2026 Volume: 02 Issue: 05 www.ijrpa.com ISSN 2456-9995 Review Article

International Journal Research Publication Analysis

Page: 01-14

MACHINE LEARNING-BASED SMART ENERGY MONITORING SYSTEM USING IOT

[1]Akshay T. Bodare, [2]Amir R. Attar, [3]Dr. Manisha V. Bhanuse

[1]Student E&TC D.Y. Patil College of Engineering and Technology Kolhapur Maharashtra.

[2]Student E& D.Y. Patil College of Engineering and Technology Kolhapur Maharashtra.

[3]Associate Prof. E& TC D.Y. Patil College of Engineering and Technology Kolhapur Maharashtra.

Article Received: 09April 2026

Article Revised: 29 April 2026

Published on: 19 May 2026

***Corresponding Author: Akshay T. Bodare**

Student E&TC D.Y. Patil College of Engineering and Technology Kolhapur Maharashtra

DOI: https://doi-doi.org/101555/ijrpa.8315

ABSTRACT

The exponential growth of global energy demand, coupled with the urgency of climate change mitigation, necessitates intelligent and adaptive energy management systems. This paper presents a comprehensive Machine Learning-Based Smart Energy Monitoring System (ML-SEMS) that leverages the Internet of Things (IoT) infrastructure to deliver real-time energy consumption analytics, predictive load forecasting, and automated demand-response mechanisms. The proposed system integrates edge-deployed IoT sensor nodes with a cloud-based analytical engine comprising ensemble machine learning models—specifically, Long Short-Term Memory (LSTM) networks for time-series forecasting, Random Forest classifiers for appliance identification, and an Isolation Forest algorithm for anomaly detection. Data collected from smart meters, current transformers, voltage sensors, and environmental sensors are transmitted over MQTT and HTTP REST protocols to a centralized data lake. A total of 18 months of real-world energy consumption data from 120 residential and 30 commercial premises were used for model training and validation. Experimental results demonstrate that the LSTM forecasting model achieves a Mean Absolute Percentage Error (MAPE) of 3.47%, compared to 8.92% for a traditional ARIMA baseline, while the Random Forest appliance classifier attains 96.3% accuracy. The anomaly detection module successfully identifies 94.7% of energy theft and equipment fault events. The system further incorporates a user-friendly dashboard and mobile application, enabling consumers to monitor consumption patterns, receive actionable energy-saving recommendations, and participate in utility demand-response programs. The proposed ML-SEMS achieves an

average energy saving of 18.6% across test premises, demonstrating significant potential for large-scale smart grid deployment.

KEYWORDS: Smart Energy Monitoring, Internet of Things (IoT), Machine Learning, LSTM, Random Forest, Anomaly Detection, Demand Response, Smart Grid, Energy Efficiency, Edge Computing.

1. INTRODUCTION

The global energy landscape is undergoing a paradigm shift driven by digitalization, decarbonization, and decentralization. According to the International Energy Agency (IEA), global electricity demand is projected to increase by 4.3% annually through 2030, placing unprecedented pressure on existing grid infrastructure. Conventional energy monitoring approaches, which rely on periodic manual meter readings and reactive fault management, are inadequate to meet the demands of modern energy systems characterized by volatile renewable generation, dynamic loads, and the proliferation of electric vehicles.

The Internet of Things (IoT) has emerged as a transformative technology enabling the deployment of pervasive sensing and connectivity infrastructure. Smart meters, embedded sensors, and edge computing devices now provide granular, real-time visibility into energy flows at the device, premise, and grid level. However, the sheer volume and velocity of IoT-generated data overwhelm traditional data processing and analytics paradigms, motivating the integration of machine learning (ML) methods capable of extracting actionable insights from high-dimensional, non-stationary energy data.

Machine learning approaches—ranging from classical statistical models to deep neural architectures—have demonstrated remarkable efficacy in energy forecasting, fault detection, and load disaggregation. Nevertheless, existing research predominantly addresses individual sub-problems in isolation, neglecting the holistic system integration challenges that arise in real-world deployments. Furthermore, few studies rigorously evaluate performance under diverse operating conditions spanning residential, commercial, and industrial environments.

