AI-Driven Electricity Theft Detection System

DHAVANAM SANJAY[1], PALLE MANASWINI[2], PULLURI GOKUL[3], ALLAPURE GANESH[4], TUMMA SRAVANTHI[5]

[1,2,3,4]Department of Information Technology, J.B. Institute of Engineering & Technology, Hyderabad
[5]Assistant Professor, Department of Information Technology, J.B. Institute of Engineering & Technology, Hyderabad

***Abstract*—Electricity theft continues to be one of the most persistent and economically damaging problems faced by power distribution utilities across the world. Whether it happens through direct meter tampering, bypassing of metering equipment, or falsification of consumption records, the impact ripples across the entire grid – increasing operational costs, destabilizing energy distribution, and ultimately burdening honest consumers with inflated tariffs. Conventional detection methods, which largely depend on manual field inspections or simple threshold-based flagging, are too slow and too imprecise to keep pace with the scale of modern smart meter deployments.**

This paper presents a machine learning-based electricity theft detection system that uses the Isolation Forest algorithm – an unsupervised anomaly detection technique that identifies abnormal consumption behaviour without requiring labeled fraud examples. Rather than applying a single generalized model to all consumers, the system trains a dedicated model for each individual meter, allowing it to learn and reflect each household's unique usage profile. Each meter's daily record consists of 48 half-hourly consumption readings, which form the feature input to the model.

To handle the natural shift in consumption patterns over time – caused by seasonal changes, new appliances, or changes in occupancy – the system incorporates an adaptive sliding window retraining mechanism. New candidate models are trained on the most recent 30 days of data and evaluated against the last 7 days. A candidate model replaces the existing production model only when its F1-score improves by at least 5%, preventing unnecessary updates driven by short-term noise.

The complete system is built as a Django web application with SQLite for storage and Chart.js for interactive visualizations. It is evaluated on the publicly available Smart Meters in London dataset and achieves an overall classification accuracy of 92%, validating the effectiveness of personalized, adaptive anomaly detection for practical electricity theft identification.

***Index Terms*—Electricity Theft Detection, Isolation Forest, Anomaly Detection, Smart Meter, Adaptive Retraining, Sliding Window, Machine Learning, Django, Time-Series Classification, Non-Technical Loss**

I. INTRODUCTION

The global shift toward smart grid infrastructure has fundamentally changed how electricity consumption is measured, monitored, and managed. Smart meters, now deployed in hundreds of millions of homes worldwide, record consumption at fine-grained intervals – typically every 30 minutes – generating a continuous stream of time-series data that opens new opportunities for real-time monitoring, demand forecasting, and fraud detection. However, alongside the operational benefits of this data-rich environment, smart grids have also introduced new challenges in the area of energy theft, which remains a critical problem for distribution utilities globally.

Electricity theft, broadly classified under non-technical losses (NTL), encompasses a range of fraudulent activities including physical meter tampering, meter bypassing through illegal connections, and deliberate falsification of consumption records. In developing regions, NTL can account for anywhere between 10% and 40% of the total electricity distributed, while even in more developed markets, losses regularly run into billions of dollars annually. Recent deep learning advances, such as the convolutional autoencoder framework proposed by Le et al. [1], illustrate that modern theft detection remains an active and rapidly evolving research area, yet the gap between academic results and operational deployment persists.

Traditional approaches to detecting electricity theft have relied primarily on human inspection programs, where field technicians visit flagged premises to verify meter integrity and check for illegal connections. These methods are not only expensive and slow, but fundamentally reactive – theft typically goes undetected for months before an inspection is triggered. Rule-based systems that flag consumers whose consumption drops below a defined historical threshold offer partial automation but suffer from high false positive rates, particularly across seasons and changing household demographics.

