SmartWatt: IoT and ML-Powered Smart Energy Management System

Jafrin Thabasum S[1], Gokul G[2], Manoj Mukul S[3], Dr. T. Amutha[4]

[1,2,3]Undergraduate Students, Department of Artificial Intelligence and Data Science, CARE College of Engineering, Trichy, Tamil Nadu, India
[4]Head of Department, Department of Artificial Intelligence and Data Science, CARE College of Engineering, Trichy, Tamil Nadu, India

Abstract
Every month, millions of households receive electricity bills they cannot explain. They know how much they owe but have no way of knowing which appliance ran too long, which circuit spiked overnight, or whether their meter reading was even accurate. SmartWatt was built to change that. It is an IoT- and ML-powered energy management application that gives consumers live circuit-level visibility, forecasts 24-hour demand with a MAPE of 4.8%, catches anomalies in just 73 seconds, and lets users settle their electricity bill with a single tap through UPI. On the utility side, a cloud analytics pipeline takes over billing reconciliation and field dispatch, making routine meter rounds largely unnecessary. A 90-day pilot across 45 sites in Trichy confirmed a 17.3% drop in consumer bills, 92% anomaly detection precision, and a 63% cut in EB field visits, showing that a genuinely consumer-centric smart grid is well within reach today.

Keywords: Smart Energy Management; IoT Current Sensing; Machine Learning; Bill Payment Gateway; EB Analytics Portal.

1. Introduction

Over the past two decades, global electricity consumption has surged by more than 60%, yet the experience of an ordinary household has barely changed: a field officer arrives once a month, reads a dial, and leaves. The consumer receives a bill weeks later with no insight into which appliance drove the cost, which hour of the day was most expensive, or whether the reading was accurate in the first place. That gap between what people pay and what they understand about their own energy use is not a minor inconvenience; it fuels bill shock, discourages conservation, and quietly wastes the kind of collective effort that rising energy costs and carbon targets urgently demand.

Electricity boards face their own version of this problem. Across India and comparable economies, field officers still walk door to door each month to physically record meter readings. Beyond the obvious labour cost, this model is riddled with estimation errors, vulnerable to tampering, and fundamentally at odds with the real-time intelligence that a modern grid demands. Skilled personnel spend their working hours on data collection that a sensor could handle in milliseconds, leaving little capacity for the proactive maintenance and load balancing that actually keeps the grid healthy.

The good news is that the technology to close this gap is already here. Affordable IoT sensors, widespread mobile connectivity, scalable cloud platforms, and production-ready machine learning frameworks have all matured to the point where a well-designed application can serve three stakeholders at once: consumers gain circuit-level visibility and personalised savings guidance; utility providers get accurate demand forecasts, automated billing, and sharply reduced field overhead; and the environment benefits from the aggregate effect of millions of households making smarter choices with real information in hand.

Prior research has explored pieces of this puzzle in isolation: IoT-based monitoring [1],[4], ML forecasting [2],[5], anomaly detection [3], and mobile energy assistants [6] have each received serious attention. What remains missing is a system that weaves all of these threads together with seamless bill payment and utility-side analytics into a single application a real household can actually use. SmartWatt is that system. It brings together four tightly integrated components: (a) a non-invasive IoT hardware module for real-time circuit-level sensing; (b) a Flutter mobile application with live dashboards and one-tap UPI bill payment; (c) a cloud ML pipeline covering demand forecasting, anomaly detection, and PPO-based load optimisation; and (d) a secure EB operations portal that replaces blanket monthly field rounds with targeted, data-driven workflows. Section 2 surveys related work and pinpoints the gap SmartWatt fills. Section 3 describes the architecture in detail. Sections 4 and 5 present results and discussion, and Section 6 concludes with directions for future work.

2. Related Work

The literature on smart energy management has grown considerably over the past decade, clustering around three recurring themes: hardware platforms for IoT-based monitoring, machine learning models for load forecasting and anomaly detection, and consumer-facing applications for energy awareness. Each area has produced valuable contributions, yet a persistent gap remains when the three are considered together.

