IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)
An International Open Access, Peer-reviewed, Refereed Journal

IOT-Based Smart Energy Monitoring System For Educational Institutions

[1]Anand Hegde, [2]Arvind Bharadwaj, [3]Kunaal Pramod
[1]Student, [2]Student, [3]Student
[1]Department Of Information Science and Engineering,
[1]BMS College Of Engineering, Bengaluru, India

Under the guidance of Prof R Indra, Assistant Professor, Department Of ISE,
BMS College Of Engineering, Bengaluru, India

Abstract: This research presents a fully implemented, end-to-end smart energy monitoring system integrating machine learning, a real-time backend service, and a Flutter-based mobile application. The system utilizes the publicly available COMBED dataset to simulate real-time smart meter readings at 30-second intervals via WebSocket streaming. Four ma- chine learning models—KMeans clustering, Isolation Forest anomaly detection, LSTM short-term forecasting, and ARIMA long-term forecasting—enable comprehensive energy analysis including pattern recognition, anomaly alerts, and multi-horizon predictions.

A FastAPI backend provides REST endpoints and WebSocket support with model inference latency below 100 ms. The Flutter application delivers a mobile-first dashboard with real-time charts, anomaly lists, bill estimation in INR, and forecast visualizations. Twilio SMS integration provides automated alerts for critical anomalies.

This work distinguishes itself from prior literature through complete implementation, real-time simulation using open datasets, mobile-first architecture, and deployable multi- model integration. The system is reproducible, scalable, and extendable to real IoT smart meter hardware.

Index Terms - Component, formatting, style, styling, insert.

1. Introduction

The rapid increase in electricity demand across educational institutions has highlighted the need for intelligent, automated, and data-driven energy management systems. Traditional building energy monitoring approaches often rely on manual meter readings, delayed reports, and lack the predictive intelligence required to optimize consumption patterns. As a result, institutions face challenges such as energy wastage, rising operational costs, and the absence of actionable insights into consumption behaviors.

Recent advances in Internet of Things (IoT), smart meters, and machine learning (ML) have enabled the development of real-time energy analytics systems capable of monitoring, forecast- ing, and controlling energy use more efficiently. However, most existing research either focuses on theoretical models, works with proprietary data, or requires costly hardware deployments. This creates a gap between academic research and deployable real-world solutions.

To address these challenges, this work proposes a fully functional, software-driven smart energy monitoring system that integrates machine learning models, a modern backend archi- tecture, and a mobile-first Flutter application. The system streams high-resolution smart meter data in real time (simulated using the publicly available COMBED dataset), enabling real-time clustering, anomaly detection, short-term and long-term forecasting, and intuitive visualization for end users.

This project bridges the gap between theory and practice by implementing production-ready components such as WebSocket-based real-time streaming, SMS alerting infrastructure, model inference APIs, and a mobile app designed for practical use by students, faculty, and building managers.

1.1 Problem Statement

Educational institutions consume significant amounts of energy, yet lack accessible tools for:

- Real-time monitoring of energy consumption
- Early detection of energy anomalies and wastage
- Predictive forecasting for operational planning
- User-friendly mobile interfaces for on-the-go access
- Low-cost reproducible research platforms without hardware requirements

The absence of real-time, predictive, and mobile-accessible systems results in delayed decision-making, energy inefficiencies, and lack of actionable insights. Existing solutions are often hardware-dependent, expensive, or non-reproducible.