This paper addresses these gaps by presenting the design, implementation, and evaluation of a Machine Learning-Based Smart Energy Monitoring System (ML-SEMS). The principal contributions of this work are:

- A unified IoT architecture spanning edge sensor nodes, a fog processing layer, and a cloud analytical backend, optimized for low-latency data acquisition and scalable model inference.

- An ensemble ML pipeline integrating LSTM-based energy forecasting, Random Forest appliance classification, and Isolation Forest anomaly detection, achieving state-of-the-art performance on real-world datasets.
- A demand-response optimization module that autonomously schedules deferrable loads to minimize peak demand and energy costs.
- An end-to-end evaluation on 18 months of data from 150 premises, demonstrating 18.6% average energy savings and superior forecasting accuracy relative to established baselines.
- A scalable, open-architecture deployment framework compatible with heterogeneous IoT hardware and cloud platforms.

The remainder of the paper is organized as follows: Section 2 reviews relevant literature. Section 3 describes the system architecture and hardware design. Section 4 presents the machine learning methodology. Section 5 details the experimental setup and results. Section 6 discusses implications and limitations. Section 7 concludes with directions for future work.

2. LITERATURE REVIEW

2.1 IoT-Based Energy Monitoring Systems

Early IoT-based energy monitoring systems focused primarily on real-time data acquisition and visualization. Fang et al. (2011) proposed a smart grid communication framework leveraging ZigBee and Wi-Fi sensor networks for residential energy monitoring. Subsequent work by Karfopoulos et al. (2013) demonstrated the feasibility of deploying distributed smart meter networks at scale, though their system lacked intelligent analytics capabilities. More recently, Zhao et al. (2019) integrated LoRaWAN-based long-range IoT communication with cloud dashboards, enabling monitoring across geographically dispersed commercial premises. However, their system relied on rule-based anomaly detection, limiting sensitivity to novel fault patterns.

Sharma and Sood (2020) proposed an edge-fog-cloud architecture for industrial energy monitoring, demonstrating significant reductions in cloud data transfer costs through edge pre-processing. Their work highlighted the importance of hierarchical system design, though ML-based analytics were not incorporated. The present work builds upon these architectural foundations while adding sophisticated ML capabilities.

2.2 Machine Learning for Energy Forecasting

Energy forecasting has been extensively studied using both classical statistical and modern ML approaches. Box-Jenkins ARIMA models remain widely used baselines (Taylor, 2003), but their linearity assumption limits performance on non-stationary, seasonally complex energy series. Support Vector Regression (SVR) was applied by Chen et al. (2015) for short-term load forecasting, yielding improvements over ARIMA but requiring careful feature engineering. Gradient Boosting methods, particularly XGBoost, demonstrated competitive performance in the GEFCom2014 forecasting competition (Hong et al., 2016).

Deep learning approaches have significantly advanced the state-of-the-art. Recurrent Neural Networks (RNNs), and specifically Long Short-Term Memory (LSTM) networks introduced by Hochreiter and Schmidhuber (1997), excel at capturing long-range temporal dependencies in energy time series. Kong et al. (2019) demonstrated that LSTM-based models outperform ARIMA by 34% on residential load forecasting tasks. More recently, Transformer-based architectures have shown promise for multi-step ahead forecasting (Wu et al., 2021), though their computational requirements can limit edge deployment feasibility.

2.3 Non-Intrusive Load Monitoring (NILM)

Non-Intrusive Load Monitoring (NILM), pioneered by Hart (1992), enables appliance-level disaggregation of aggregate electricity consumption without per-appliance sensing. Early approaches relied on steady-state power signatures and combinatorial optimization. Machine learning-based NILM methods have since demonstrated substantial improvements: Kolter and Johnson (2011) applied Factorial Hidden Markov Models, while Kelly and Knottenbelt (2015) demonstrated the efficacy of deep learning for NILM using the UK-DALE dataset. Random Forest classifiers have been employed for appliance identification from smart meter data, with Rashid et al. (2019) reporting 94.2% classification accuracy across 10 appliance categories.

