International Journal of Power Electronics and Drive System (IJPEDS)
Vol. 17, No. 4, December 2026, pp. 3075~3090
ISSN: 2088-8694, DOI: 10.11591/ijpeds.v17.i4.pp3075-3090

Voltage forecasting in IoT systems: classical and AI models

Sandro Carneiro Ribeiro[1], Leonardo de Carvalho Vidal[1,2]
[1]Graduate Program in Industrial Assembly, Fluminense Federal University (UFF), Niterói, Brazil
[2]Graduate Program in Industrial Automation and Robotics, Federal Institute of Rio de Janeiro (IFRJ), Volta Redonda, Brazil

Article Info

Article history:

Received Jul 22, 2026
Revised Sep 22, 2026
Accepted Oct 4, 2026

Keywords:

Electrical voltage monitoring
ESP32
Internet of things
Lagged regression
Machine learning
Time series analysis
Voltage forecasting

ABSTRACT

This paper presents a low-cost internet of things (IoT)-based electrical monitoring system with short-term voltage forecasting using an ESP32 microcontroller. The proposed system integrates electrical measurement, wireless communication, cloud storage, statistical analysis, and forecasting in a real residential installation. Voltage, current, active power, frequency, and power factor were continuously acquired, and the measurement system was experimentally validated against a calibrated reference instrument. The voltage time series was statistically characterized using the Augmented Dickey-Fuller test, autocorrelation function, and partial autocorrelation function, followed by the evaluation of six forecasting approaches: ARIMA, moving average, lagged regression, XGBoost, random forest, and long short-term memory (LSTM). A common chronological testing procedure was adopted to ensure a consistent out-of-sample comparison. Among the evaluated models, XGBoost achieved the best forecasting performance, with a root mean square error (RMSE) of 0.735 V, while random forest and lagged regression produced closely comparable results. The findings indicate that, for the monitored residential voltage series, relatively simple lag-based and machine-learning models can provide competitive short-term forecasting performance without the need for more computationally demanding architectures. These results highlight the relevance of temporal persistence and model complexity when designing IoT-based voltage forecasting systems.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Sandro Carneiro Ribeiro
Graduate Program in Industrial Assembly, Fluminense Federal University (UFF)
Niterói, Brazil
Email: sandro.ribeiro@inpe.br

1. INTRODUCTION

The continuous growth in global electricity demand, accelerated electrification, and the rapid expansion of artificial intelligence (AI) applications have increased the need for reliable, resilient, and energy-efficient electrical infrastructures [1]. This demand has been further intensified by the proliferation of AI-driven data centers and cloud computing services, whose substantial energy consumption places increasing pressure on modern power systems [2]. At the same time, the large-scale integration of renewable energy sources, particularly photovoltaic and wind generation, has introduced greater variability and operational complexity into electrical networks. These challenges have reinforced the need for continuous monitoring, accurate measurements, and data-driven decision-making to ensure reliable and efficient system operation [3].

Electrical monitoring has become an essential component of modern industrial, commercial, and residential infrastructures. It contributes to improving operational efficiency, reducing energy waste, maintaining power quality, facilitating the integration of renewable energy resources, and supporting intelligent energy management strategies [4]. In this context, continuous monitoring of electrical quantities,

particularly voltage, provides an important basis for characterizing the behavior of electrical systems and supporting forecasting and operational assessment.

Recent advances in embedded systems, cloud computing, and internet of things (IoT) technologies have enabled the development of low-cost monitoring platforms capable of continuously acquiring, transmitting, and analyzing electrical parameters in real time [5]. The integration of IoT technologies into electrical monitoring systems has significantly expanded the capabilities of conventional energy management infrastructures [4]. In particular, low-cost microcontrollers such as the ESP32 have become attractive solutions for real-time applications because of their integrated wireless communication, high processing capability, and reduced implementation cost [6]. Consequently, numerous studies have proposed IoT-based monitoring systems for smart grids, residential energy management, and industrial automation applications [4]. These distributed architectures also provide an appropriate foundation for moving part of data processing and intelligent analysis closer to the measurement device, an approach that is particularly relevant when real-time electrical monitoring and low computational overhead are required.

The increasing availability of large-scale electrical datasets has also stimulated the application of statistical methods, machine learning, and deep learning techniques for forecasting electrical variables [7]. Classical statistical approaches, including autoregressive and moving average models, remain widely used because of their interpretability, robustness, and low computational complexity [8]. More recently, machine learning algorithms such as XGBoost and random forest, together with deep learning architectures including long short-term memory (LSTM) neural networks, have shown promising results in forecasting nonlinear temporal processes and smart energy systems [7]. Although a substantial part of the recent forecasting literature focuses on energy and load forecasting, the same data-driven perspective is relevant to voltage time series, whose short-term behavior can contain temporal information suitable for predictive analysis. Despite the growing adoption of artificial intelligence techniques for electrical forecasting, many studies employ sophisticated machine learning models without systematically comparing their performance against simpler mathematical approaches [9]. This issue is particularly relevant for voltage forecasting, where the predictive benefit of increasing model complexity depends on the temporal characteristics of the measured voltage series.

Likewise, although numerous IoT-based electrical monitoring systems have been reported in the literature, most investigations focus on isolated aspects, such as hardware implementation, cloud connectivity, or forecasting algorithms [10]. Comprehensive studies integrating long-term real-world monitoring, rigorous experimental validation using calibrated industrial instrumentation, statistical characterization of voltage behavior, and systematic comparison of classical statistical, machine learning, and deep learning forecasting methods remain scarce. The integration of these aspects is particularly relevant to electrical energy systems because it connects physical measurement, power-quality observation, time-series characterization, and predictive intelligence within the same experimental framework. Furthermore, selecting the most appropriate forecasting model continues to be an open challenge because predictive performance strongly depends on the statistical characteristics of the monitored data and the specific application scenario [7]. Consequently, the relationship between the statistical properties of electrical voltage time series and forecasting performance remains insufficiently explored in practical IoT-based monitoring environments [11].

Accurate voltage monitoring is fundamental for electrical energy analysis, operational stability assessment, and power quality evaluation [12]. According to the IEEE 1159 [13] and IEEE 1459 [14] standards, the reliable acquisition of electrical quantities is essential for the proper characterization of electrical systems and monitoring infrastructures. Therefore, voltage monitoring and forecasting should be considered not only as data-analysis tasks, but also as complementary functions within intelligent electrical monitoring and power-quality assessment systems. To address these gaps, this work proposes a low-cost IoT-based electrical monitoring system using an ESP32 microcontroller for the near-real-time monitoring and transmission of electrical parameters. The proposed architecture is experimentally validated using a calibrated HP 3478A industrial multimeter as reference instrumentation. The study is based on a long-term real-world dataset comprising approximately 772924 voltage measurements collected over an eight-month monitoring period. In addition, classical statistical methods, machine learning algorithms, and deep learning models are systematically compared for voltage forecasting. Statistical analyses, including stationarity testing, autocorrelation analysis, and forecasting performance evaluation, are conducted to investigate how the statistical properties of voltage time series influence predictive performance.