1.2 Objectives

This project aims to design and implement a smart energy analytics system that:

- Provides end-to-end real-time monitoring using simulated smart meter data
- Performs consumption pattern clustering using KMeans
- Detects anomalies using Isolation Forest with spike, drop, and unusual pattern classifica- tion
- Generates short-term (6-hour) and long-term (24+ hour) energy forecasts using LSTM and ARIMA
- Displays consumption trends, anomalies, and forecasts through a mobile-first dashboard
- Implements SMS-based alerting for critical anomalies
- Ensures reproducibility using publicly available COMBED dataset
- Provides a cost estimation tool using Indian electricity tariffs (Rs.160/kWh)

1.3 Scope of the Project

The scope of this work includes:

- Data preprocessing, cleaning, and feature engineering from COMBED dataset
- Implementation of four ML models: KMeans, Isolation Forest, LSTM, and ARIMA
- Design of REST API and WebSocket server using FastAPI
- Mobile application development using Flutter with three main pages:
  - Dashboard (real-time usage + anomalies)
  - Forecasting Page (6-24 hour predictions with confidence bounds)
  - Bill Calculator (INR-based cost estimation)
- Real-time simulation of 30-second interval smart meter data
- SMS alert system using Twilio for anomaly notifications

The system focuses exclusively on software implementation using open datasets, without deploying real IoT hardware. However, the architecture is designed to be extendable to real smart meter devices in future work.

2. Literature Review

A comprehensive survey of existing work in smart energy monitoring, IoT systems, and machine-learning-based energy analytics was conducted to identify advances, gaps, and research oppor- tunities. The literature spans IoT hardware systems, cloud-integrated monitoring, predictive analytics, anomaly detection, and smart campus energy management frameworks.

2.1 IoT-Based Energy Monitoring Systems

Several studies focus on developing IoT-based smart meters and monitoring infrastructures for educational institutions and commercial buildings and introduced a real-time IoT monitoring and control system for school buildings using NodeMCU and Firebase, demonstrating reductions in energy wastage. An ESP8266- based system was implemented with alert mechanisms using ThingSpeak, improving operational efficiency. A low-cost IoT power monitoring system using ESP32 and ACS712 with cloud dashboards, emphasized occupancy- based automated switching to reduce energy waste in educational institutions.

Large-scale campus deployments were explored, demonstrating 15- 20% energy savings from smart campus dashboards and wireless sensors. a campus-wide monitoring case study integrating IoT sensors for Brazilian universities.

2.2 Machine Learning for Energy Prediction and Optimization

Machine learning models have been widely applied to building energy consumption predic- tion. Elsisi et al. [**?**] used Random Forest and KNN models for load prediction in educational buildings, achieving improved energy savings via predictive control. Jain et al. [**?**] focused on AI-driven predictive alerts for smart campuses using occupancy data and plug-level monitoring.

Advanced ML-based building management systems were proposed, incorporating real-time decision-making for intelligent energy control and provided insights into machine learning approaches for building energy prediction, contrasting single vs. ensemble models.

Several review papers emphasize the growing importance of AI and IoT integration for energy efficiency, recommending hybrid models and multi-modal data integration.

2.3 Anomaly Detection and Real-Time Fault Identification

Isolation Forest, clustering, and statistical methods have been widely used for anomaly detection. Liu et al. [**?**] introduced the foundational Isolation Forest approach, which inspired several building monitoring systems. Rivera et al. [**?**] emphasized the need for scalable, multi-building IoT monitoring frameworks with anomaly detection capabilities.

Recent studies such as the edge-computing solution by NREL (COMBED dataset) [**?**] and the open-source edge architectures by international researchers highlight improved anomaly detection through fog/edge computing layers and demonstrated how building digital twins enhance anomaly detection, predictive maintenance, and real-time monitoring.

2.4 Smart Energy Management in Educational Institutions

Multiple works specifically target schools and universities like classroom- level IoT monitoring systems using DHT22 and BH1750 sensors. Mehta et al. proposed centralized cloud-IoT frameworks for campus analytics across departments. Smart city compatible BEMS architectures were explored while a low-cost smart meter was developed using cloud computing technologies.

Several papers emphasize the need for energy awareness and student engagement, including survey-based studies highlighting behavioral aspects rather than technical ones.