2.4 Anomaly Detection in Smart Grids

Anomaly detection in energy systems encompasses energy theft detection, equipment fault identification, and demand-side irregularity monitoring. Statistical process control methods have been widely applied but suffer from high false-positive rates in dynamic environments. Jokar et al. (2016) employed Support Vector Machines for electricity theft detection, achieving 93.1% detection accuracy. Unsupervised methods, particularly the Isolation Forest algorithm (Liu et al., 2008), have gained traction for their computational efficiency and robustness to class imbalance, which is inherent in anomaly detection problems. Zheng et al. (2020) demonstrated Isolation Forest superiority over One-Class SVM on smart meter anomaly datasets.

2.5 RESEARCH GAPS

A systematic review of the literature reveals several critical gaps: (1) the lack of end-to-end system architectures integrating IoT hardware, communication protocols, and ML analytics; (2) limited evaluation on diverse, real-world datasets spanning multiple building types and climatic conditions; (3) insufficient attention to system scalability and deployment practicality; and (4) the absence of integrated demand-response optimization in energy monitoring systems. This paper directly addresses these gaps through the holistic ML-SEMS design and rigorous empirical evaluation.

3. SYSTEM ARCHITECTURE

The ML-SEMS adopts a three-tier hierarchical architecture comprising: (1) the Edge Sensing Layer, (2) the Fog Processing Layer, and (3) the Cloud Analytics Layer. This stratified design enables localized real-time control at the edge, intermediate aggregation and preprocessing in the fog tier, and computationally intensive ML analytics in the cloud.

3.1 Edge Sensing Layer

The edge sensing layer constitutes the physical sensing and actuation infrastructure deployed at monitored premises. Each sensing node comprises the following hardware components:

- Microcontroller Unit (MCU): ESP32-S3 dual-core Xtensa LX7 processor operating at 240 MHz with 512 KB SRAM and 8 MB PSRAM, supporting Wi-Fi 802.11 b/g/n and Bluetooth 5.0 LE.
- Current Sensing: ACS712 Hall-effect current sensors (±30 A range, 66 mV/A sensitivity) for per-circuit current measurement with 0.5% typical accuracy.
- Voltage Sensing: ZMPT101B precision voltage transformer with resistive divider network, providing 0–250 V AC measurement with ±1% accuracy.
- Power Quality Monitoring: ATM90E32 three-phase energy metering IC providing active power, reactive power, apparent power, power factor, frequency, and THD measurements at 50/60 Hz.
- Environmental Sensors: BME680 integrated sensor providing temperature (±1°C), humidity (±3% RH), atmospheric pressure (±1 hPa), and indoor air quality (IAQ) index.

- Real-Time Clock: DS3231 TCXO-compensated RTC with ±2 ppm accuracy for precise timestamping.
- Local Storage: 32 GB microSD card for data buffering during network outages, ensuring data integrity.

Energy data is sampled at 1-second intervals for high-fidelity transient capture, with 15-minute interval aggregates transmitted to the fog layer. The node firmware, developed in C++ using the ESP-IDF framework, implements a finite-state machine for robust operation under intermittent connectivity. A watchdog timer and automatic OTA (Over-The-Air) firmware update mechanism ensure continuous, unsupervised field operation.

3.2 Fog Processing Layer

The fog layer is instantiated on Raspberry Pi 4 Model B single-board computers (4 GB RAM, quad-core ARM Cortex-A72 @ 1.8 GHz) serving as local fog nodes, each handling up to 32 edge sensing nodes. Key functions of the fog layer include:

- Data Aggregation: Collection and time-alignment of sensor streams from multiple edge nodes via MQTT (Mosquitto broker v2.0).
- Lightweight Preprocessing: Real-time signal conditioning, outlier filtering (3-sigma Winsorization), and feature computation (rolling statistics, FFT-based harmonics).
- Local Inference: Deployment of quantized ML models for sub-second anomaly detection and appliance state estimation, enabling rapid local alerting without cloud round-trip latency.
- Protocol Translation: Conversion between MQTT (local) and HTTPS REST (cloud upload) protocols with TLS 1.3 encryption.
- Edge Buffering: Local data persistence (SQLite) ensuring zero data loss during cloud connectivity interruptions.

3.3 Cloud Analytics Layer

The cloud analytics layer is deployed on Amazon Web Services (AWS), leveraging a microservices architecture for modularity and scalability. The primary cloud components are:

- Data Ingestion: AWS IoT Core handles MQTT connections from fog nodes, routing messages to AWS Kinesis Data Streams for real-time processing and Amazon S3 for long-term data lake storage.