The main contributions of this work are summarized as follows: i) Development and experimental validation of a low-cost ESP32-based IoT electrical monitoring system using calibrated industrial instrumentation; ii) Acquisition and analysis of a long-term real-world voltage dataset comprising approximately 772924 measurements collected over an eight-month monitoring period; iii) Statistical characterization of voltage behavior through stationarity and autocorrelation analyses, providing evidence

of strong short-term temporal dependence relevant to voltage predictability; iv) Systematic comparison of classical statistical, machine learning, and deep learning approaches for voltage forecasting; and v) Demonstration that increasing model complexity does not necessarily provide a substantial forecasting advantage for voltage time series exhibiting strong short-term temporal persistence.

2. METHOD

2.1. System architecture

The proposed system is a low-cost IoT-based electrical monitoring architecture developed around an ESP32 microcontroller for the near-real-time monitoring, processing, and transmission of electrical quantities [6]. The architecture integrates voltage and current sensing circuits, embedded signal processing, wireless communication, cloud storage, and data analysis modules. The overall methodological workflow of the proposed monitoring and voltage forecasting system is illustrated in Figure 1.

The novelty of the monitoring system does not rely solely on the use of the ESP32 or IoT communication, but on the integration of low-cost embedded acquisition, calibrated instrumentation, and long-term continuous monitoring in a real residential installation. This configuration provides a real-world dataset for subsequent voltage forecasting analysis. The forecasting contribution is distinct from the hardware contribution. Rather than proposing a new forecasting algorithm, this work systematically compares classical statistical, lagged regression, machine-learning, and deep-learning approaches using the same long-term residential voltage dataset. The analysis investigates whether increasing model complexity provides a meaningful forecasting advantage for a voltage series with strong temporal persistence.

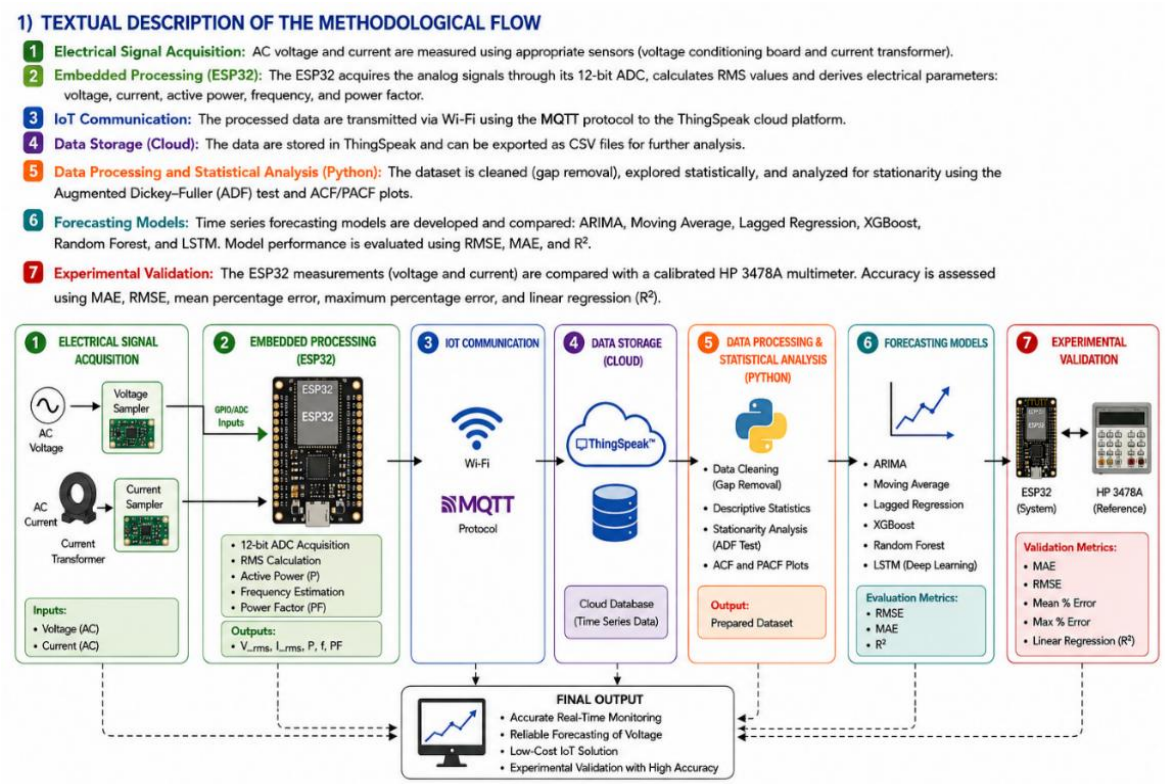

Figure 1. Overall methodological workflow of the proposed IoT-based electrical monitoring and voltage forecasting system

2.2. Electrical signal acquisition

The proposed system acquires voltage and current signals from the electrical network through dedicated sensing and signal-conditioning modules [15], [16]. The voltage measurement module is based on the ZMPT101B potential transformer, followed by a conditioning circuit using the LM358N operational amplifier. The ZMPT101B provides galvanic isolation between the electrical network and the acquisition circuit, approximately unity transformation ratio, and is specified for applications involving voltages up to

1 kV, with isolation withstand voltage of up to 4 kV. The conditioning stage adapts the transformer output to voltage levels compatible with the ESP32 ADC. The current measurement module employs the AC1025 current transformer (Talema Corporation), with a nominal transformation ratio of 1000:1, nominal primary current of 25 A, and maximum specified current of 75 A. The acquisition circuit comprises a burden resistor for conversion of the secondary current into voltage, biasing and coupling networks, and two LM358 operational-amplifier conditioning stages, providing an output compatible with the ESP32 acquisition system.

The conditioned analog signals are acquired by the 12-bit ESP32 ADC, with an input range of 0–3.3 V and 4096 quantization levels, corresponding to a theoretical resolution of approximately 0.806 mV/bit. The acquired samples are digitally processed to estimate RMS voltage and RMS current, as well as active power, apparent power, power factor, and frequency. The average interval between recorded measurements is approximately 20 s, enabling continuous long-term monitoring with low communication overhead. The monitoring system was deployed in a real residential electrical installation and operated continuously for approximately eight months. The hardware prototype developed by the authors is shown in Figure 2. The schematics, firmware, source code, and datasets generated during this study are publicly available in the following GitHub repository [17].



Figure 2. Prototype of the proposed ESP32-based IoT monitoring hardware

2.3. Embedded firmware processing

The embedded firmware performs real-time data acquisition, digital signal processing, and estimation of the monitored electrical quantities. Analog voltage and current signals are periodically sampled by the internal 12-bit ADC. The acquired samples are digitally processed to estimate RMS voltage, RMS current, active power, apparent power, power factor, and frequency.

The RMS voltage is calculated as (1).

$$V_{RMS} = \sqrt{\frac{1}{N}} \sum_{i=1}^{N} Vi^2 \quad (1)$$

Where: Vi represents the sampled voltage values; N is the total number of samples. Similarly, the RMS current is obtained as (2).

$$I_{RMS} = \sqrt{\frac{1}{N}} \sum_{i=1}^{N} Ii^2 \quad (2)$$

The active power is estimated by (3).

$$P = V_{RMS} \times I_{RMS} \times PF \quad (3)$$

Where: P is the active power; PF denotes the power factor. The apparent power is calculated as (4):

$$S = V_{RMS} \times I_{RMS} \quad (4)$$

and the power factor is determined by (5).

$$PF = \frac{P}{S} \quad (5)$$

After signal processing, the calculated electrical quantities are packaged into data packets and periodically transmitted to the ThingSpeak cloud platform through the integrated Wi-Fi interface for remote storage and visualization.

