and the lowest error metrics. Consequently, it has been confirmed that this proposed data-driven machine learning framework can be utilized as a highly reliable decision-support tool in energy budgeting, baseline creation, and planning processes for manufacturing facilities.

Keywords

Energy Forecasting, Time Series Analysis, Univariate Modeling, Machine Learning, Deep Learning

1. Introduction

In today's world, where global sustainability goals and carbon footprint reduction strategies have taken center stage, energy management in industrial facilities has become one of the most critical engineering challenges. In particular,

the effective implementation of energy management system standards such as ISO 50001 requires facilities to continuously monitor and improve their energy performance using data-driven approaches. At the heart of this process lies the ability to predict future energy consumption with high accuracy. A successful forecasting model enables the early detection of potential anomalies, the optimization of operational costs, and the implementation of proactive energy planning by preventing waste. Accurate forecasts serve as the cornerstone for managers to develop data-driven strategies rather than making reactive decisions.

Industrial electricity consumption has a nonlinear and highly complex structure due to the cyclical nature of production trends and temporal dependencies inherent within it. Traditional statistical methods (e.g., ARIMA) are often insufficient for capturing these dynamic and complex patterns. Given the fast-paced nature of industrial operations, there is a significant need for robust and practical models that can extract maximum predictive power from the facility's own historical consumption data—without always requiring complex external variables to feed into forecasting models. With the advancement of data collection infrastructure, the historical consumption data collected has made the use of modern machine learning algorithms both possible and necessary to solve this problem.

In time series analysis, decision tree-based ensemble algorithms (Random Forest, XGBoost) and deep learning architectures (LSTM) capable of learning long-term dependencies in sequential data are leading approaches for identifying complex patterns. The primary objective of this research is to develop ARIMA, Random Forest, XGBoost, and LSTM models using historical electricity consumption data from an industrial facility and to conduct a comparative analysis of their forecasting performance. The study will examine the algorithms' accuracy metrics (RMSE, MAE, etc.) and their behavior on industrial data. It is intended that the findings will guide energy managers and engineers in establishing a practical and reliable decision-support mechanism capable of operating with high accuracy without requiring complex external data. The primary research objectives of this study are designed to enable industrial facilities to make reliable energy forecasts without being dependent on multi-variable and complex data collection processes. The unique research contributions of this study to the literature and industrial applications are as follows:

Developing a Feasible Univariate Forecasting Framework for Industrial Facilities: While feeding forecasting models with external variables (production volume, shift, weather conditions, etc.) is a common approach in the theoretical literature, real-world industrial applications face serious data integration challenges, such as mismatches in the collection frequencies of these data, differences in start and end periods, and the difficulty of establishing meaningful relationships between variables. This study aims to overcome this data acquisition bottleneck in industrial settings by establishing a single-variable forecasting infrastructure that is data-efficient and much more feasible at the facility level, relying solely on historical electricity consumption data generated by the meter without depending on complex external data sources.

Comparative Analysis of Results Obtained Using Machine Learning and Deep Learning Architectures: Going beyond traditional statistical methods, directly comparing the success of tree-based ensemble algorithms (Random Forest and XGBoost) with a deep learning architecture specialized in time series (LSTM) in capturing industrial energy consumption patterns (using RMSE, MAE, and MAPE metrics).

Practical Decision-Support Mechanism and Compliance with Standards: Based on the findings, to provide energy managers with a straightforward, sustainable, and practical methodology they can use in the processes of reducing their carbon footprint and establishing an energy baseline in accordance with ISO 50001 standards.

The remaining sections of this study are organized as follows: Chapter 2 presents a comprehensive literature review of energy consumption forecasting approaches and highlights the research gap addressed by this study. Chapter 3 establishes the methodological framework of the study; it details the data definition and preprocessing steps, the machine learning and deep learning model architectures used, the restructuring of the time series using the sliding window approach, and the performance evaluation metrics used to compare the models. Chapter 4 covers the application of the developed univariate forecasting models to Manufacturing and Assembly buildings within an industrial facility and includes a comparative analysis of the obtained metrics. Finally, Chapter 5 summarizes the study's key findings, provides evaluations for industrial applications, and concludes the paper.