- Stream Processing: Apache Flink on AWS Kinesis Data Analytics performs windowed aggregations, anomaly pre-screening, and feature extraction on live data streams.
- ML Training Pipeline: Amazon SageMaker orchestrates model training, hyperparameter optimization (Bayesian optimization), and A/B deployment of updated models on a weekly retraining schedule.
- Model Serving: Amazon SageMaker Endpoints host trained LSTM, Random Forest, and Isolation Forest models behind an auto-scaling API Gateway, providing <50ms inference latency at the 99th percentile.
- Relational Database: Amazon RDS (PostgreSQL) stores processed metrics, model predictions, alerts, and user configuration data.
- Dashboard Backend: Django REST Framework API serves the web dashboard and mobile application, providing time-series queries, alert management, and demand-response scheduling endpoints.

3.4 Communication Protocols

MQTT (Message Queuing Telemetry Transport) is employed for edge-to-fog communication due to its lightweight publish-subscribe model and low overhead (2-byte fixed header), making it suitable for resource-constrained IoT devices. HTTPS REST with JSON payloads is used for fog-to-cloud communication, providing compatibility with standard cloud API gateways. All inter-layer communications are secured with TLS 1.3 and certificate-based mutual authentication. Wi-Fi (IEEE 802.11n) serves as the primary connectivity medium for edge nodes, with LoRaWAN as a fallback option for remote premises.

4. MACHINE LEARNING METHODOLOGY

4.1 Data Preprocessing Pipeline

Raw sensor data undergoes a multi-stage preprocessing pipeline prior to ML model ingestion. Missing values, occurring at a rate of 0.3% in the deployment dataset, are imputed using forward-fill for durations under 5 minutes, and linear interpolation for longer gaps. Outlier detection using the Modified Z-score method (threshold $|M_i| > 3.5$) flags implausible readings for replacement. Feature normalization employs Min-Max scaling to [0, 1] for LSTM inputs and Standard Score normalization for Random Forest and Isolation Forest features.

Temporal features engineered from timestamps include: hour of day, day of week, month, public holiday indicator, and cyclic encoding (sin/cos transformations) of periodic features. Lagged consumption values (1-hour, 24-hour, 168-hour lags) are incorporated to capture autocorrelation structure. Environmental features (temperature, humidity) are included as exogenous variables, as they exhibit strong correlation (Pearson $r = 0.73$) with cooling/heating loads.

4.2 LSTM-Based Energy Forecasting

Energy consumption forecasting is formulated as a multi-step time-series prediction problem: given a historical window of $W = 168$ hours, predict aggregate consumption for the subsequent $H = 24$ hours at 1-hour resolution. The LSTM architecture comprises:

- Input Layer: $W \times F$ tensor where $F = 14$ features (consumption lags, temporal indicators, environmental variables).
- LSTM Layer 1: 256 units with hyperbolic tangent activation, returning full sequences.
- Dropout Layer: Rate = 0.25 to mitigate overfitting.
- LSTM Layer 2: 128 units with hyperbolic tangent activation, returning final hidden state.
- Dropout Layer: Rate = 0.20.
- Dense Layer: 64 units with ReLU activation.
- Output Layer: $H = 24$ units with linear activation for multi-step output.

The model is trained using the Adam optimizer (learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$) with mean squared error (MSE) loss. Early stopping (patience = 15 epochs) and learning rate reduction on plateau (factor = 0.5, patience = 7) are applied to prevent overfitting. Training is conducted on AWS SageMaker ml.p3.2xlarge instances (NVIDIA V100 GPU) with a batch size of 64 and a maximum of 200 epochs.

4.3 Random Forest Appliance Classification

Appliance identification employs a multi-class Random Forest classifier that maps aggregate load signatures to individual appliance states. The classifier operates on 2-second windows of smart meter data, extracting a 47-dimensional feature vector comprising: active power, reactive power, apparent power, power factor, current magnitude, voltage magnitude, fundamental frequency, and harmonic distortion indices up to the 11th harmonic.