2.4. IoT communication and cloud storage

Following local processing, the calculated electrical quantities are transmitted through the ESP32 integrated Wi-Fi interface using IoT communication protocols. The monitored variables are stored in the ThingSpeak cloud platform, allowing remote visualization, long-term historical storage, and real-time access. The acquired dataset includes: timestamp, voltage, current, active power, accumulated energy, frequency, and power factor. The stored data are exported as CSV files for subsequent offline statistical processing and forecasting analyses using Python. ThingSpeak cloud services were employed for long-term data storage, remote visualization, and dataset exportation.

2.5. Dataset description and preprocessing

Before statistical and forecasting analyses, the acquired dataset underwent a preprocessing procedure to improve data quality and preserve the temporal structure of the monitored voltage series. The dataset contained 772924 measurement records. Timestamp information was converted to a consistent temporal format and the records were arranged in chronological order.

Temporal continuity was evaluated by calculating the time difference between consecutive observations. Temporal gaps greater than 120 s were used to identify interruptions in the acquisition sequence. Continuous segments containing fewer than 1000 observations were excluded as a conservative data-quality criterion. Records with missing voltage values were also removed. A sensitivity analysis of the minimum segment-length criterion showed that this threshold did not materially alter the retained dataset or the forecasting results. After preprocessing, 772,772 valid voltage observations remained available for statistical and forecasting analyses.

For forecasting evaluation, the valid observations were divided chronologically, without random shuffling, into 80% for training (618217 observations) and 20% for testing (154555 observations). Thus, the testing set represents the final 20% of the monitored time series and was kept completely out of model training. This chronological procedure prevents information from future observations from being used to train the forecasting models and provides an out-of-sample evaluation under temporal conditions representative of practical forecasting.

The forecasting task was defined as one-step-ahead prediction, in which the objective is to predict the next recorded voltage observation. Because the monitoring system records measurements at intervals of approximately 20 s, this corresponds to a short-term forecasting horizon of approximately 20 s ahead. For the lag-based models, five previous voltage observations were used as input variables. The five-lag configuration was adopted to represent the short-term temporal dependence identified in the statistical analysis while maintaining a compact input representation. The same five lagged voltage variables were used by the lagged regression, XGBoost, and random forest models, while the LSTM used the corresponding five-observation sequence as its input. ARIMA and moving average retained their respective temporal formulations.

The construction of the five lagged variables reduced the number of usable samples from 772772 to 772767, resulting in 618212 training samples and 154555 testing samples. The lag construction was performed using only preceding observations, so that no future voltage value was used as an input to predict an earlier observation. All forecasting models were evaluated using the same chronological out-of-sample testing period and the same evaluation metrics (RMSE, MAE, and R^2). Model-specific input representations were retained, while the underlying dataset, temporal split, forecasting target, and evaluation procedure were kept consistent across the six approaches.

2.6. Statistical analysis

The statistical properties of the monitored voltage series were investigated through descriptive and inferential analyses. Stationarity was assessed using the Augmented Dickey–Fuller (ADF) test [18], while temporal dependencies were analyzed through the autocorrelation function (ACF) and partial autocorrelation function (PACF) [19]. These analyses were performed to characterize the temporal behavior and dependence structure of the voltage time series and to support the selection of appropriate forecasting models.

Descriptive statistics, including the mean, standard deviation, minimum value, and maximum value, were also computed. The mean voltage is defined as (6).

$$\mu = \frac{1}{N} \sum_{i=1}^{N} x_i \quad (6).$$

Where: x_i represents each voltage observation; and N is the total number of observations. The sample standard deviation was calculated according to (7).

$$\sigma = \sqrt{\frac{1}{N-1}} \sum_{i=1}^{N} (x_i - \bar{x})^2 \quad (7)$$

Where x_i represents each voltage observation; $\bar{x}$ is the sample mean voltage; and N is the total number of observations.

2.7. Forecasting models

Six forecasting approaches were comparatively evaluated: ARIMA [19], moving average [19], linear regression with lagged features [20], XGBoost [21], random forest [22], and (LSTM) neural networks [23]. These models were grouped into three categories: i) Classical statistical methods (ARIMA and moving average); ii) Machine learning methods (lagged regression, XGBoost, and random forest); and iii) Deep learning methods using LSTM.

The forecasting task was defined as one-step-ahead prediction, corresponding to approximately 20 s ahead according to the acquisition interval. For the lag-based models, five previous voltage observations were used as input features. The same five lagged variables were provided to the lagged regression, XGBoost, and random forest models, while the LSTM received the corresponding five-observation sequence. ARIMA and moving average retained their respective temporal formulations.

For the machine learning models, the input variables consisted exclusively of the lagged voltage observations, without introducing additional exogenous variables. The models were trained using the chronological training subset and evaluated on the subsequent testing period, as described in section 2.5. No random shuffling of the temporal observations was performed.

The LSTM model consisted of a single LSTM layer with 50 units, followed by a fully connected dense layer with one output neuron for the predicted voltage value. The network was trained using 10 epochs and a batch size of 512. The LSTM input sequence comprised five consecutive voltage observations, consistent with the lag configuration adopted for the other lag-based models. The voltage data used by the LSTM were represented as 32-bit floating-point values and normalized using the minimum and maximum voltage values calculated from the training data. The same chronological training/testing principle was maintained during normalization so that information from the testing period was not used to determine the scaling parameters.

The XGBoost, random forest, and LSTM models were implemented using their respective computational frameworks, with the model configurations defined prior to evaluation. The experiments were performed on a computer equipped with an NVIDIA GeForce 940MX GPU, which was used to accelerate LSTM training through TensorFlow and compute unified device architecture (CUDA) Deep Neural Network Library (cuDNN). The remaining forecasting models were executed using the corresponding CPU-based implementations.

All models were evaluated using the same out-of-sample testing period and the same performance metrics: root mean square error (RMSE), mean absolute error (MAE), and R^2. This procedure provides a consistent basis for comparing forecasting performance while preserving the methodological characteristics of each forecasting approach. For reproducibility, the main configuration parameters used in the final implementation of each forecasting model are summarized in Table 1.

2.8. Forecasting performance metrics

Forecasting performance was evaluated using the MAE, RMSE, and coefficient of determination R^2 [23]. These complementary metrics quantify the magnitude of prediction errors and the agreement between measured and predicted voltage values, providing a consistent basis for comparing all forecasting models under the same out-of-sample testing conditions.

The MAE is defined as (8).

$$MAE = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i| \quad (8)$$

The RMSE is calculated as (9).

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2} \quad (9)$$

The coefficient of determination is given by (10).

$$R^2 = 1 - \frac{\sum_{i=1}^{N} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{N} (y_i - \bar{y})^2} \quad (10)$$

Where: y_i represents the measured voltage value; $\hat{y}_i$ denotes the corresponding predicted voltage value; $\bar{y}$ is the mean of the measured voltage values in the evaluation dataset; and N represents the number of observations in the testing dataset.