2. Literature Review

Forecasting energy consumption in industrial facilities is a topic that researchers have been studying for many years, and which has received extensive coverage in the literature. The complex nature of energy consumption profiles has driven researchers to continuously develop new methodologies capable of achieving higher accuracy rates. Studies on energy forecasting in the literature is generally divided into three main categories: traditional statistical methods, machine learning algorithms, and deep learning architectures. In a comprehensive systematic review on energy consumption forecasting conducted by Khalil et al. (2022), the evolution of these three fundamental approaches in the literature was examined, and the advantages of modern algorithms over traditional methods in handling complex data structures were detailed.

Early studies in the literature primarily focused on traditional statistical methods such as the Autoregressive Integrated Moving Average (ARIMA) and its variants. Mijic et al. (2024) examined the role of ARIMA-based models in general energy forecasting and noted that these methods are effective in modeling linear trends. However, Khan (2025) emphasized that traditional statistical methods fall short when dealing with the high-frequency and nonlinear patterns specific to electricity consumption. Supporting this limitation, Kachalla et al. (2025), in their examination of the limitations of traditional time series approaches, highlight the challenges faced by classical ARIMA models, particularly when dealing with nonlinear datasets. The study emphasizes that for the model's autoregressive structure to function effectively, the data must satisfy the stationarity assumption (constant variance and autocorrelation over time). Although ARIMA can eliminate linear trends by applying a differencing operation due to its integrated structure, the model's forecasting flexibility and performance remain limited in systems with complex and nonlinear dynamics.

To overcome these limitations of statistical methods, the focus in the literature has shifted toward decision tree-based machine learning algorithms. Amasyali and El-Gohary (2018) note that these algorithms use a hierarchical tree structure to map data samples to predictions; in this structure, the attributes of internal nodes represent the features, the branches represent the values these features take, and the leaf nodes represent the final prediction class. Zhu et al. (2024) have demonstrated that ensemble models such as Random Forest and XGBoost, built upon the flexible structure of decision trees that can scale with increasing training data, exhibit high performance in capturing input-output relationships in energy forecasting problems. The vast majority of current studies aiming to maximize prediction accuracy adopt multivariate approaches. For example, Lozano et al. (2023) successfully integrated external factors such as temperature, humidity, and production volume into machine learning models, thereby reducing error rates. This trend in the literature has increased the popularity of algorithms like Random Forest and XGBoost due to their ability to effectively handle a large number of external variables.

In recent years, however, with advancements in hardware capabilities, deep learning techniques have redefined standards, particularly in time series analysis. Oyewola et al. (2024) demonstrated that recurrent neural network (RNN) architectures, such as Long Short-Term Memory (LSTM), outperform traditional machine learning methods in energy demand forecasting due to their ability to learn both long-term and short-term temporal dependencies in sequential data.

The general trend in the literature regarding energy consumption forecasting is to improve accuracy by continuously adding new external variables to models. However, in the practical operations of buildings with different characteristics (production, warehouse, office, etc.) within industrial facilities, it is generally not possible to access clean, uninterrupted, and real-time external data. A recent systematic literature review by Geller et al. (2025) highlighted the success of deep learning models in electricity consumption forecasting; however, it noted that studies applying these powerful algorithms solely based on historical consumption data independent of external data still present a specific research need in the literature.

This study aims to contribute to filling this significant research gap by offering a practical alternative to the multivariate trend in the literature; it seeks to comparatively analyze the capacity of ARIMA, Random Forest, XGBoost, and LSTM algorithms to capture purely temporal patterns on univariate datasets from buildings with different dynamics within an industrial facility.

3. Methods

This section presents the methodological framework of forecasting models developed using univariate electricity consumption data from buildings with different characteristics within an industrial facility. The study details data preprocessing, the restructuring of time series for machine learning algorithms, the implementation of model architectures (Random Forest, XGBoost, and LSTM), and performance evaluation metrics.