2.5 Identified Research Gaps

From the literature, several gaps were consistently observed:

- Limited reproducibility**:** Most research relies on proprietary datasets or hardware se- tups.
- Lack of real-time mobile dashboards**:** Existing systems predominantly use web inter- faces.
- Absence of combined analytics**:** Few systems integrate clustering, anomaly detection, and forecasting into a unified platform.
- No mobile-first design**:** Most systems do not support smartphone-based analytics.
- Limited focus on Indian electricity tariffs and local contexts.
- No integration with open datasets like COMBED for fully reproducible research.

2.6 Summary of Literature Insights

The survey shows extensive work on IoT monitoring, ML-based prediction, and anomaly de- tection, but very few solutions combine all aspects into a single deployable system. Most importantly, there is a lack of mobile-first smart energy dashboards, reproducible real-time systems using open datasets, and an integrated ML pipeline usable in practical institutional environments.The proposed work addresses these gaps by developing a complete end-to-end software platform integrating real-time analytics, machine learning, and mobile interfaces.

3. METHODOLOGY

This section details the complete methodology used in the development of the smart energy monitoring and analytics system. The methodology is divided into two major components:

(1) the software pipeline, which includes data preprocessing, machine learning model devel- opment, backend services, and mobile application design; and (2) the hardware prototype built using microcontroller-based sensing components to demonstrate real-world applicability.

3.1 Overview of System Architecture

The proposed system integrates multiple machine learning models, a FastAPI backend, and a Flutter-based mobile application. High-resolution smart meter data from the COMBED dataset is processed through the machine learning pipeline to generate clustering labels, anomaly de- tections, and short- and long-term forecasts. The backend exposes REST and WebSocket inter- faces, which are consumed by a cross-platform mobile application for real-time visualization.

3.2 Software Pipeline

The software methodology is structured into five main stages: (a) Data preprocessing, (b) Fea- ture engineering, (c) Model training, (d) Backend deployment, (e) Mobile app integration.

Data Preprocessing Pipeline

The preprocessing pipeline transforms the raw 30-second COMBED dataset into structured, model-ready input. The steps include:

1) Data Cleaning

- Remove missing timestamps and invalid consumption values
- Replace negative power values with zero
- Interpolate gaps larger than 60 seconds

Cleaning Logic

```
FOR each record in raw_dataset:
IF timestamp is NULL OR energy_consumption is NULL: MARK record for removal
IF energy_consumption < 0:
    SET energy_consumption = 0 IF timestamp delta > 60 seconds:
MARK for interpolation
```

2) Temporal Feature Engineering

From each timestamp, the following are extracted:

- Hour of day
- Day of week
- Weekend indicator
- Business-hours indicator

3) Rolling Statistical Features

- 30-minute rolling mean
- 30-minute rolling standard deviation
- Rate of change between consecutive readings

4) Normalization Min-Max scaling is applied:

$$X_{scaled} = \frac{X - X_{min}}{\max X - \min X}$$

5) Temporal Resampling To support different analysis levels:

- Minute-level averages (2 readings)
- Hourly totals (120 readings)
- Daily totals (2,880 readings)

3.2.1 Machine Learning Pipeline

Four ML models form the analytical engine of the system:

1) KMeans Clustering

- Groups power consumption into 4 clusters
- Human-readable labels: Idle, Low, Moderate, High
- Fit on single feature: power (kW)

Model Objective

$$J = \sum_{i=1}^{k} \sum_{x \in C_i} \|x - \mu_i\|^2$$

2) Isolation Forest for Anomaly Detection

- Flags anomalies as Spike, Drop, or Unusual Pattern
- Contamination = 0.05
- Operates on power (kW) alone for real-time inference

$$s(x) = 2^{\left(-\frac{E(h(x))}{c(n)}\right)}$$

3) LSTM Short-Term Forecaster

- Predicts 6-hour ahead consumption
- Sequence length = 48 (24 hours)
  - Horizon = 12 time steps LSTM cell formulation:

$$f_t = \sigma(W_f [h_{t-1}, x_t] + b_f)$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \quad \tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C)$$

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad h_t = o_t \cdot \tanh(C_t)$$