The Random Forest ensemble comprises 500 decision trees with the following hyperparameters determined via 5-fold cross-validated grid search: maximum tree depth = 20, minimum samples per leaf = 5, maximum features per split = $\sqrt{F}$ (square root criterion), and bootstrap sampling with replacement. Class imbalance is addressed through SMOTE

(Synthetic Minority Over-sampling Technique) applied to minority appliance classes in the training set.

4.4 Isolation Forest Anomaly Detection

Anomaly detection targets three operational scenarios: (i) energy theft via meter tampering or unauthorized bypassing, (ii) equipment malfunction manifesting as anomalous load patterns, and (iii) data transmission errors producing implausible readings. The Isolation Forest algorithm is selected for its linear time complexity O(n log n) and superior performance under high-dimensional, multi-modal data distributions.

The model is trained on 12 months of verified normal-operation data. The contamination parameter (expected fraction of anomalies) is set to 0.02 based on domain expertise. Features include: 15-minute interval energy consumption, power factor, voltage deviation, current unbalance (3-phase premises), and contextual features (time-of-day, temperature). The anomaly score threshold is calibrated to achieve a false-positive rate below 1% on the validation set while maximizing recall of confirmed anomaly events.

4.5 Demand-Response Optimization

The demand-response module formulates load scheduling as a Mixed-Integer Linear Programming (MILP) problem, solved using the PuLP library with the CBC solver. Decision variables represent the on/off state and operating schedule of N deferrable loads (dishwashers, washing machines, EV chargers, water heaters) over a 24-hour planning horizon. Constraints encode appliance operational windows, minimum run-time requirements, and occupant comfort preferences. The objective function minimizes total energy cost under time-of-use (ToU) tariffs while respecting a peak demand contract limit, with LSTM forecasts providing the 24-hour load baseline.

5. EXPERIMENTAL RESULTS AND EVALUATION

5.1 Dataset Description

The evaluation dataset comprises 18 months (January 2023 – June 2024) of energy consumption data collected from 150 premises: 120 residential households and 30 commercial buildings across Mumbai, Pune, and Bangalore, India. A total of 3,240 IoT sensing nodes were deployed. The dataset encompasses approximately 2.1 billion measurement records totaling 840 GB of raw data. Ground-truth appliance labels were obtained for a 30-premise annotated subset using supervised monitoring during a 4-week labeling campaign.

Table 1: Forecasting Performance Comparison (MAPE %).

Model	Residential	Commercial	Overall	24-hr Ahead	Time (ms)
ARIMA	9.21%	8.34%	8.92%	12.47%	42
SVR	7.83%	6.91%	7.51%	10.23%	28
XGBoost	5.47%	4.89%	5.24%	7.81%	35
LSTM (Proposed)	3.21%	3.81%	3.47%	5.13%	67
Transformer	3.08%	3.44%	3.21%	4.89%	312

Note: Bold indicates best performance. Inference time measured on fog node hardware (Raspberry Pi 4). Transformer model excluded from fog deployment due to latency constraints.

5.2 Forecasting Results

Table 1 presents a comparative evaluation of energy forecasting models. The proposed LSTM model achieves a MAPE of 3.47% overall, representing a 61.1% improvement over the ARIMA baseline and a 33.8% improvement over XGBoost. While the Transformer model marginally outperforms LSTM (3.21% vs. 3.47% MAPE), its 312 ms fog-node inference latency exceeds the 100 ms real-time control requirement, making it unsuitable for edge deployment. The LSTM model achieves 67 ms inference latency on the Raspberry Pi 4 fog node, satisfying real-time constraints.

Seasonal analysis reveals MAPE degradation during festival seasons (Diwali, Holi) due to atypical consumption patterns, with peak MAPE of 8.3% on festival days. Incorporating a holiday indicator feature reduces festival-day MAPE to 5.7%. The model demonstrates superior performance for commercial premises during weekday business hours (MAPE = 2.8%) compared to weekends (MAPE = 5.1%), reflecting regular commercial load patterns.

5.3 Appliance Classification Results

The Random Forest classifier achieves an overall appliance identification accuracy of 96.3% across 12 appliance categories on the 30-premise annotated dataset. Per-class F1-scores range from 0.923 (refrigerator) to 0.987 (air conditioner), with lower performance on appliances exhibiting variable power signatures (iron: F1 = 0.931). Comparison with a deep learning NILM baseline (CNN-LSTM) yields comparable accuracy (96.7%) but with 4× higher computational cost, motivating the Random Forest selection for resource-constrained fog deployment.