3.1. Data Description and Preprocessing

The data used in the study were obtained from the central energy monitoring system (SCADA / energy analyzers) of an industrial facility operating in accordance with energy management system standards. The dataset consists of sequential, single-variable time series representing the historical electricity consumption of a total of 103 buildings located within the facility and categorized into 6 different functional categories (Social Facilities, Warehouse Buildings, Support Facilities, Office Buildings, Manufacturing and Assembly Buildings, Process and Test Buildings). The distribution of these buildings is as follows: 7 Social Facilities, 8 Warehouse Buildings, 14 Support Facilities, 16 Office Buildings, 28 Manufacturing and Assembly Buildings, and 30 Process and Test Buildings. However, since this study focuses on industrial production dynamics, the 28 Manufacturing and Assembly Buildings—which account for the largest share of the facility’s total electricity consumption and are the most energy-intensive—have been included in the analysis. The dataset for these selected buildings covers a broad 6-year period from January 1, 2020, to December 31, 2025. The time series frequency is defined as daily electricity consumption, and the measurement unit for the data is kilowatt-hours (kWh).

Data preprocessing steps were meticulously carried out to optimize model performance. It was observed that communication interruptions between the SCADA system and the meters resulted in missing data in the time series; on the first day the connection was restored, the total consumption for the missing days was reflected in the system as a single cumulative value.

At this stage, no data loss occurred in the actual consumption data, and no data was deleted. To maintain the system’s overall energy balance and avoid disrupting the sequential structure of the time series, a uniform distribution approach was applied at points where cumulative reading were recorded. This cumulative consumption value was divided by the sum of the number of days with missing data and the day the data was entered into the system ($n + 1$), and then distributed equally across all days in the relevant period. In this way, the gaps in the data for days with communication outages were filled using average consumption figures, ensuring data continuity for analysis and resulting in a lossless daily net consumption dataset.

To identify outliers, the Interquartile Range (IQR) method—a statistical technique—was employed, considering that energy consumption data typically do not follow a normal distribution and instead exhibit skewed distributions. According to this method, the lower (Q_1) and upper (Q_3) quartiles of the data were determined, and the IQR was calculated as $IQR = Q_3 - Q_1$. Values outside the range Q_1 and Q_3 were considered outliers. To avoid completely losing the dynamic structure of the time series and its operational load characteristics, these identified outliers were corrected using a clipping method rather than being removed. With this approach, values falling below the lower limit, Q_1 were directly set to this lower limit, while extreme load values exceeding the upper limit (Q_3) were set equal to the upper limit, thereby limiting the disruptive effect of these extremes on the model’s error functions.

3.2. Model Architectures

The study utilized four core algorithms representing statistical methods, machine learning algorithms, and deep learning architectures to learn patterns based on historical data without the need for external variables.

Autoregressive Integrated Moving Average (ARIMA): A traditional econometric method widely used in time series analysis, particularly effective in modeling linear relationships and trends. The ARIMA model consists of three main components: the autoregressive (AR) component, which is based on the series’ own past lagged values; the differencing operation applied to make the time series stationary; and the moving average (MA) component, which uses the weighted average of past forecast errors. In this study, the ARIMA model has been integrated into the system as a robust reference model to objectively compare the forecasting performance of nonlinear complex machine learning and deep learning models.

Random Forest (RF): As noted by Zolfaghari and Golabi (2021), Random Forest is an ensemble learning algorithm based on the Bagging method, consisting of a large number of decision trees. It creates different decision trees by selecting random subsets from the training data and performs the final prediction by averaging the results produced by these trees in regression problems. It is resistant to overfitting and is effective at reducing variance in consumption data.

Extreme Gradient Boosting (XGBoost): A powerful tree-based model that operates on the boosting principle. The trees are not built independently but sequentially; each new tree is trained to minimize the prediction errors made by the previous tree. Thanks to advanced regularization techniques, it can model complex energy consumption patterns with high accuracy and low computational cost.

Long Short-Term Memory (LSTM): A deep learning architecture developed as a solution to the problem that traditional Recurrent Neural Networks (RNNs) cannot learn long-term dependencies. LSTM cells use forget, input, and output gates to determine which historical data to retain or discard. This architecture offers the advantage of being able to process sequential data directly without requiring a state matrix.