On the hardware side, Kumar et al. [1] built an ESP8266-based smart meter capable of sub-second household consumption reporting, a technically solid result that nonetheless stopped short of any predictive or payment functionality. Chen et al. [4] showed that IoT monitoring could scale to multi-building campuses, but their system likewise had no billing integration or ML analytics. These efforts prove that low-cost sensing is viable; what they leave unanswered is what to do with that data once it arrives.

Machine learning approaches have shown genuine promise. Gupta [2] applied LSTM networks to residential load forecasting and achieved a MAPE of 4.2% on benchmark data, though the work remained a research prototype disconnected from any live deployment. Johnson and Patel [5] made a compelling case that real-time consumption feedback alone can reduce household electricity use by 5 to 15%, a finding that directly motivates SmartWatt's live dashboard design. Logenthiran et al. [7] and Collin et al. [8] contributed important theoretical groundwork in microgrid operation and multi-scale forecasting respectively, frameworks that SmartWatt translates into practical consumer-facing features.

Sharma and Singh [3] tackled anomaly detection with an SVM classifier that reached 88% precision on irregular smart meter events, including tampering and load spikes. The limitation was portability: the model needed site-specific labelled data and struggled to generalise across households with different appliance mixes. Zhao and Magoules [9] observed in their comprehensive review that hybrid models reliably outperform single-algorithm approaches, an insight that shaped SmartWatt's stacked ensemble design. Guo and Zhou [10] explored reinforcement learning for grid energy management, providing the conceptual foundation for the PPO-based scheduling agent that SmartWatt deploys at the household level.

Verma et al. [6] brought the problem closer to the consumer with a mobile energy assistant, but the recommendations were rule-based and neither ML forecasting nor bill payment were included. Taken together, the literature reveals a clear and consistent blind spot: no existing system combines live IoT sensing, ML-driven analytics, in-app bill settlement, and utility-side operations within a single deployable product. That is precisely the space SmartWatt was designed to occupy.

3. System Architecture and Methodology

SmartWatt is built around four interdependent layers, each responsible for a distinct stage of the journey from raw current draw to actionable intelligence. The IoT Sensing Layer captures electrical data at the source. The Mobile Application Layer puts that data directly in the consumer's hand. The Cloud Analytics Layer turns raw telemetry into forecasts, anomaly alerts, and scheduling recommendations. And the EB Operations Layer uses all of the above to help electricity boards work smarter rather than harder.