Table 1. Main configuration parameters of the forecasting models

Model	Main configuration
ARIMA	Order (1,0,0)
Moving average	Window size = 5
Lagged regression	5 voltage lags; linear regression
XGBoost	n_estimators = 100; max_depth = 6; learning_rate = 0.1; objective = reg:squarederror; random_state = 42
Random forest	n_estimators = 100; max_depth = 10; random_state = 42; n_jobs = −1
LSTM	Sequence length = 5; 50 LSTM units; Dense(1); Adam; MSE; batch size = 512; 10 epochs

2.9. Experimental validation

The proposed monitoring system was experimentally validated using a calibrated HP 3478A industrial multimeter as reference instrumentation [24]. A total of 80 voltage measurements were acquired and compared with those obtained by the proposed monitoring system. Measurement accuracy and linearity were evaluated using MAE, RMSE, and the coefficient of determination R^2, mean absolute percentage error (MAPE), maximum percentage error, and linear regression analysis.

The percentage error was calculated as (11).

$$Error(\%) = \frac{|V_{ref} - V_{iot}|}{V_{ref}} \times 100 \quad (11)$$

Where: V_{ref} represents the voltage measured by the HP 3478A reference instrument; and V_{iot} corresponds to the voltage measured by the proposed ESP32-based monitoring system. The calculated validation metrics were used to assess the measurement agreement, accuracy, and linearity of the proposed monitoring architecture.

3. RESULTS AND DISCUSSION

3.1. Experimental validation results

The proposed ESP32-based monitoring system was experimentally validated using a calibrated HP 3478A industrial multimeter as the reference instrument. A total of 80 voltage measurements were acquired under the operating conditions of the monitored residential electrical installation and compared with the corresponding measurements obtained by the proposed system. The validation was performed using the same electrical installation and measurement configuration described in section 2.9.

The comparison was performed using the reference voltage as the measurement baseline. The measurement error was calculated for each observation according to the difference between the reference instrument and the ESP32-based system. Measurement agreement was subsequently evaluated using MAE, RMSE, MAPE, maximum percentage error, R^2, and linear regression analysis. Representative measurements are presented in Table 2, while the complete set of 80 observations was used to calculate the validation metrics.

The quantitative validation results are summarized in Table 3. The proposed system achieved an MAE of 0.0246 V, RMSE of 0.0276 V, MAPE of 0.0199%, maximum percentage error of 0.0405%, and $R^2 = 0.9989$. Linear regression produced a slope coefficient of 0.995 and a correlation coefficient of 0.9994,

indicating a strong agreement between the proposed system and the reference instrument. The experimental results demonstrate that the proposed monitoring architecture achieved excellent measurement accuracy and linearity [25].

Table 2. Comparison between the reference instrument and the proposed monitoring system

Observation	HP 3478A (V)	ESP32 system (V)
1	122.37	122.39
20	124.98	124.94
40	123.22	123.20
60	122.16	122.15
80	124.08	124.06

These results indicate that the proposed low-cost acquisition system provides consistent voltage measurements under the tested residential operating conditions. The observed errors were considerably smaller than the natural voltage variations present in the monitored electrical network, indicating that the measurement error had a limited effect on the subsequent statistical and forecasting analyses. Figure 3 compares the voltage measurements acquired by the proposed monitoring system with those obtained from the HP 3478A reference instrument. The scatter distribution exhibits an almost ideal linear relationship, demonstrating excellent agreement between the two measurement systems.

The validation should nevertheless be interpreted within its experimental scope. The comparison was performed using one calibrated reference instrument, one residential installation, and the operating voltage conditions available during the validation measurements. Therefore, the results demonstrate measurement agreement and accuracy under the tested conditions rather than constituting a complete metrological uncertainty characterization or a multi-level voltage calibration study. Further validation under controlled voltage levels and different electrical loading conditions would be required to establish the behavior of the system over a broader operating range.

Table 3. Experimental validation results

Metric	Result
MAE	0.0246 V
RMSE	0.0276 V
MAPE	0.0199 %
Maximum percentage error	0.0405 %
Experimental R^2	0.9989

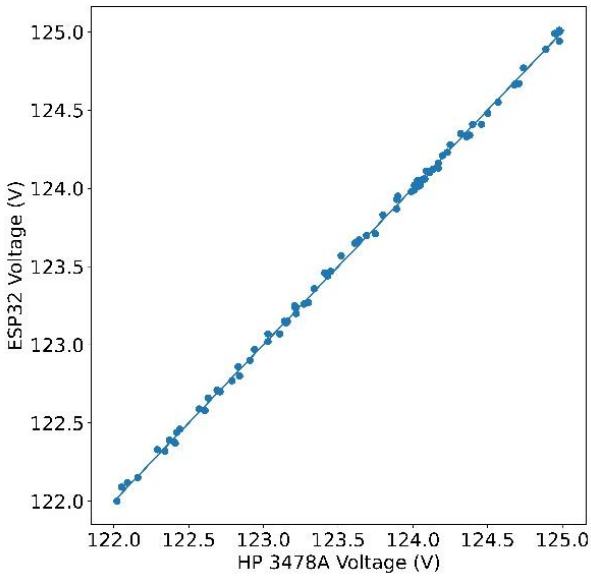

Figure 3. Experimental validation: HP 3478A versus ESP32

3.2. Statistical analysis of voltage series

Following the experimental validation of the proposed monitoring architecture, the long-term voltage measurements acquired by the IoT system were statistically analyzed to characterize the temporal behavior of the monitored electrical network. Although the proposed monitoring system continuously measures voltage, current, active power, frequency, and power factor, the forecasting analysis in this study focuses specifically on voltage. This selection was made because voltage is the primary variable of interest for evaluating short-term voltage predictability and for investigating the temporal structure of the electrical network. Focusing the forecasting analysis on a single electrical variable also allows the statistical characterization and comparison of all forecasting models to be performed under a clearly defined and controlled target.

The other measured electrical quantities remain relevant to the broader monitoring capability of the proposed system. Current and active power can provide information about load behavior and variations in electrical demand, while frequency and power factor provide complementary information about operating conditions and power-quality characteristics. These variables could therefore be incorporated as explanatory or exogenous variables in future forecasting studies, potentially providing additional information for voltage prediction and enabling broader analyses of electrical energy and power-quality behavior. However, their inclusion is outside the forecasting scope of the present study.

The monitoring campaign generated a large real-world voltage dataset collected over approximately eight months. The large number of observations provides a suitable basis for statistical characterization and subsequent forecasting evaluation. The descriptive statistical parameters of the monitored voltage series are summarized in Table 4.

The monitored voltage series exhibited relatively low dispersion around the mean value, indicating predominantly stable short-term behavior, although voltage excursions were observed during the monitoring period. Figure 4 illustrates the temporal evolution of the monitored voltage series. The measurements exhibit smooth transitions and strong continuity between consecutive observations, providing evidence of significant temporal dependence and short-term predictability. The observed voltage dynamics are consistent with the expected behavior of electrical distribution systems, in which short-term voltage variations generally occur gradually as a consequence of load fluctuations, network impedance characteristics, and voltage regulation mechanisms [26]. As a result, consecutive voltage samples retain substantial predictive information, making the series particularly suitable for short-term forecasting applications.