For machine learning and deep learning models to learn complex patterns in a dataset without overfitting and to achieve optimal generalization, it is necessary to optimize the external model settings known as hyperparameters. The functions of the key hyperparameters of the models used in this study are described in detail in Table 1:

Table 1. Hyperparameters of the Models

Model	Parameters	Definition
ARIMA Parameters	Order (p, d, q)	This is the basic structure that defines the time series components of the model. p (Auto-Regressive - AR) specifies how many consecutive time steps (lag) from the past are included in the analysis to forecast current consumption. d (Integrated - I) indicates how many times the difference of the data itself is taken to make the series stationary. q (Moving Average - MA) dampens future deviations by determining the lag with which past forecast errors are integrated into the model.
Tree-Based Model Parameters (Random Forest & XGBoost)	n_estimators	The total number of independent decision trees to be created in the ensemble model. Increasing this number improves the model's prediction stability and variance-reducing effect but also increases computational cost.
	max_depth	Limits the maximum level at which each decision tree's nodes can branch. Optimizing this value is the most critical regularization step, preventing the tree from memorizing outlier noise in the training data.
	learning_rate	The weighting factor applied by successively built trees when correcting the error of the previous tree. Keeping this value low ensures the model reaches the global minimum more slowly but with much greater accuracy and robustness.
Deep Learning Parameters (LSTM)	units	The number of memory cells (neurons) in the hidden layers of the LSTM architecture. This number determines the network's ability to learn complex, nonlinear energy consumption relationships.
	dropout_rate	An effective thinning technique used to prevent the network from overfitting. It involves randomly deactivating a certain percentage (e.g., 20%) of the neurons in the network during the training phase. This prevents neurons from becoming overly dependent on one another and forces each neuron to perform independent feature extraction.
	epochs	This specifies the total number of complete passes through the entire training dataset via the feedforward and backpropagation optimization processes.
	batch_size	The size of the sub-set of data that the neural network processes before calculating the loss function and updating its weights. Processing data in

		batches ensures efficient use of hardware memory and enables the training process to proceed more stably.
--	--	---

3.3 Time Series Framing (Sliding Window Approach)

Machine learning algorithms such as Random Forest and XGBoost are inherently designed for supervised learning problems. To apply these algorithms to univariate time series, the data has been restructured using the “Sliding Window” technique. In this approach, the time steps from the past (lag features) are defined as the independent variable vector, while the consumption value at time t is defined as the dependent variable. Mathematically, the $\{y_{t-1}, y_{t-2}, \dots, y_{t-n}\}$ sequence is used as input to predict the consumption of a building at time t .

3.4. Evaluation Metrics

To objectively compare the prediction performance of the developed models across different building types, three fundamental statistical error metrics widely accepted in the literature were used: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). The equations representing the actual consumption value, the predicted value, and the total number of observations are provided below:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$$

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

$$MAPE = \frac{100}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|$$

4. Application and Discussion

This section presents the implementation steps, training and testing processes, hyperparameter optimizations, and comparative performance results of machine learning and deep learning models developed to predict energy consumption in manufacturing and assembly facilities. The industrial applicability of the models is discussed based on the fundamental principle that “the model with the lowest error is the best predictor.”

4.1. Implementation Details and Model Training

As part of the analysis, the daily electricity consumption data for 28 manufacturing and assembly buildings—which underwent the data preprocessing steps outlined in Section 3 (SCADA data interpolation and IQR outlier suppression)—were aggregated into a single total consumption time series to represent the facility’s overall production load. Considering the dynamics of physical growth and capacity expansion over time in industrial facilities, this aggregation process was performed dynamically, considering each building’s operational lifecycle (commissioning date). For each calendar day, the actual consumption of only those buildings actively generating data on that date was summed to obtain the facility’s total production load for that day. This approach ensured that the permanent capacity increases resulting from the commissioning of new buildings were learned by the prediction models as an organic pattern.

In accordance with the standard approaches accepted in the literature for time-series energy forecasting studies, the dataset was divided chronologically into two parts. The first 80% of the data was used as a training set for the models to learn past patterns, while the remaining 20% was used as a test set to measure the models’ performance on future periods they had never seen before. As noted in the study by Chinmayi et al. (2025) on machine learning algorithms, the optimal ratio for splitting training and test data falls within the range of 80:20 to 70:30 and allocating too little data to the training set can lead to overfitting.