Table 2: Anomaly Detection Performance.

Anomaly Type	Precision	Recall	F1-Score	AUC-ROC
Energy Theft	93.2%	91.8%	92.5%	0.974
Equipment Fault	95.7%	96.1%	95.9%	0.988
Data Transmission Error	98.4%	97.9%	98.2%	0.997
Overall	95.2%	94.7%	94.9%	0.986

5.4 Anomaly Detection Results

Table 2 summarizes anomaly detection performance across three anomaly categories on a held-out test set comprising 2,340 labeled anomaly events and 89,460 normal observation windows. The Isolation Forest model achieves an overall F1-score of 94.9% and AUC-ROC of 0.986, outperforming the One-Class SVM baseline (F1 = 88.3%, AUC = 0.941) and Autoencoder-based methods (F1 = 92.1%, AUC = 0.967). Equipment fault detection achieves the highest performance (F1 = 95.9%), while energy theft detection is marginally lower (F1 = 92.5%) due to sophisticated meter tampering techniques that partially mimic normal consumption patterns.

5.5 Energy Savings Evaluation

The demand-response optimization module was evaluated on 45 premises equipped with smart controllable loads over a 3-month pilot period. Participating households achieved an average energy cost reduction of 22.4% under time-of-use tariffs, with peak demand curtailment of 31.2%. Overall energy consumption reduction across the 150-premise deployment, attributable to behavioral feedback from the monitoring dashboard, averaged 18.6% (residential: 21.3%, commercial: 13.7%). These results are consistent with prior literature reporting 15–25% savings from smart monitoring interventions (Fischer, 2008), validating the real-world efficacy of ML-SEMS.

5.6 System Performance

End-to-end data latency from sensor measurement to cloud dashboard update averages 2.3 seconds under normal network conditions (edge → fog: 0.8s, fog → cloud: 1.1s, cloud processing: 0.4s). System uptime achieved 99.2% availability over the 18-month deployment, with downtime primarily attributable to planned maintenance windows. The cloud analytics platform successfully scales to process 50,000 concurrent fog node connections with <100 ms API response time (95th percentile), validated through AWS load testing.

6. DISCUSSION

6.1 Key Findings

The ML-SEMS achieves state-of-the-art performance across all evaluated sub-tasks, validating the efficacy of the integrated ML pipeline. The LSTM forecasting model's 3.47% MAPE represents a meaningful improvement over classical baselines, enabling accurate demand-response scheduling and capacity planning. The Random Forest appliance classifier's 96.3% accuracy supports non-intrusive disaggregation without per-appliance sensing infrastructure, substantially reducing deployment costs. The Isolation Forest anomaly detector's 94.9% F1-score enables proactive fault management and energy theft prevention, delivering direct economic value to utilities.

6.2 Practical Implications

From a utility perspective, ML-SEMS enables enhanced situational awareness of distribution network loads, facilitating more accurate transformer loading assessments and deferred capital expenditure on grid reinforcement. The demand-response module's 31.2% peak curtailment capability could contribute significantly to utility peak shaving programs, reducing reliance on expensive peaking generation. For residential consumers, the average 21.3% energy cost reduction provides compelling economic justification for smart monitoring adoption.

From a policy perspective, the demonstrated energy saving of 18.6% underscores the role of intelligent monitoring in achieving national energy efficiency targets. The system's cloud-based architecture enables seamless integration with emerging smart grid standards (IEC 61968, OpenADR 2.0), facilitating utility-scale deployment within existing regulatory frameworks.

6.3 Limitations and Future Work

Several limitations warrant acknowledgment. First, the dataset is geographically constrained to three Indian cities, limiting generalizability to other climatic regions and tariff structures. Future work will expand evaluation to diverse geographic and regulatory contexts. Second, the current LSTM architecture requires weekly retraining to adapt to seasonal consumption shifts; online learning approaches (e.g., reservoir computing) could enable continuous adaptation with lower computational overhead. Third, the Isolation Forest anomaly detector exhibits limited performance against sophisticated, temporally distributed energy theft attacks that gradually deviate from normal patterns; incorporation of graph neural network-based spatial correlation analysis could enhance detection capability. Fourth, privacy-preserving federated learning has not yet been implemented; future iterations will adopt federated LSTM training to enable collaborative model improvement across premises without centralizing raw consumption data.