Table 4. Statistical analysis of the monitored voltage series

Statistical parameter	Value
Mean voltage	126.15 V
Standard deviation	2.86 V
Minimum voltage	107.10 V
Maximum voltage	135.10 V

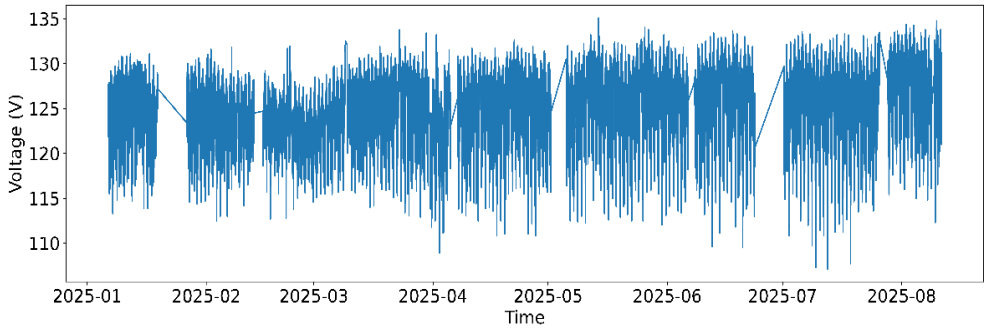

Figure 4. Temporal evolution of the monitored voltage series

3.3. Stationarity and temporal dependency analysis

To characterize the statistical properties and temporal dependence of the monitored voltage series, an augmented dickey-fuller (ADF) test was initially applied to the 772772 voltage observations used in the statistical analysis. The test yielded a statistic of −48.7662, with a p-value below 0.001, providing strong evidence against the null hypothesis of a unit root and supporting the stationarity of the monitored voltage series. To assess the robustness of this result with respect to the lag specification, the ADF test was additionally evaluated using fixed lag orders of 5, 10, 20, 40, and 60. In all cases, the null hypothesis of a

unit root was strongly rejected ($p < 0.001$), indicating that the stationarity conclusion was not dependent on a single lag specification.

The ADF test, however, provides information about stationarity and does not by itself establish whether the voltage series is deterministic, nonlinear, or predictable [27]. Therefore, the temporal dependence of the series was further examined using the autocorrelation function (ACF) and partial autocorrelation function (PACF). The ACF measures the correlation between observations separated by different lags, whereas the PACF measures the direct correlation at a given lag after accounting for the effects of shorter lags. The ACF results are presented in Figure 5, and the PACF results in Figure 6.

The ACF exhibited strong positive autocorrelation over the initial lags and remained high over an extended lag range. For example, the ACF was 0.971369 at lag 1 and remained 0.851200 at lag 10, indicating substantial persistence in the voltage series. In contrast, the PACF exhibited a markedly dominant first lag, with PACF (1) = 0.971369, followed by a substantially smaller PACF(2) = 0.141664 and PACF(3) = 0.069649. This pattern indicates that the immediately preceding voltage observation represents the dominant direct temporal dependency, while the contribution of subsequent lags is considerably smaller after accounting for the preceding observations. The five-lag configuration subsequently adopted in the forecasting models was therefore selected to represent the short-term temporal dependence identified in the series while maintaining a compact input representation. The selection was also consistent with the forecasting objective of predicting the next recorded voltage observation rather than modeling long-term seasonal behavior.

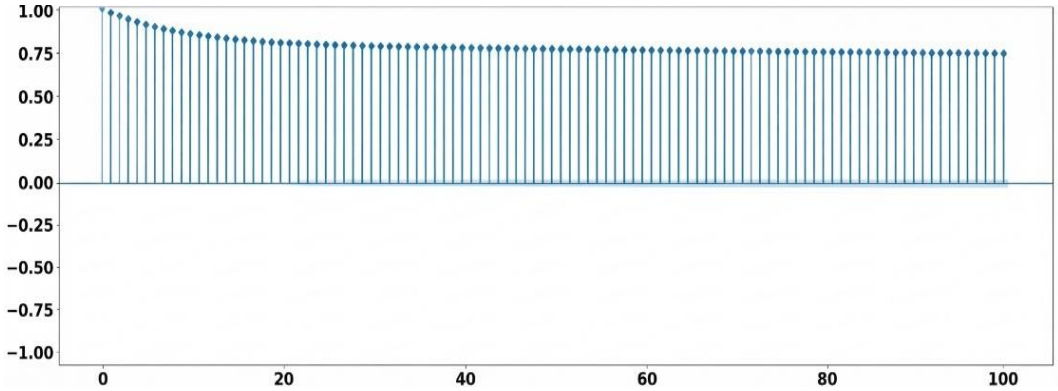

Figure 5. Autocorrelation function (ACF)

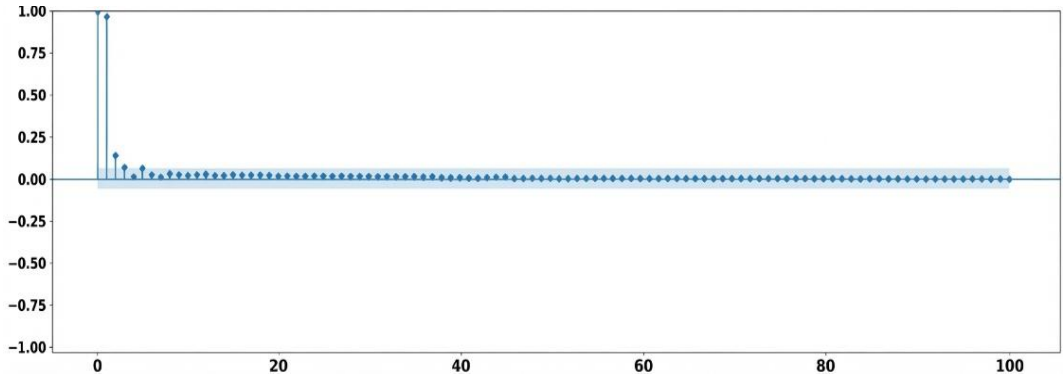

Figure 6. Partial autocorrelation function (PACF)

The eight-month monitoring period provides sufficient temporal coverage to observe the natural variation of the residential electrical installation. However, the statistical analyses presented here do not establish the existence or absence of formal daily or seasonal components. Such patterns may be influenced by household operating conditions and other external factors and would require dedicated seasonal decomposition or longer-term comparative analysis. The present study therefore focuses on the short-term temporal dependence relevant to the forecasting task [28].

Taken together, the stationarity, ACF, and PACF results demonstrate that the monitored voltage series possesses strong and persistent short-term temporal dependence. They do not, by themselves, demonstrate that the process is exclusively deterministic or that its nonlinear complexity is limited. Rather, the evidence for a predominantly predictable component emerges from the combination of these statistical characteristics with the forecasting results presented in the subsequent sections.

3.4. Forecasting results

3.4.1. ARIMA forecasting results

The ARIMA model was selected as the first classical statistical forecasting baseline based on the statistical characteristics of the monitored voltage series described in section 3.3. The ADF test indicated stationarity, while the ACF and PACF analyses showed strong temporal dependence and a markedly dominant partial autocorrelation at lag 1. Based on these characteristics, an ARIMA(1,0,0) specification was adopted, corresponding to a first-order autoregressive model without differencing or moving-average terms.

The choice of ARIMA(1,0,0) was therefore guided by the observed temporal structure of the voltage series rather than by an automatic order-selection procedure. A systematic search over alternative (p,d,q) configurations using AIC, BIC, or an automatic ARIMA procedure was not performed in the original experimental protocol. The model was intentionally retained as a simple autoregressive baseline, allowing its forecasting performance to be compared with the more complex machine-learning and deep-learning approaches under the same experimental conditions. The ARIMA(1,0,0) model was initially implemented as a classical statistical approach for voltage forecasting. The model achieved an RMSE of 3.394223 V, an MAE of 2.844255 V, and a negative coefficient of determination $R^2 = -0.370432$.