Feature engineering was applied in addition to the sliding window dimension so that machine learning and deep learning algorithms could better understand industrial shift dynamics. In addition to 7-day backward consumption lags, the input matrix of the models includes calendar variables such as day, month, and day of the week, as well as binary indicators for the official Turkish holidays (is_holiday) and weekends (is_weekend).

4.2. Hyperparameter Optimization

Hyperparameter optimization was performed to ensure that each algorithm achieves high accuracy in accordance with its architectural structure and to prevent overfitting on the training data. Hyperparameters were determined through iterative trial-and-error processes and empirical performance observations. For example, in tree-based models, the tree depth was limited to improve generalization ability; in the LSTM model, memorization was prevented by randomly deactivating 20% of the neurons (Dropout). This approach is also directly consistent with the strategy implemented by Ait Lahoussine Ouali et al. (2026), which integrates a dropout layer before the fully connected layers in the LSTM architecture to reduce the risk of overfitting and improve generalization capacity. Additionally, an Early Stopping mechanism, which automatically halts training when the model's development stalls, was employed during LSTM training. The final hyperparameter configurations applied to all models are summarized in Table 2.

Table 2. Hyperparameters of the Applied Forecasting Models

Model	Hyperparameter	Value
ARIMA	Order (p, d, q)	(7, 1, 2)
Random Forest	n_estimators	200
	max_depth	15
	random_state	42
XGBoost	n_estimators	200
	learning_rate	0.05
	max_depth	6
	random_state	42
LSTM	seq_length (Lookback)	7
	units (layer 1, layer 2)	64, 32
	dropout_rate	0.2
	epochs	100
	batch_size	32

4.3. Experimental Results and Discussion

The fundamental principle in evaluating the performance of prediction models is to achieve the highest possible explanation rate by minimizing error metrics. Accordingly, the MAE, RMSE, and MAPE values of the models were calculated, and their success rates were determined. The reliability of a developed machine learning model is measured not only by its success on the training data but also by its consistency on test data it has never encountered before. To identify overfitting in the models and analyze whether they achieved “Good Fit,” the training and test performances are presented in a comparative manner in Table 3, while a visual comparison of the success rates obtained for manufacturing and assembly buildings is provided in Figure 1. When the numerical metrics in Table 3 and the graphical model behaviors in Figure 1 are examined holistically, the key findings that emerge are as follows:

Table 3. Train and Test Performance Metrics of the Models

Model	Dataset	MAE	RMSE	MAPE	Success Rate
ARIMA	Train	14761.53	23255.90	7.96%	92.04%
	Test	43715.88	51382.34	16.96%	83.04%
Random Forest	Train	4298.76	6843.65	2.35%	97.65%
	Test	15872.93	23314.59	5.90%	94.10%
XGBoost	Train	4448.77	6507.47	2.31%	97.69%
	Test	16852.46	24532.85	6.29%	93.71%
Multi-LSTM	Train	12636.82	19782.45	7.10%	92.90%
	Test	17466.96	24635.29	6.87%	93.13%




Figure 1. Model Comparisons

The Superiority of Tree-Based Models (RF and XGBoost): The Random Forest (94.10%) and XGBoost (93.71%) algorithms achieved the highest accuracy rates on the test set. The fact that the difference between the training and test performance of these models remained within the 3% to 4% range proves that the algorithms did not simply memorize the training data; rather, they successfully learned operational patterns—such as weekend drops, public holidays, and delayed production effects—to build a robust model architecture. Random Forest, which had the lowest MAE (15,872.93) and RMSE (23,314.59) values on the test set, was the most successful predictor in the study. The study conducted by Shwartz-Ziv and Armon (2022) supports the effectiveness of the Random Forest and XGBoost models obtained in this research. The authors note that, over the past decade, traditional machine learning methods have continued to outperform deep learning algorithms in modeling tabular data and have played a dominant role in the field.