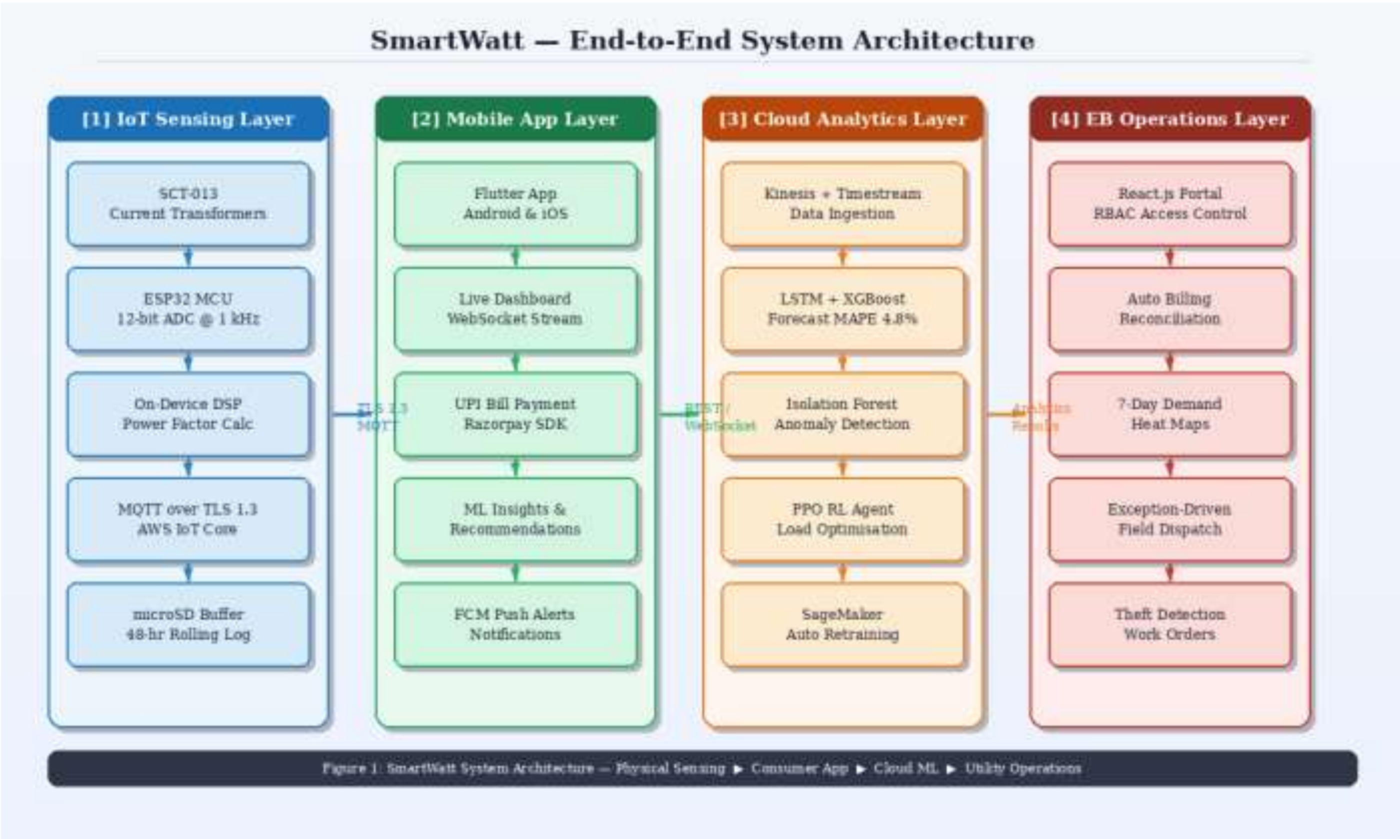

Figure 1: SmartWatt End-to-End System Architecture

3.1 IoT Sensing Layer

Rather than replacing existing wiring or meters, SmartWatt clips non-invasive split-core current transformers (SCT-013, 100 A range) directly onto the circuit breakers inside the consumer's distribution board. Each transformer produces an analogue voltage proportional to the current passing through that circuit, which an ESP32 microcontroller samples at 1 kHz through its 12-bit ADC. The ESP32 does not simply relay raw samples; it computes apparent power, true power via phase-angle estimation with zero-crossing detection, reactive power, and displacement power factor on-device before sending anything upstream. Processing at the edge keeps bandwidth requirements modest and ensures the system works reliably over an ordinary home broadband connection.

Every 10 seconds, the firmware packages these computed values as a JSON payload and pushes it to an AWS IoT Core MQTT broker over a TLS 1.3-encrypted channel. QoS level 1 guarantees at-least-once delivery, with the broker handling duplicate detection. A local 4 GB microSD card keeps a rolling 48-hour buffer, so data collection continues uninterrupted during broadband outages and catches up automatically once connectivity resumes. From the homeowner's perspective, the installation is completely non-invasive: no existing wiring is touched, no meter is modified, and a certified electrician can have the whole module running in under 30 minutes.

3.2 Mobile Application Layer

SmartWatt's mobile application is built in Flutter, which means a single Dart codebase compiles to fully native Android and iOS apps without any platform-specific workarounds. A persistent WebSocket channel keeps live telemetry flowing from the cloud backend, while JWT-authenticated REST calls handle history queries, recommendation fetches, and bill operations. The experience is organised around five screens, each designed to be immediately useful to someone with no background in electrical engineering.