This choice is also consistent with the PACF results, in which the first lag presented a markedly dominant direct temporal dependence, while the contribution of subsequent lags was substantially smaller. The ARIMA(1,0,0) model therefore provides a parsimonious baseline for evaluating how much of the short-term voltage behavior can be captured by a simple autoregressive relationship. The model was trained using the chronological training dataset and evaluated on the subsequent testing period described in section 2.5. The resulting forecasting performance was evaluated using RMSE, MAE, and R^2, allowing direct comparison with the other forecasting approaches.

3.4.2. Moving average forecasting results

The moving average model was implemented as a simple baseline forecasting approach based on short-term temporal smoothing. The model achieved an RMSE of 0.956132 V, an MAE of 0.526209 V, and a coefficient of determination of $R^2=0.891254$. Compared with the ARIMA model, the moving average approach exhibited substantially better predictive performance. The low forecasting error indicates that the monitored voltage series possesses strong short-term temporal continuity, allowing local averaging techniques to generate accurate short-term forecasts.

The satisfactory performance achieved by this computationally efficient method indicates that voltage variations evolve gradually over time, with limited abrupt changes between consecutive observations. Such behavior is consistent with the expected dynamics of electrical distribution systems operating under relatively stable loading conditions, where voltage fluctuations are typically smooth and highly correlated over short time intervals. These findings provide additional evidence of the strong short-term temporal dependence identified in the statistical analyses. At the same time, they suggest that accurate voltage forecasting can be achieved using relatively simple models capable of exploiting the local temporal persistence of the monitored series.

3.4.3. Lagged regression forecasting results

The lagged regression model achieved highly competitive forecasting performance among all evaluated approaches, despite its very low computational complexity. The model obtained an RMSE of 0.740982 V, an MAE of 0.375461 V, and a coefficient of determination of $R^2=0.934688$. The outstanding predictive performance of the lagged regression model indicates that the monitored voltage dynamics are predominantly governed by immediate short-term temporal dependencies. Analysis of the regression coefficients revealed that the first temporal lag (lag_1) provides the greatest predictive contribution, whereas the influence of subsequent lags decreases progressively.

The competitive performance achieved by this relatively simple linear model suggests that the monitored electrical voltage series exhibits strong short-term temporal dependence and predominantly linear behavior, with limited evidence of additional nonlinear predictive structure. From a practical perspective, this finding indicates that future voltage values can be accurately estimated using only the most recent observations, reflecting the strong temporal persistence previously identified through the statistical analyses.

Another important advantage of the lagged regression approach is its very low computational complexity. Compared with more sophisticated machine learning and deep learning techniques, the model achieved highly competitive predictive accuracy while requiring minimal computational resources and training

time. This characteristic makes the approach particularly attractive for real-time forecasting applications in embedded and edge-computing environments, where processing capability and energy consumption are often constrained. Overall, these findings demonstrate that simple linear lagged-regression models can provide highly accurate short-term voltage forecasts when the monitored electrical system exhibits stationary behavior, strong temporal persistence, and predominantly linear dynamics.

3.4.4. XGBoost and random forest forecasting results

Machine learning forecasting models based on the XGBoost and random forest algorithms were implemented to investigate whether nonlinear relationships could further improve voltage forecasting performance. The XGBoost model achieved an RMSE of 0.734667 V, an MAE of 0.369024 V, and a coefficient of determination of R^2=0.935797, whereas the random forest model achieved an RMSE of 0.738314 V, an MAE of 0.369364 V, and an R^2=0.935158.

Both machine learning approaches demonstrated excellent predictive capability and high forecasting stability. However, their improvement over the simpler lagged regression approach was relatively small despite their substantially greater modeling complexity. Feature importance analysis indicated that the first temporal lag accounted for approximately 97% and 98% of the total feature importance in the XGBoost and random forest models, respectively. This remarkable concentration of predictive importance demonstrates that the forecasting models relied predominantly on the most recent voltage observation.

An important outcome of this analysis is that two fundamentally different ensemble learning algorithms independently identified the same dominant predictive feature. This convergence strongly supports the statistical evidence obtained from the ACF, PACF, and lagged regression analyses, reinforcing the conclusion that short-term voltage forecasting is primarily driven by immediate temporal persistence rather than by complex nonlinear relationships. The limited improvement obtained by the nonlinear machine learning models suggests that additional nonlinear structure did not provide a substantial predictive advantage for the monitored voltage series. Consequently, increasing model sophistication does not necessarily translate into improved forecasting accuracy when the underlying time series is predominantly stationary, highly persistent, and characterized predominantly by short-term linear dynamics.

3.4.5. LSTM forecasting results

In the original experimental implementation, the LSTM model was trained using a reduced subset of 20,000 samples due to the computational limitations of the available CPU-based environment. This restriction was computational rather than methodological. For the revised experimental analysis, the computational limitation was removed through GPU acceleration, and all forecasting models were re-evaluated using the complete available preprocessed dataset under the same standardized experimental protocol. For the LSTM, this resulted in 772767 sequences, comprising 618212 training sequences and 154555 testing sequences, generated from the 772772 valid voltage observations. The LSTM consisted of one layer with 50 units followed by a Dense layer with one output neuron. The model was trained for 10 epochs using a batch size of 512. The experiment was executed using an NVIDIA GeForce 940MX GPU with TensorFlow and cuDNN acceleration.

The resulting LSTM performance was RMSE = 0.750965 V, MAE = 0.391630 V, and R^2=0.932916. The use of the complete dataset removes the sample-size restriction present in the original implementation and provides a consistent basis for comparison with the other forecasting models. Although the LSTM network demonstrated satisfactory forecasting capability, its predictive performance remained inferior to that of XGBoost, random forest, and lagged regression, although it outperformed the moving average and ARIMA baselines [29].

3.5. Comparative analysis of forecasting models

Table 5 summarizes the forecasting performance of all evaluated approaches. The computational complexity categories are reported qualitatively based on the model structure and the computational requirements observed during the experiments. The comparative analysis shows that the three best-performing approaches—XGBoost, random forest, and lagged regression—produced very similar forecasting errors. XGBoost achieved the lowest RMSE (0.734667 V), followed by random forest (0.738314 V) and lagged regression (0.740982 V). The differences among these three models are therefore relatively small compared with the substantially larger error observed for ARIMA.

The results also indicate that increasing model complexity did not produce a substantial forecasting advantage in the present application. XGBoost and random forest achieved slightly lower errors than the very-low-complexity lagged regression model, while the LSTM produced a higher RMSE despite its considerably greater computational complexity. The first lag accounted for approximately 97% and 98% of the total feature importance in the XGBoost and random forest models, respectively. In the present experiment, increasing model complexity did not necessarily improve forecasting performance [30].