4) ARIMA Long-Term Forecaster Used for 24+ hour predictions with:

$$y_t = c + \sum_{i=1}^{p} \varphi_i y_{t-i} + \sum_{j=1}^{q} \theta_j \epsilon_{t-j} + \epsilon_t$$

3.2.2 Unified Model Training Procedure

```
PROCEDURE

TrainModels(data):

    processed_data = preprocess(data)
        X_simple = processed_data['power'] / 1000
        kmeans.fit(X_simple)
        iso_forest.fit(X_simple)

        LSTM.fit(power_series) ARIMA.fit(power_series)

        RETURN kmeans, iso_forest, LSTM, ARIMA
```

Training takes 2-3 minutes (first run) due to dataset download, and 10 seconds later due to caching.

3.3 Backend Implementation (FastAPI)

The backend provides:

- REST endpoints: /summary, /forecast
- WebSocket endpoint: /ws/live
- Twilio SMS integration
- Real-time anomaly evaluation

3.4 Flutter Mobile Application

The mobile app consists of:

1) Dashboard Page

- Real-time graph (neon theme)
- Cluster label
- Anomaly count
- Cost indicator (Rs.160/kWh)

2) Forecast Page Displays:

- Historical vs. predicted graph
- RMSE and MAPE
- 6-24 hour forecasting controls

3) Bill Calculator Page

Direct calculation using:

$$\text{Cost} = \text{kWh} \times 160$$

3.5 Hardware Prototype

Although the software system uses the COMBED dataset for experimentation, a working hard- ware prototype was developed to emulate real-time energy measurement.

3.5.1 Components Used

- Arduino Uno Microcontroller
- ESP8266 Wi-Fi module
- AC current sensor (ACS712 or equivalent)
- Multimeter module for calibration
- 230V AC bulb as load

3.5.2 Working Principle

1) Sensing The ACS712 sensor measures load current:

$$\mathbf{I} = \frac{V_{out} - 2.5}{0.185}$$

2) Power Calculation

$$P = V \times \mathbf{I}$$

3) Microcontroller Processing

- Arduino samples current at fixed intervals
- Noise is filtered using moving averages
- Power readings sent to ESP8266 over UART

4) Data Transmission

- ESP8266 transmits readings to serial monitor or cloud (for demonstration)

5) Output Display

- Laptop serial monitor shows live voltage, current, and power
- Bulb acts as load to demonstrate consumption variation

3.5.3 Role of Hardware Prototype

While the project is software-driven, the hardware prototype:

- Demonstrates how real energy meters capture data
- Bridges understanding between sensing and analytics
- Validates the concept for future IoT integration

4. Results and Analysis

This section presents the outcomes of the implemented machine learning models, real-time analytics pipeline, and mobile application interface. The evaluation includes clustering inter- pretation, anomaly detection performance, forecasting accuracy, and system-level performance metrics.

4.1 Clustering Results (KMeans)

The KMeans clustering algorithm successfully grouped energy consumption patterns into four distinct behavioral categories using a simple single-feature model (power in kW). This approach enables fast real-time inference while maintaining interpretability.

Cluster Interpretation

- **Cluster 0: Idle Usage** Represents baseline nighttime or unoccupied periods, typically between 0–1.5 kW.
- **Cluster 1: Low Usage** Early mornings or late evenings with partial occupancy, typically 1.5–3.0 kW.
- **Cluster 2: Moderate Usage** Regular working hours with stable equipment usage, typically 3.0–5.0 kW.
- **Cluster 3: High/Peak Usage** Maximum operational load or HVAC peaks, exceeding 5.0 kW.

The cluster assignments are displayed in the mobile app as human-readable labels (”Idle”, ”Low”, ”Moderate”, ”High”) updated every 30 seconds during live streaming.

4.2 Anomaly Detection Results (Isolation Forest)

The Isolation Forest model identifies unusual consumption events and classifies them into three anomaly categories: spike, drop, and unusual pattern. The model was configured with a contamination rate of 0.05, assuming approximately 5% anomalies in the dataset.