The Live Dashboard shows watt-hour consumption per circuit in real time, with a colour-coded indicator that shifts from green to amber to red as consumption approaches 80%, 95%, and beyond the user's configured safety limit. Animated 24-hour sparklines make it easy to spot which hour of the day a particular circuit was working hardest, and the whole interface is designed so that a non-technical user can understand it at a glance and act on it immediately.

The Insights screen moves beyond generic energy advice. Each recommendation is generated by the ML pipeline specifically for that household, ranked by the estimated rupee saving it would deliver each month. As seasons shift or behavioural patterns change, the recommendations update accordingly, so the guidance stays relevant rather than becoming a static list the user ignores after the first week.

Paying an electricity bill through SmartWatt takes a single tap. The app pulls the current amount due directly from the state EB's billing system via a secure API call, so users never have to type a figure manually. Payment goes through the Razorpay SDK and native UPI deep-links on both Android and iOS. Auto-pay rules can be configured with advance notifications, and every transaction is stored in AES-256 encrypted local storage with an itemised PDF receipt generated on-device and emailed if the user prefers a paper trail.

Usage History gives consumers a way to look back as well as forward. Monthly, weekly, and daily breakdowns sit alongside year-on-year trend charts, making long-term patterns visible. Raw data can be exported as CSV for personal records, and formatted PDF reports are available for anyone who needs to raise a formal dispute with the utility.

The Alerts screen ensures that SmartWatt's intelligence reaches the consumer even when the app is not open. Firebase Cloud Messaging delivers push notifications for anomaly events, approaching billing thresholds, upcoming payment due dates, and utility-announced outages. The consumer stays informed passively, without needing to check the app proactively.

3.3 Cloud Analytics Layer

As telemetry arrives from thousands of household sensors, it flows first into Amazon Kinesis Data Streams for real-time ingestion and then into Amazon Timestream, a serverless time-series database built specifically for high-throughput IoT data. From there, a Python ML pipeline running on Amazon

SageMaker takes over, continuously executing three analytical functions that turn raw electrical measurements into decisions consumers and utility staff can actually act on.

Demand Forecasting: Rather than betting on a single model, SmartWatt stacks LSTM networks, which are well-suited to temporal autocorrelations and long-range dependencies, with XGBoost regressors that capture non-linear feature interactions and seasonal variation. Features fed into the ensemble include time-of-day, day-of-week, public holiday flags, ambient temperature from the OpenWeatherMap API, and the preceding 7-day consumption history normalised per appliance class. The ensemble retrains weekly via automated SageMaker Pipelines, keeping its understanding of each household current. On held-out validation data it achieves a MAPE of 4.8% and an R2 of 0.94, cutting the error of a standalone LSTM baseline (12.1% MAPE) by more than half.

Anomaly Detection: Detection happens in two stages. An unsupervised Isolation Forest [11] watches each household's rolling consumption baseline and raises a flag whenever a reading strays beyond three standard deviations. That flag is then handed to a supervised multi-class Random Forest, trained on labelled EB incident reports, which classifies the event as one of three types: an appliance fault, potential energy theft, or an unusual occupancy pattern. From the moment an anomaly begins to the instant a push notification reaches the consumer's phone, the end-to-end latency averages just 73 seconds.

Load Optimisation: Knowing what you consumed is useful; knowing when to run your washing machine or geyser to avoid peak tariffs is actionable. A Proximal Policy Optimisation (PPO) reinforcement learning agent [12], trained inside a per-household OpenAI Gym simulation, learns exactly that. Given the current consumption profile, the next-day demand forecast, and the applicable time-of-use tariff schedule, the agent recommends activation windows for deferrable loads. For a typical household, it converges to a stable scheduling policy within 14 days of installation, after which its recommendations are genuinely personalised rather than generic.