Table 5. Summary of forecast model analyses

Model	RMSE (V)	MAE (V)	R^2	Relative computational complexity	Potential application
XGBoost	0.734667	0.369024	0.935797	Moderate	Short-term voltage forecasting and nonlinear data analysis
Random forest	0.738314	0.369364	0.935158	Moderate	Short-term voltage forecasting and feature analysis
Lagged regression (5)	0.740982	0.375461	0.934688	Very low	Lightweight real-time and embedded forecasting
LSTM (GPU)	0.750965	0.391630	0.932916	High	Complex temporal forecasting and offline analysis
Moving average (5)	0.956132	0.526209	0.891254	Very low	Simple monitoring and baseline forecasting
ARIMA	3.394223	2.844255	−0.370432	Low	Classical statistical baseline

From a practical perspective, the forecasting errors should be interpreted relative to the voltage magnitude of the monitored installation. The best-performing model produced an RMSE of 0.734667 V, corresponding to approximately 0.58% of a nominal 127 V voltage level. The LSTM RMSE of 0.750965 V corresponds to approximately 0.59%, while the lagged regression RMSE corresponds to approximately 0.58%. These errors are small relative to the nominal voltage magnitude and indicate that the evaluated models can provide useful short-term predictions of the monitored voltage under the tested residential operating conditions. The forecasting horizon is approximately 20 s, corresponding to the measurement interval of the monitoring system.

The forecasting results should nevertheless be distinguished from a formal power-quality compliance assessment. The forecasting RMSE represents the error between predicted and measured voltage values and should not be interpreted directly as a voltage-quality violation or as an equivalent to the limits specified for power-quality assessment. The results instead indicate the accuracy with which the short-term voltage behavior can be predicted from the monitored time series.

The computational comparison is also relevant to practical implementation. The lagged regression and moving average approaches require very low computational resources, whereas XGBoost and random forest provide slightly better predictive performance at moderate computational cost. LSTM requires substantially greater computational resources, although graphics processing unit (GPU) acceleration enabled training on the complete dataset within the available computational environment. This trade-off is relevant for embedded and edge-oriented applications in which computational resources, memory, and energy consumption may be constrained.

Because all models were evaluated on the same out-of-sample observations, the prediction errors are paired across models. A formal statistical significance analysis would require the individual prediction errors from each model rather than only the aggregate RMSE, MAE, and R^2 values reported in Table 5. Therefore, the present comparison is interpreted descriptively, and the relatively small differences among XGBoost, random forest, and lagged regression should not be described as statistically significant without an additional paired hypothesis test or confidence-interval analysis. Overall, the results indicate that XGBoost achieved the lowest numerical RMSE, but its advantage over random forest and lagged regression was small. Considering the combination of predictive performance and computational complexity, the lagged regression model remains particularly attractive for lightweight short-term voltage forecasting, while XGBoost provides the best numerical accuracy among the evaluated approaches.

From a practical perspective, these findings demonstrate that accurate short-term voltage forecasting can be achieved using lightweight forecasting models with minimal computational requirements. This characteristic is particularly attractive for embedded and edge-computing IoT applications, where processing power, memory, and energy consumption are inherently limited [31]. Consequently, low-cost IoT monitoring systems combined with computationally efficient forecasting models provide an effective solution for real-time voltage monitoring without requiring complex artificial intelligence architectures.

4. CONCLUSION

This study presented the development and experimental evaluation of a low-cost IoT-based electrical monitoring system using an ESP32 microcontroller, together with a comparative analysis of statistical, machine learning, and deep learning approaches for short-term voltage forecasting. The monitoring architecture enabled continuous acquisition of electrical quantities over approximately eight months in a real residential installation. Experimental comparison with a calibrated HP 3478A multimeter showed close agreement between the proposed system and the reference instrument under the tested operating conditions. These results support the suitability of the developed platform for the monitoring and forecasting analyses performed in this study, although they should not be interpreted as a complete metrological characterization of the measurement system.

Statistical analysis revealed stationarity and strong short-term temporal persistence in the monitored voltage series. The ACF and PACF analyses showed substantial temporal dependence, with the immediately preceding voltage observation representing the dominant direct dependency. These statistical characteristics were consistent with the forecasting results, in which models based primarily on recent voltage observations achieved high predictive performance.

Among the evaluated models, XGBoost achieved the lowest numerical forecasting error, with an RMSE of 0.734667 V, followed closely by random forest (0.738314 V), lagged regression (0.740982 V), and LSTM (0.750965 V). The relatively small differences among these models indicate that, for the monitored residential voltage series and the approximately 20-s forecasting horizon considered in this study, increasing model complexity provided only limited improvement in predictive accuracy. In this specific experimental context, lightweight lag-based approaches therefore offer an attractive balance between forecasting performance and computational complexity.

The findings should be interpreted within the scope of the experimental conditions. The dataset was obtained from a single residential electrical installation, and the results cannot be directly generalized to industrial, commercial, or other distribution-network environments. Electrical systems containing distributed renewable generation, rapidly changing loads, significant voltage disturbances, different sampling intervals, or different operating characteristics may exhibit temporal behavior substantially different from that observed in this study. Consequently, the relative performance of the forecasting models may also change under these conditions.

Although the forecasting models evaluated in this study used only lagged voltage observations, the additional electrical quantities acquired by the monitoring platform may provide complementary information for voltage prediction. Current can reflect changes in the loading condition of the residential installation and may therefore help identify voltage variations associated with load changes. Active power provides information about electrical demand, while power factor may indicate changes in the characteristics of the connected loads. Frequency may also provide information about broader operating conditions of the electrical system. Incorporating lagged measurements of these variables as exogenous predictors could allow multivariate forecasting models to capture relationships that are not represented by the voltage series alone. However, their actual contribution to predictive accuracy should be established empirically, particularly because the strong short-term voltage persistence observed in the present study already provides substantial predictive information.

Future work should therefore evaluate the proposed methodology using multivariate forecasting approaches and different electrical networks and operating scenarios, including installations with distributed generation and more dynamic load profiles. Such studies could determine whether the additional electrical variables improve forecasting accuracy and whether the observed dominance of short-term voltage persistence is maintained under different operating conditions.

Overall, the results demonstrate that the combination of long-term IoT monitoring, statistical time-series characterization, and comparative forecasting analysis provides a practical framework for investigating short-term residential voltage predictability. Rather than establishing the general superiority of a particular forecasting class, the results emphasize the importance of selecting model complexity according to the temporal characteristics of the monitored electrical system and the requirements of the intended application.

FUNDING INFORMATION

The authors confirm that the research was carried out independently without financial influence.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Sandro Carneiro Ribeiro	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			✓	
Leonardo de Carvalho Vidal	✓	✓			✓	✓			✓	✓	✓	✓	✓	

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

The dataset supporting the findings of this study, comprising electrical measurements collected over approximately eight months in a residential installation, is publicly available at reference [17].

REFERENCES

[1] International Energy Agency., "Demand - electricity 2026," IEA. Accessed: Jul. 08, 2026. [Online]. Available: https://www.iea.org/reports/electricity-2026/demand.

[2] Y. Van Fan, C. Wilson, G. Kamiya, and A. Mastrucci, "Data centre energy demand projections within shared socioeconomic pathways," *Energy and Climate Change*, vol. 7, 2026, doi: 10.1016/j.egycc.2026.100253.

[3] M. Khalid, "Smart grids and renewable energy systems: perspectives and grid integration challenges," *Energy Strategy Reviews*, vol. 51, Jan. 2024, doi: 10.1016/j.esr.2024.101299.

[4] D. E. A. Mansour *et al.*, "Applications of IoT and digital twin in electrical power systems: a comprehensive survey," *IET Generation, Transmission and Distribution*, vol. 17, no. 20, pp. 4457–4479, 2023, doi: 10.1049/gtd2.12940.