Types of Anomalies Observed

- **Off-Hours Spikes (42%)** High power usage during nighttime or weekends. Example: 145.7 kWh at 03:45 AM, well above typical baseline values.
- **Sudden Load Changes (31%)** Rapid increases or decreases in power within 30 seconds. Example: Jump from 35 kWh to 92 kWh.
- **Extended Duration High Usage (18%)** Elevated consumption sustained for more than 30 minutes.
- **Unusually Low Consumption (9%)** Below-baseline usage during active hours, indicat- ing outages or sensor issues.

4.2.2 Performance Metrics

Manually labeled validation data (200 samples) produced:

- **Precision**: 0.87
- **Recall**: 0.82
- **F1-Score**: 0.845
  - **False Positive Rate**: 6.2% The system also provides:
- Real-time anomaly detection during WebSocket streaming
- Top 5 anomalous events displayed in the Dashboard
- SMS alerts via Twilio when 5 anomalies accumulate (30-minute cooldown)

4.3 Forecasting Results (LSTM)

The LSTM time-series forecasting model was evaluated on a 7-day test set, with prediction horizons of 6 and 24 hours.

4.3.1 Short-Term Forecasting (6 Hours Ahead)

- **RMSE**: 4.73 kWh
- **MAPE**: 8.24%
- **MAE**: 3.61 kWh
- **R^2 Score**: 0.923

4.3.2 Medium-Term Forecasting (24 Hours Ahead)

- **RMSE**: 8.91 kWh
- **MAPE**: 14.67%
- **MAE**: 6.82 kWh
- **R^2 Score**: 0.874

4.3 Forecast Interpretation

Visual inspection shows:

- Predictions closely follow actual values during consistent operational periods.
- Peak loads are predicted within ±30 minutes accuracy.
- Larger forecast errors occur during transition periods and weekends due to variability.
- Forecast confidence intervals widen appropriately during volatile periods.

4.4 Mobile Application Results

The Flutter mobile application displays results through three main pages:

- **Dashboard**: Real-time consumption chart, cluster label, anomaly count, cost estimate.
- **Forecast Page**: Historical vs. predicted consumption with confidence bounds and accu- racy metrics.
- **Bill Calculator**: Direct cost estimation using INR tariff of Rs.160/kWh.

Toast notifications appear automatically when:

- A new anomaly is detected
- Usage pattern changes (e.g., Moderate → Peak)

4.5 System Performance Summary

Table 1: System Performance Characteristics

Metric	Value
Startup Time	2–3 minutes (first run), ∼10 seconds later
Real-Time Latency	< 100 ms for summary endpoint
WebSocket Update Rate	30-second intervals (real-time)
Memory Usage	500–800 MB (with models in memory)
Concurrent Users Tested	1 (scalable via load balancing)
Platform Support	Windows, macOS, Linux, iOS, Android, Web

These results demonstrate that the system performs reliably across clustering, anomaly de- tection, and forecasting tasks, providing an end-to-end analytical solution for energy monitor- ing applications.

5. Novelty and Key Contributions

This work presents several novel elements that distinguish it from existing research in smart energy monitoring, machine learning analytics, and IoT-based building management systems. Unlike prior literature, which typically focuses on isolated algorithms, hardware prototypes, or non-reproducible datasets, the proposed system delivers a complete, deployable, software- driven solution integrating multiple models, real-time analytics, and a mobile-first user inter- face.

5.1 End-to-End Practical Implementation

A key contribution of this work is the complete implementation of a production-ready smart energy analytics pipeline that includes:

- A fully functional **FastAPI backend** with REST and WebSocket support.
- A complete **multi-model ML pipeline** including KMeans, Isolation Forest, LSTM, and ARIMA.
- A **Flutter mobile application** providing real-time visualization, forecasting, billing, and anomaly alerts.
- **SMS-based alerting system** using Twilio for anomaly notifications.
- **Integrated INR support** with $160/kWh tariff, tailored for the Indian energy market.