Limitations of the Statistical Method (ARIMA): While the training accuracy of the ARIMA model was 92.04%, this rate dropped to 83.04% during the testing phase, and the MAE value (43,715.88) showed a significant increase. This situation demonstrates that traditional statistical methods, which can only capture linear relationships, are insufficient for forecasting the nonlinear, complex fluctuations in production and consumption at industrial facilities.

Stability of the LSTM Model and the Effect of Data Size: The fact that the test accuracy of the LSTM deep learning model (93.13%) is marginally higher than its training accuracy (92.90%) demonstrates that the Dropout mechanism integrated into the network architecture during training has been highly effective in suppressing overfitting and preventing memorization. However, despite the LSTM demonstrating stable performance, it lagged tree-based algorithms. Energy consumption profiles in industrial facilities do not exhibit continuous and smooth fluctuations; instead, they follow step functions characterized by sudden discontinuities such as shift changes, weekend shutdowns, and holidays. In a comprehensive study by Grinsztajn et al. (2022), it was demonstrated that while deep learning architectures excel with unstructured data such as images or text, tree-based models achieve the highest performance on tabular datasets containing numerical and categorical variables such as past consumption, day, month, and holidays. The fact that tree models partition the data space using logical rules (if-then-else) enabled them to isolate the sudden drops in consumption caused by the binary variables (weekend or holiday) in this study with much higher accuracy compared to LSTM's continuous functions. Furthermore, as emphasized in Ghannam's (2026) study, while deep learning models offer superior performance when processing large and high-frequency (minute/hourly) datasets, this advantage diminishes in smaller and annual datasets. The 6-year daily dataset from the industrial facility used in the study was at an optimal saturation level for tree-based algorithms to reach their maximum learning capacity; however, it remained relatively limited in volume for the deep learning model to outperform the tree algorithms by addressing its data hunger.

In summary, based on the performance evaluations and literature comparisons conducted, it has been demonstrated that tree-based algorithms (particularly Random Forest) represent the most feasible approach for estimating the daily electricity consumption of industrial facilities, offering both low error rates and high explainability. Achieved using only single-variable historical consumption data without the need for complex external datasets, these high accuracy rates provide a highly reliable foundation for facility managers to make data-driven decisions and establish a standards-compliant energy baseline.

5. Conclusion

The primary objective of this study is to optimize energy management processes in industrial facilities and to develop a sustainable, data-driven, and practically feasible single-variable forecasting framework by overcoming the challenges of external data integration. To this end, daily electricity consumption data spanning six years from 28 manufacturing and assembly buildings—which account for the highest energy consumption within the facility—were dynamically aggregated based on their commissioning dates; following rigorous data preprocessing steps, the forecasting performance of ARIMA, Random Forest, XGBoost, and LSTM models was analyzed comparatively. The findings clearly demonstrated the superiority of tree-based algorithms in modeling the sudden fluctuations, weekend shutdowns, and effects of public holidays inherent in industrial operations. While the traditional ARIMA model proved inadequate in forecasting nonlinear dynamics, and the robust deep learning architecture LSTM lagged tree-based models on the tabulated daily dataset despite its stable structure, the Random Forest model—which partitions the data space using logical rules—emerged as the study's most successful predictor with a test accuracy rate of

94.10%. This success achieved using only historical consumption data and temporal feature engineering without the need for complex external sensor data (weather, production volume, etc.), demonstrates the practical value of the proposed methodology in the field. In conclusion, it has been confirmed that machine learning algorithms with minimized error metrics can be used as highly reliable decision-support tools in energy budgeting, establishing an energy baseline under ISO 50001, and planning processes for manufacturing facilities. In future studies, the goal is to investigate the behavioral generalizability of the algorithms by applying them to other functional facility categories (Process and Test Buildings, Offices, etc.) that were excluded from the scope of this study.