[5] D. Witczak and S. Szymoniak, "Review of monitoring and control systems based on internet of things," *Applied Sciences (Switzerland)*, vol. 14, no. 19, 2024, doi: 10.3390/app14198943.

[6] H. J. El-Khozondar *et al.*, "A smart energy monitoring system using ESP32 microcontroller," *e-Prime - Advances in Electrical Engineering, Electronics and Energy*, vol. 9, 2024, doi: 10.1016/j.prime.2024.100666.

[7] R. Wazirali, E. Yaghoubi, M. S. S. Abujazar, R. Ahmad, and A. H. Vakili, "State-of-the-art review on energy and load forecasting in microgrids using artificial neural networks, machine learning, and deep learning techniques," *Electric Power Systems Research*, vol. 225, 2023, doi: 10.1016/j.epsr.2023.109792.

[8] S. Ali, S. Bogarra, M. N. Riaz, P. P. Phyo, D. Flynn, and A. Taha, "From time-series to hybrid models: advancements in short-term load forecasting embracing smart grid paradigm," *Applied Sciences (Switzerland)*, vol. 14, no. 11, 2024, doi: 10.3390/app14114442.

[9] A. B. Nassif, B. Soudan, M. Azzeh, I. Attilli, and O. Almulla, "Artificial intelligence and statistical techniques in short-term load forecasting: a review," *International Review on Modelling and Simulations*, vol. 14, no. 6, pp. 408–430, 2021, doi: 10.15866/iremos.v14i6.21328.

[10] N. S. e. Silva, R. Castro, and P. Ferrão, "Smart grids in the context of smart cities: a literature review and gap analysis," *Energies*, vol. 18, no. 5, 2025, doi: 10.3390/en18051186.

[11] A. Mohanty *et al.*, "Smart grid and application of big data: opportunities and challenges," *Sustainable Energy Technologies and Assessments*, vol. 71, 2024, doi: 10.1016/j.seta.2024.104011.

[12] T. Ravi and K. Sathish Kumar, "Analysis, monitoring, and mitigation of power quality disturbances in a distributed generation system," *Frontiers in Energy Research*, vol. 10, 2022, doi: 10.3389/fenrg.2022.989474.

[13] IEEE Standards Association, "IEEE Standard 1159," [Online]. Available: https://standards.ieee.org/ieee/1159/6124 . Accessed: Jul. 13, 2026.

[14] IEEE Standards Association, "IEEE Standard 1459," [Online]. Available: https://standards.ieee.org/ieee/1459/7578 . Accessed: Jul. 13, 2026.

[15] I. Abu-Mahfouz, *Instrumentation: theory and practice, part 1*. in Synthesis Lectures on Mechanical Engineering. Cham: Springer International Publishing, 2022, doi: 10.1007/978-3-031-15246-7.

[16] K. Tomczyk, M. Sieja, K. Ostrowska, and D. Owczarek, "Review of accuracy assessment methods for current transformers: errors, uncertainties, and dynamic performance," *Energies*, vol. 18, no. 18, Sep. 2025, doi: 10.3390/en18184995.

[17] SandroRibeiro-PY1SAN, "Advanced power quality monitoring system with IoT, cloud, and AI embedded." Accessed: Jul. 21, 2026. [Online]. Available: https://github.com/SandroRibeiro-PY1SAN/Advanced-Power-Quality-Monitoring-System-with-IoT-Cloud-and-AI-Embedded.

[18] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, vol. 74, no. 366a, pp. 427–431, Jun. 1979, doi: 10.1080/01621459.1979.10482531.

[19] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time series analysis*. in Wiley Series in Probability and Statistics. Wiley, 2008, doi: 10.1002/9781118619193.

[20] G. James, D. Witten, T. Hastie, and R. Tibshirani, "An introduction to statistical learning: with applications in R," in *Springer Texts in Statistics*, Springer, 2017.

[21] T. Chen and C. Guestrin, "XGBoost," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA: ACM, Aug. 2016, pp. 785–794, doi: 10.1145/2939672.2939785.

[22] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001, doi: 10.1023/A:1010933404324.

[23] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[24] H. Carstens, X. Xia, and S. Yadavalli, "Measurement uncertainty in energy monitoring: present state of the art," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 2791–2805, Feb. 2018, doi: 10.1016/j.rser.2017.10.006.

[25] N. Venkata *et al.*, "Challenges and opportunities in calibrating low-cost environmental sensors," *Sensors*, vol. 24, no. 11, pp. 1–22, 2024.

[26] R. Cleenwerck, R. Claeys, T. Coosemans, and J. Desmet, "Distribution networks in transition: a European survey on challenges and opportunities," *Renewable and Sustainable Energy Reviews*, vol. 234, Jul. 2026, doi: 10.1016/j.rser.2026.116885.

[27] M. Hasan *et al.*, "A state-of-the-art comparative review of load forecasting methods: characteristics, perspectives, and applications," *Energy Conversion and Management: X*, vol. 26, Apr. 2025, doi: 10.1016/j.ecmx.2025.100922.

[28] L. Baur, K. Ditschuneit, M. Schambach, C. Kaymakci, T. Wollmann, and A. Sauer, "Explainability and interpretability in electric load forecasting using machine learning techniques – a review," *Energy and AI*, vol. 16, May 2024, doi: 10.1016/j.egyai.2024.100358.
[29] Y. Eren and İ. Küçükdemiral, "A comprehensive review on deep learning approaches for short-term load forecasting," *Renewable and Sustainable Energy Reviews*, vol. 189, Jan. 2024, doi: 10.1016/j.rser.2023.114031.
[30] V. I. Kontopoulou, A. D. Panagopoulos, I. Kakkos, and G. K. Matsopoulos, "A review of ARIMA vs. machine learning approaches for time series forecasting in data driven networks," *Future Internet*, vol. 15, no. 8, Jul. 2023, doi: 10.3390/fi15080255.
[31] R. Singh and S. S. Gill, "Edge AI: a survey," *Internet of Things and Cyber-Physical Systems*, vol. 3, pp. 71–92, 2023, doi: 10.1016/j.iotcps.2023.02.004.

BIOGRAPHIES OF AUTHORS



Sandro Carneiro Ribeiro received the M.Sc. degree in automation engineering from the University of Taubaté and is currently pursuing the Ph.D. degree in electrical engineering at Fluminense Federal University (UFF), Brazil. He is a technologist at the National Institute for Space Research (INPE), where he works with radio-frequency ground station systems for satellite telemetry, tracking, and control. He has more than 33 years of experience in embedded systems, electronic instrumentation, automation, and radio-frequency communication systems. His current research interests include the internet of things (IoT), embedded systems, electrical monitoring, power quality, artificial intelligence, machine learning, and smart energy systems. He can be contacted at email: sandro.ribeiro@inpe.br.



Leonardo de Carvalho Vidal is a full professor at the Federal Institute of Rio de Janeiro (IFRJ), Brazil, where he coordinates the graduate program in industrial automation and robotics. He is also a permanent faculty member of the graduate program in industrial assembly at Fluminense Federal University (UFF), Brazil. He received the Ph.D. degree in Electrical Engineering from the Federal University of Itajubá (UNIFEI), Brazil, and the M.Sc. degree in mechanical engineering with emphasis on automation. His research interests include industrial automation, industrial instrumentation, dynamic systems modeling and identification, control systems, computational intelligence, robotics, smart cities, and smart factory technologies. He can be contacted at email: leonardo.vidal@ifrj.edu.br.