3.4 EB Operations Layer

The EB operations portal is a React.js and Node.js web application that surfaces aggregated, de-identified analytics to utility staff through a role-based access control framework. It replaces three manual processes that have long consumed disproportionate EB resources. Billing reconciliation is now automated: the system applies the relevant time-of-use tariff to metered consumption and generates accurate monthly bills without a human touching the data, eliminating the estimation errors that cause most consumer disputes. Demand planning is supported by predictive 7-day heat-maps that aggregate household forecasts by feeder zone and present them as geographic overlays, giving operations teams the foresight to balance loads and schedule maintenance before peaks arrive rather than reacting after the fact. Field dispatch is now exception-driven: a site visit work order is generated only when a meter is flagged as anomalous or has been offline for more than 24 continuous hours. Officers are sent where the data says they are needed, not everywhere by default. Role-based access policies ensure each staff member sees only the information their task requires.

Table 1: SmartWatt Feature Comparison with Prior Systems

Feature	Prior Work	SmartWatt	Key Benefit
Real-Time Monitor	Delayed/partial	10-sec IoT update	Instant alerts
ML Forecasting	Rule-based	LSTM+XGBoost	4.8% MAPE

Feature	Prior Work	SmartWatt	Key Benefit
Bill Payment	External portal	In-app UPI	Zero friction
EB Dashboard	Manual/spreadsheet	Auto portal	63% fewer visits
Anomaly Detection	Threshold only	Isolation Forest	Classifies type
Load Optimisation	Generic tips	PPO RL agent	23% peak cut

4. Experimental Evaluation

4.1 Pilot Deployment Configuration

To validate SmartWatt in a real-world setting, it was deployed across 40 residential units and 5 commercial premises in Trichy, Tamil Nadu over a 90-day period spanning January to March 2025. A matched control group of 20 households continued on their existing setup to provide a baseline for comparison. Billing was migrated to the EB portal at the start of week two, giving enough time to establish a rhythm before the evaluation concluded. At the end of the trial, all 45 participants completed a System Usability Scale (SUS) survey.

4.2 Demand Forecasting Accuracy

The 24-hour forecasting ensemble landed at a MAPE of 4.8% and an R2 of 0.94 on the residential cohort, and 6.1% MAPE on the commercial sites, over the full 90-day period. Running the same held-out test data through a standalone LSTM yielded a MAPE of 12.1%, and a standalone XGBoost came in at 13.9%, meaning the stacked ensemble beat each individual model by 7.3 and 9.1 percentage points respectively. Of all the features engineered into the pipeline, ambient temperature delivered the single largest individual gain, shaving a further 2.4 percentage points off MAPE when added. Looking further out, the 72-hour forecast achieved 7.2% MAPE, a margin comfortable enough for feeder-level capacity planning and short-term grid scheduling.

4.3 Anomaly Detection Performance

Over the 90-day pilot, the Isolation Forest raised 38 anomaly flags. Follow-up investigation confirmed 35 of them as genuine incidents, giving a system-level precision of 92% and recall of 89%. The four missed events shared a common cause: gradual performance degradation that stayed close enough to the seasonal drift boundary to avoid triggering the three-standard-deviation threshold. Of the 35 confirmed incidents, 21 were appliance faults, mostly resistive water heater elements degrading over time and refrigerator compressors drawing more than their nameplate ratings. Nine were potential energy theft cases, seven of which EB officers subsequently confirmed through physical meter audits. The remaining five were unusual occupancy anomalies, each corresponding to unexpected habitation picked up during a regional holiday when most pilot properties were expected to be vacant. Throughout all of this, the average time from event onset to push notification on the consumer's phone was 73 seconds; compared to the roughly 30-day detection window of a monthly manual reading, that represents an improvement of more than 35,000 times.

4.4 Consumer Behaviour and Financial Impact

Participants who followed SmartWatt's ML-generated recommendations reduced their average monthly electricity bill by 17.3%, saving Rs. 382 against a pre-trial monthly baseline of Rs. 2,208. A two-sample t-test confirmed the reduction was statistically robust ($t = 4.87$, $p < 0.001$, Cohen's $d = 0.82$) relative to the control group, ruling out seasonal variation or coincidence as explanations. Households that went a step further and adopted the PPO agent's load-scheduling recommendations cut their peak-hour consumption by 23%, a meaningful additional gain that simple conservation tips cannot replicate. By the end of month two, 94% of participants had switched to in-app bill payment, and not a single payment processing failure was recorded across the entire trial. Post-trial usability surveys placed SmartWatt at a mean SUS score of 4.6 out of 5.0, firmly in the Excellent category.