This practical integration of backend, ML models, and a cross-platform mobile app is seldom achieved in academic works, where prototypes are often limited to theoretical descriptions or offline experiments.

5.2 Reproducible Real-Time Simulation using the COMBED Dataset

A significant innovation is the real-time simulation engine built on the publicly available COMBED dataset. The system:

- Automatically downloads and caches high-resolution smart meter data on first run.
- Streams the dataset in real time through WebSockets to mimic smart meter behaviour.
- Allows accelerated playback (0.5-second intervals) for rapid testing of real-time ML algorithms.
- Ensures reproducibility, allowing researchers and developers to replicate results without proprietary hardware.

This approach bridges the gap between hardware-dependent IoT systems and purely theoretical research, making the system accessible and deployable in software-only environments.

5.3 Mobile-First Dashboard Architecture

While most smart energy monitoring systems rely on web dashboards, this work adopts a mobile-first architecture using Flutter, providing:

- Native-level performance on Android, iOS, and Web.
- Real-time animated charts with neon-glass morphism design.
- Live clustering labels, anomaly indicators, and cost estimations.
- Offline caching for improved reliability.

This enhances accessibility for energy managers, enabling monitoring on the move rather than relying on desktop systems.

5.4 Proactive Toast Notification System

A unique feature of the system is the real-time, event-driven toast notification mechanism:

- Alerts triggered automatically for anomalies and pattern transitions.
- Severity-based color coding (red, orange, blue, green).
- Interactive notifications redirecting users to detailed anomaly views.
- Background push notifications for critical events.

Most existing systems rely on email or static threshold alerts, whereas this work provides mobile-native, low-latency alerts driven by ML outputs.

5.5 Modular and Scalable System Architecture

The backend follows a model-agnostic and microservices-ready design:

- Each ML model implements a standardized predict() interface.
- The architecture supports future decomposition into microservices.
- Data ingestion is abstracted, allowing easy integration with real IoT devices.
- The system is ready for database extensions such as PostgreSQL or TimescaleDB.

These design choices allow seamless scaling to support multi-building or campus-wide deployments.

Table 2: Comparison of Analytical Capabilities

Use Case	**Model**	**Output**
Pattern Recognition	KMeans	Idle/Low/Moderate/High clusters
Anomaly Detection	Isolation Forest	Spike/Drop/Unusual alerts
Short-Term Forecasting	LSTM	6-hour predictions
Long-Term Forecasting	ARIMA	24+ hour predictions
Cost Estimation	Statistical Model	INR-based billing

5.6 Comprehensive Analytical Coverage

Unlike prior works focusing on a single capability (clustering, forecasting, or anomaly detec- tion), this system integrates the full spectrum of energy analytics in one platform.This multi-model integration makes the system more practical and valuable for end users.

5.7 Academic and Real-World Impact Academic Contributions

- Provides a reproducible benchmark using the COMBED dataset.
- Demonstrates multi-model ML integration in real time.
- Serves as a full-stack learning resource spanning ML, backend, and mobile UI.

Practical Contributions

- Offers a deployable energy analytics system for institutions.
- Provides real-time anomaly alerts and cost estimations.
- Eliminates dependence on hardware for testing and research.

The combination of reproducibility, real-time streaming, mobile-first design, and multi- model integration forms the core novelty of this research work.

6. Conclusion and Future Work

6.1 Summary

This work presents a complete end-to-end smart energy monitoring and analysis system that integrates machine learning models, backend services, and a mobile application into a unified, fully functional software solution. Using the publicly available COMBED dataset, the system demonstrates real-time data streaming, clustering, anomaly detection, forecasting, cost estimation, and mobile visualization, without requiring specialized hardware deployments.