7. CONCLUSION

This paper presented ML-SEMS, a comprehensive machine learning-based smart energy monitoring system that integrates IoT sensing infrastructure with advanced analytics to deliver real-time energy visibility, accurate consumption forecasting, non-intrusive appliance classification, and robust anomaly detection. The system's hierarchical edge-fog-cloud architecture balances real-time local responsiveness with scalable cloud-based ML analytics, achieving end-to-end data latency of 2.3 seconds and 99.2% system availability.

Experimental evaluation on 18 months of real-world data from 150 premises demonstrates that the proposed LSTM forecasting model achieves MAPE of 3.47%, outperforming ARIMA by 61.1%. The Random Forest appliance classifier attains 96.3% accuracy, while the Isolation Forest anomaly detector achieves 94.9% F1-score. The integrated demand-response module enables average energy cost reductions of 22.4%, with 18.6% overall energy savings attributable to the system's behavioral feedback mechanisms.

ML-SEMS makes a compelling case for the deployment of intelligent energy monitoring systems at scale, contributing to the realization of sustainable, efficient, and resilient smart grids. The proposed framework provides a practical blueprint for utility operators, building managers, and policymakers seeking to leverage IoT and ML technologies for energy transition objectives. Future research will focus on federated learning for privacy preservation, online learning for continuous model adaptation, and extension to industrial energy management scenarios.

REFERENCES

1. Box, G.E.P., Jenkins, G.M., Reinsel, G.C., and Ljung, G.M. (2015). Time Series Analysis: Forecasting and Control (5th ed.). John Wiley & Sons, Hoboken, NJ.
2. Chen, Y., Xu, P., Chu, Y., Li, W., Wu, Y., Ni, L., Bao, Y., and Wang, K. (2017). Short-term electrical load forecasting using the Support Vector Regression (SVR) model to calculate the demand response baseline for office buildings. Applied Energy, 195, 659–670.

3. Fang, X., Misra, S., Xue, G., and Yang, D. (2011). Smart Grid—The New and Improved Power Grid: A Survey. IEEE Communications Surveys & Tutorials, 14(4), 944–980.
4. Fischer, C. (2008). Feedback on household electricity consumption: a tool for saving energy? Energy Efficiency, 1(1), 79–104.
5. Hart, G.W. (1992). Nonintrusive appliance load monitoring. Proceedings of the IEEE, 80(12), 1870–1891.
6. Hochreiter, S. and Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780.
7. Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., and Hyndman, R.J. (2016). Probabilistic energy forecasting: Global Energy Forecasting Competition 2014 and beyond. International Journal of Forecasting, 32(3), 896–913.
8. International Energy Agency (IEA). (2024). World Energy Outlook 2024. IEA Publications, Paris.
9. Jokar, P., Arianpoo, N., and Leung, V.C.M. (2016). Electricity Theft Detection in AMI Using Customers' Consumption Patterns. IEEE Transactions on Smart Grid, 7(1), 216–226.
10. Karfopoulos, E.L., Tsekouras, G.J., Koutroumpezis, G., and Hatziargyriou, N.D. (2013). Distributed energy storage system control for energy management in microgrids. IET Generation, Transmission & Distribution, 7(6), 638–645.
11. Kelly, J. and Knottenbelt, W. (2015). Neural NILM: Deep Neural Networks Applied to Energy Disaggregation. Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments (BuildSys '15), Seoul, Korea, 55–64.
12. Kolter, J.Z. and Johnson, M.J. (2011). REDD: A Public Data Set for Energy Disaggregation Research. Proceedings of the Workshop on Data Mining Applications in Sustainability (SustKDD), San Diego, CA.
13. Kong, W., Dong, Z.Y., Jia, Y., Hill, D.J., Xu, Y., and Zhang, Y. (2019). Short-Term Residential Load Forecasting Based on LSTM Recurrent Neural Network. IEEE Transactions on Smart Grid, 10(1), 841–851.