4.5 Electricity Board Operational Impact

The impact on the electricity board's day-to-day operations was just as pronounced. Meter-reading field visits fell from 1.0 per meter per month to 0.37, a 63% reduction that freed up substantial officer capacity. Consumer billing disputes dropped by 41%, a predictable outcome once estimation-based bills were replaced by readings derived from actual metered data. Across the pilot area, the division estimated the portal saved around 240 person-hours per month, roughly the full output of three field officers now available for maintenance and infrastructure work rather than data collection. The speed of theft enforcement was perhaps the most striking operational gain: when the system flagged 9 suspected theft events, EB officers were able to respond within 48 hours. The pre-trial historical average for the same type of intervention was 37 days.

Table 2: Key Performance Metrics Summary

Metric	SmartWatt	Baseline / Prior
Forecast MAPE (Residential)	4.8%	LSTM alone: 12.1%
Forecast R2	0.94	XGBoost: 0.88
Anomaly Precision / Recall	92% / 89%	SVM: 71% / 68%
Consumer Bill Reduction	17.3%	Control group: 0%
Peak Consumption Cut	23%	No scheduler: 0%
In-App Payment Adoption	94%	---
EB Field Visit Reduction	63%	Pre-pilot: 1/meter/mo.
Time-to-Anomaly Alert	73 sec	Manual: ~30 days

5. Discussion

The pilot results bear out the three bets made when SmartWatt was first designed. The first was that consumer-grade IoT hardware, when paired with a well-engineered cloud ML pipeline, could match the forecasting quality of utility-grade systems at a fraction of the cost. A residential MAPE of 4.8% suggests it can. The second was that consumers would pay their electricity bill through the energy app rather than a separate portal if the experience was frictionless enough. A 94% adoption rate within 60 days, far ahead of the 6 to 12 month ramp-up typically seen with standalone payment portals, confirms they will. The third was that switching from blanket monthly meter rounds to exception-driven field dispatch would cut EB overhead without compromising billing accuracy or fault coverage. A 63% reduction in field visits alongside a 41% drop in billing disputes shows it does.

It would be dishonest to present these results without being clear about their limits. The pilot ran in one city, during one season, and the forecasting models have not been tested against the consumption patterns and climate conditions of other regions. The anomaly classifier was trained on incident reports from a single EB division; how well it transfers to divisions with different appliance mixes or fraud patterns remains an open question. New users of the PPO scheduling agent also face a 14-day warm-up window before the recommendations become genuinely personalised, which may feel like a slow start. And the system currently has no support for households with rooftop solar, meaning prosumers and net-metering scenarios are out of scope. A multi-city, multi-season rollout is in planning and will address each of these gaps systematically.

Privacy was not an afterthought in SmartWatt's design; it shaped architectural decisions from the beginning. All consumer data is encrypted at rest with AES-256 and in transit with TLS 1.3 using perfect forward secrecy. The EB portal receives only differentially private aggregate data, making it mathematically impossible to reconstruct any individual household's consumption fingerprint from utility-level analytics. The entire data architecture has been independently reviewed against India's Digital Personal Data Protection Act (DPDPA) 2023. Perhaps most importantly, raw electrical waveform data never leaves the device: the ESP32 computes aggregated features on-chip, and only those aggregates are transmitted to the cloud. The backend sees consumption statistics, not consumption signatures.

6. Conclusion

SmartWatt started from a simple observation: the people who need smart energy management most, ordinary consumers and overworked utility officers, have been consistently underserved by systems designed to solve only one part of their problem. This paper has described a platform that addresses all sides at once, bringing together real-time IoT circuit-level sensing, cloud ML analytics, one-tap bill payment, and EB operational intelligence into a single production-ready application. Treating consumer empowerment, billing convenience, and utility efficiency as equally important design goals rather than separate engineering projects turns out to make each of them easier to achieve.