The major accomplishments of this work include:

- Development of a **multi-model ML pipeline** combining K-Means clustering, Isolation Forest anomaly detection, and both LSTM and ARIMA forecasting models.
- Implementation of a **FastAPI backend** with REST endpoints and real-time WebSocket streaming with $<$ 100 ms latency.
- Design of a **Flutter mobile application** providing real-time dashboards, forecasting visualizations, and a bill calculator with INR tariff integration.
- Automatic COMBED dataset integration with one-time download, caching, and rapid subsequent startups.
- Implementation of **SMS alerting** via Twilio when multiple anomalies accumulate.
  - Indian market alignment with a tariff of Rs.160/kWh and carbon footprint estimation.

Overall, the system demonstrates that sophisticated smart meter analytics can be implemented effectively using open-source tools and public datasets, providing a deployable and reproducible framework suitable for research, education, and prototype development.

6.2 System Limitations

Although the system performs efficiently in real-time scenarios, several limitations are noted:

- **Dataset dependence**: The COMBED dataset represents commercial buildings; accuracy may vary for residential or industrial environments.
- **Feature simplification**: Real-time predictions use only power (kW) as a feature, omit- ting weather, occupancy, and multi-sensor inputs.
- **Single-user load testing**: The system has been tested for individual users; large-scale deployments require additional optimization.
- **Static model training**: Models are trained once during startup; continuous learning or adaptive retraining is not yet implemented.
- **UI scope**: The current mobile UI includes three pages and displays anomalies on the dashboard but lacks a dedicated anomaly analytics screen.

6.3 Future Work

Several enhancements can be pursued to extend the functionality and real-world applicability of the system:

1) Integration with Real IoT Smart Meters

Future versions can incorporate meter hardware (DLMS/COSEM, Modbus, MQTT sensors) for live building data, enabling validation beyond simulated datasets.

2) Campus-Scale Multi-Building Deployment

A hierarchical dashboard structure (campus → building → floor → device) can support institution- wide monitoring.

3) Weather-Based Forecasting

Integrating APIs such as OpenWeatherMap can enrich forecasting, especially for HVAC-driven loads.

4) Load Shifting and Demand Response

Optimization modules can be added to recommend off-peak scheduling and participation in demand response programs.

5) Advanced Anomaly Diagnosis

Root-cause analysis using causal inference, sub-metering, or explainable AI (XAI) can auto- mate troubleshooting.

6) Energy Efficiency Benchmarking

The system can compare building performance against peer benchmarks and generate improve- ment recommendations.

7) Long-Term Trend Analysis and Reporting

Automated monthly/quarterly reporting with trend decomposition and natural-language sum- maries can be added.

8) ML Model Enhancements

Future work may explore:

- Transformer-based forecasting models (e.g., TFT)
- Ensemble prediction architectures
- Online learning for adaptive model updates
- Federated learning for multi-building privacy preservation

6.4 Concluding Remarks

This research demonstrates a complete, deployable smart energy monitoring system built en- tirely from open-source tools and publicly available data. By combining real-time analytics, anomaly detection, forecasting, mobile visualization, and user-centric features such as billing and notifications, the system serves as both a practical tool and a valuable academic contribu- tion.

The modular architecture, reproducible design, and mobile-first approach make this sys- tem suitable for further research, classroom use, and real-world prototyping. As buildings increasingly adopt smart and sustainable technologies, systems like this one provide a strong foundation for energy optimization, cost reduction, and environmental impact reduction.

Acknowledgment

The authors express their sincere gratitude to Prof. R. Indra, Assistant Professor, Department of Information Science and Engineering, B.M.S. College of Engineering, Bengaluru, for her invaluable guidance, continuous motivation, and constructive feedback.

The authors further acknowledge the National Renewable Energy Laboratory (NREL) for providing the COMBED dataset.

REFERENCES

[1] F. Maqbool, A. Shaikh, and S. Basha, “A novel smart energy monitoring and controlling system for school buildings using IoT,” International Journal of Innovative Technology and Exploring Engineering, vol. 11, no. 2, pp. 89–95, 2022.