References

- Ait Lahoussine Ouali, H., Abida, O., Essalhi, M. and Moukhtar, I., Long short-term memory and metaheuristic algorithm-based deep learning model approach for forecasting green hydrogen production, *International Journal of Hydrogen Energy*, vol. 202, p. 152970, 2026.
- Amasyali, K., and El-Gohary, N. M., A review of data-driven building energy consumption prediction studies, *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1192-1205, 2018.
- Chinmayi, H. K., Flynn, K. C., Baath, G. S., Gowda, P., Northup, B. and Ashworth, A., Monitoring legume nutrition with machine learning: The impact of splits in training and testing data, *Applied Soft Computing*, vol. 176, pp. 113186, 2025.
- Geller, M. T. B., Silva, D. G. and Meneses, A. A. S., Deep learning applied to univariate electricity consumption time series: A systematic literature review, *Journal of Engineering and Computer Science*, vol. 5, pp. 100-115, 2025.
- Ghannam, S., An explainable comparative study of statistical, machine learning, deep learning, and hybrid models for CO2 emissions forecasting in Australia, *Array*, vol. 29, pp. 100639, 2026.
- Grinsztajn, L., Oyallon, E. and Varoquaux, G., Why do tree-based models still outperform deep learning on typical tabular data?, *Advances in Neural Information Processing Systems*, vol. 35, pp. 507-520, 2022.
- Kachalla, I.A., Ghiaus, C. and Baseer, M., Comparative analysis of machine learning models for prediction and forecasting of electric water boilers energy consumption, *Applied Thermal Engineering*, vol. 267, p. 125799, 2025.
- Khalil, M., McGough, A. S., Pourmirza, Z., Pazhoohesh, M. and Walker, S., Machine learning, deep learning and statistical analysis for forecasting building energy consumption—A systematic review, *Engineering Applications of Artificial Intelligence*, vol. 115, pp. 105287, 2022.
- Khan, M. A., Household energy consumption forecasting using LSTM and ARIMA models: A comprehensive analysis, *Research Square Preprint*, pp. 1-20, 2025.
- Lozano, A., Perez, M. and Gomez, R., A comparative study of energy consumption forecasting machine learning models, *Chemical Engineering Transactions*, vol. 103, pp. 115-120, 2023.
- Mijic, D., Electricity consumption forecasting using ARIMA and LSTM, *Zbornik radova Elektrotehničkog instituta Nikola Tesla*, vol. 34, pp. 1-15, 2024.
- Oyewola, D. O., Time series forecasting of energy demand using machine and deep learning approach, *International Research Journal of Engineering and Technology*, vol. 11, pp. 45-55, 2024.
- Shwartz-Ziv, R. and Armon, A., Tabular data: Deep learning is not all you need, *Information Fusion*, vol. 81, pp. 84-90, 2022.
- Zhu, J., Zhang, Y. and Wang, L., Comparative analysis of deep learning and tree-based models in power demand prediction, *Scientific Reports*, vol. 14, pp. 1-12, 2024.
- Zolfaghari, M. and Golabi, M.R., Modeling and predicting the electricity production in hydropower using conjunction of wavelet transform, long short-term memory and random forest models, *Renewable Energy*, vol. 170, pp. 1367-1381, 2021.

Biographies

Ebru YÜKSEL HALILOĞLU currently works as an Associate Professor at the Department of Industrial Engineering, Gazi University. She received her B.S. in Industrial Engineering from Middle East Technical University and PhD in Economics from Bilkent University. Her research focuses on macroeconomic modeling, energy policy and energy economics, applied macroeconomics, time series analysis and applied statistical analysis. She has published extensively in SSCI- and SCI-indexed international journals on oil price dynamics, electricity demand forecasting, financial markets, and macroeconomic modeling. Her recent research integrates advanced econometric and machine learning techniques into energy forecasting, trade modeling, and sectoral policy analyses.

Mehmet KABAK is a Professor of Gazi University, Department of Industrial Engineering. He has received MS degree from School of Science, Engineering and Technology, Istanbul Technical University in 2001. Kabak has received his PhD from Operations Research (Marmara University) at 2006. His research interests are multi criteria decision making, optimization, modeling, energy, fuzz logic.

Melike BİLDİRİCİ received her B.S. in Industrial Engineering from Ankara Yildirim Beyazit University and is currently pursuing an M.S. in Industrial Engineering at Gazi University. Professionally, she works as an Industrial Engineer in the energy sector within the defense industry. Her research interests include data analysis, machine learning, deep learning, and the development predictive models.