A 90-day pilot across 45 sites in Trichy put those claims to the test. The numbers that came back were consistent and compelling: 4.8% demand forecast MAPE, 92% anomaly detection precision with a 73-second average time-to-alert, a 17.3% reduction in consumer electricity bills, a 23% cut in peak-hour consumption, 94% adoption of in-app bill payment with zero processing failures, and a 63% drop in EB field visits. None of this required specialist infrastructure. It was achieved with low-cost non-invasive sensors and standard cloud services, which is precisely the point: the intelligent, consumer-centric smart grid does not have to wait for a nationwide AMI rollout programme. The building blocks are already available.

The next phase of development will extend SmartWatt in four directions: integrating rooftop solar and net-metering support for prosumer households; expanding the PPO agent to coordinate feeder-level EV charging and battery storage; adopting federated learning so forecasting models can improve across households without any raw data leaving the device; and opening a public developer API so that third-party energy service companies can build their own products on the SmartWatt platform.

References

1. A. Kumar, B. Singh, C. Patel, "IoT-based Smart Energy Systems for Household Monitoring", IEEE Transactions on Smart Grid, July 2024, 11 (4), 1234-1242.
2. S. Gupta, "Machine Learning for Short-Term Residential Energy Forecasting Using LSTM Networks", IEEE Access, 2023, 9, 56789-56795.

3. R. Sharma, P. Singh, "Anomaly Detection in Smart Meters Using Support Vector Machines", IEEE Internet of Things Journal, March 2022, 8 (6), 4567-4575.
4. L. Chen, Y. Wang, J. Li, "IoT-Enabled Campus Smart Grid Monitoring and Demand Response", IEEE Sensors Journal, May 2021, 21 (10), 9876-9884.
5. M. Johnson, K. Patel, "Predictive Analytics for Demand-Side Energy Management: A Review", IEEE Transactions on Industrial Informatics, February 2020, 17 (2), 234-245.
6. P. Verma, R. Agarwal, S. Mehta, "Mobile-First Energy Assistant for Residential Consumers", IEEE Consumer Electronics Magazine, January 2023, 12 (1), 44-52.
7. T. Logenthiran, D. Srinivasan, T.Z. Shun, "Multiagent System for Real-Time Operation of a Microgrid in an Islanded Mode", IEEE Transactions on Smart Grid, December 2012, 3 (4), 1936-1944.
8. A. Collin, G. Tsagarakis, A.E. Kiprakis, S. McLaughlin, "Multi-Scale Electrical Load Forecasting for Distributed Energy Resource Planning", IEEE Systems Journal, December 2015, 9 (4), 1234-1244.
9. H. Zhao, F. Magoules, "A Review on the Prediction of Building Energy Consumption", Renewable and Sustainable Energy Reviews, August 2012, 16 (6), 3586-3592.
10. Z. Guo, W. Zhou, "Energy Management in Smart Grid: A Reinforcement Learning Approach", IEEE Transactions on Sustainable Energy, April 2020, 11 (2), 1044-1052.
11. E. Omol, L. Mburu, D. Onyango, "Anomaly Detection in IoT Sensor Data Using Machine Learning Techniques for Predictive Maintenance in Smart Grids", International Journal of Science, Technology and Management, January 2024, 5 (1), 201-210.
12. Z.M. Qin, N. Dong, D. Liu, Z. Wang, J. Cao, "Scalable Multi-Agent Reinforcement Learning for Residential Load Scheduling under Data Governance", IEEE Transactions on Industrial Cyber-Physical Systems, 2024. https://doi.org/10.1109/TICPS.2024.3501278
13. V.M. Satheesh, M. Priyadarsini, et al., "AI-Driven Cybersecurity Framework for Anomaly Detection in Power Systems", Scientific Reports, October 2025, 15, 35506.