[2] S. Nagarajan and R. Jayashree, “IoT based energy monitoring and management system for educational institutes,” International Journal of Engineering Research and Technology, vol. 9, no. 5, pp. 551–555, 2020.

[3] M. Elsisi et al., “A novel approach for energy management in educational buildings based on IoT and machine learning,” IEEE Access, vol. 9, pp. 116173–116191, 2021.

[4] A. Siddique, M. H. U. Rehman, and M. Hussain, “IoT-based smart energy management system for university campuses,” Procedia Computer Science, vol. 178, pp. 264–272, 2020.

[5] R. Jain, A. Shukla, and A. Sharma, “Energy efficiency enhancement in smart campus using AI and IoT,” Sustainable Computing: Informatics and Systems, vol. 37, p. 100738, 2023.

[6] V. Marinakis and H. Doukas, “An advanced IoT-based system for intelligent energy man- agement in buildings,” Sensors, vol. 18, no. 2, p. 610, 2018.

[7] N. Rivera et al., “An IoT-based solution for monitoring a fleet of educational buildings focusing on energy efficiency,” Sensors, vol. 18, no. 11, p. 3837, 2018.

[8] R. Manikandan and R. Ramya, “IoT-based real-time power monitoring system with cloud integration,” International Journal of Scientific Research and Engineering Development, vol. 4, no. 2, pp. 124–128, 2022.

[9] S. Mahalle and S. Kolhe, “IoT based smart energy management system for educational institutes,” in Proceedings of the International Conference on Intelligent Computing and Control Systems (ICICCS), pp. 1411–1415, 2021.

[10] J. Yasuoka et al., “IoT solution for energy management and efficiency on a Brazilian university campus - a case study,” in IEEE Global Humanitarian Technology Conference (GHTC), pp. 1–6, 2020.

[11] M. Desnanjaya et al., “Improving energy efficiency in buildings with an IoT-based smart monitoring system,” in International Conference on Electrical Engineering and Computer Science (EECS), pp. 56–60, 2023.

[12] A. S. Cespedes-Cubides and M. Jradi, “A review of building digital twins to improve energy efficiency in the building operational stage,” Energy Informatics, vol. 7, no. 1, 2024.

[13] S. Faddel, Y. Elrayes, and H. Elmarakbi, “A smart IoT-based energy management system for residential buildings,” Energies, vol. 13, no. 23, p. 6101, 2020.

[14] G. Kaur and A. Verma, “Smart energy management in educational institutions using IoT and ML: A review,” International Journal of Engineering Trends and Technology, vol. 69, no. 10, pp. 123–129, 2021.

[15] M. Lu and F. Li, “Survey on Lie group machine learning,” Big Data Mining and Analytics, vol. 3, no. 4, pp. 235–258, 2020.

[16] NREL, “Commercial Building Energy Dataset (COMBED): High-Resolution Smart Meter Data,” Available online: https://www.nrel.gov/buildings/combed. html.

[17] L. Breiman, “Random forests,” Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[18] F. T. Liu, K. M. Ting, and Z. H. Zhou, “Isolation forest,” in Proceedings of the IEEE International Conference on Data Mining (ICDM), pp. 413–422, 2008.

[19] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794, 2016.

[20] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[21] J. MacQueen, “Some methods for classification and analysis of multivariate observa- tions,” in Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, pp. 281–297, 1967.

[22] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in International Conference on Learning Representations (ICLR), 2015.

[23] S. Ram´ırez, “FastAPI: Modern Web APIs with Python,” Available at: https:// fastapi.tiangolo.com, 2024.

[24] Google, “Flutter: Build Apps for Any Screen,” Available at: https://flutter.dev, 2024.

[25] Z. Wang et al., “A review of artificial intelligence based building energy use prediction: Contrasting the capabilities of single and ensemble prediction models,” Renewable and Sustainable Energy Reviews, vol. 75, pp. 796–808, 2017.