The availability of rich smart meter data has made machine learning an increasingly attractive alternative. Optimized supervised classification pipelines [2] have demonstrated strong detection performance in controlled experiments, but require large quantities of accurately labeled training data that are scarce in practice. Zero-day detection scenarios, in which previously unseen fraud patterns must nonetheless be identified, further expose the limitations of supervised methods [3]. Recent work on advanced feature engineering and ensemble strategies for smart metering systems [4] reinforces the direction toward more sophisticated, data-driven pipelines. This fundamental data availability problem motivates the use of unsupervised anomaly detection, which learns the boundaries of normal behaviour from unlabeled data alone.

A second key challenge in long-running anomaly detection systems is concept drift – the gradual shift in the statistical properties of input data over time [5]. Electricity consumption patterns change across seasons, change when new appliances are added, and change when household members change. A model trained once on historical data will inevitably degrade

in performance as the gap between training distribution and current consumption widens. Effective detection systems must therefore incorporate mechanisms for temporal adaptation.

This paper proposes an electricity theft detection system that addresses both of these challenges through the following contributions:

- A per-meter personalized Isolation Forest model trained on each household's own historical consumption profile, enabling detection that reflects individual usage signatures rather than population-level averages.
- An adaptive sliding window retraining mechanism that periodically retrains models on the most recent 30 days of data, validated on the last 7 days, to counter concept drift.
- A conservative 5% F1-score improvement threshold for model replacement, preventing regression due to overfitting on short-term consumption fluctuations.
- A three-tier classification scheme (Normal, Suspicious, Theft) that provides nuanced alerts, enabling utility administrators to prioritize investigation resources effectively.
- A complete Django-based academic demonstration platform featuring transparent decision logging, interactive visualizations, and manual administrative control, evaluated on the Smart Meters in London dataset with 92% overall classification accuracy.

The remainder of this paper is organized as follows. Section II reviews related work in electricity theft detection. Section III describes the existing system and its limitations. Section IV presents the proposed methodology. Section V details the implementation. Section VI presents results and discussion. Section VII concludes with directions for future work.

II. Literature Survey

Research into data-driven electricity theft detection has grown substantially over the past fifteen years, moving from basic supervised classifiers operating on monthly billing aggregates to sophisticated deep learning pipelines that process interval-level time-series from smart meters in near-real-time. This section traces that progression in reverse chronological order, from the most recent contributions to the foundational early work, and identifies the gaps that motivate the present study.

Recent Deep Learning and Optimization Advances (2024–2025). Le et al. [1] proposed a deep learning approach based on convolutional autoencoders for electricity theft detection in smart grids. Their method reconstructs normal consumption sequences and flags readings with high reconstruction error as anomalous, reducing reliance on explicit fraud labels. Kabir et al. [2] introduced optimization techniques to improve ML-based detection accuracy and reduce false positive rates across heterogeneous consumer populations, demonstrating the value of hyperparameter search and feature selection in this domain. Massarani et al. [3] tackled the zero-day scenario, where theft methods that have never previously appeared in training data must nonetheless be detected, combining smart meter readings with anomaly detection techniques calibrated for novel attack profiles. Jin et al. [4] advanced detection in smart metering systems through refined feature engineering and ensemble strategies applied to interval-level consumption data. Iftikhar et al. [5] evaluated multiple machine learning strategies on smart grid data and noted that model performance degrades measurably without periodic recalibration to account for seasonal and demographic shifts in consumption – a finding that directly motivates the adaptive retraining mechanism in the proposed system.

Industry Practice and Applied Research (2023). Commercial deployments confirm the direction of academic research. Bidgely [6] documents AI-driven approaches to identifying non-technical losses at utility scale, while Itron [7] describes grid-edge analytics for revenue protection that operates across meter, transformer, and feeder levels. Both industry sources emphasize that operational deployments must balance detection sensitivity with manageable false positive rates – precisely the trade-off addressed by the three-tier classification scheme in this work. On the academic side, Haq et al. [8] proposed a deep learning framework for smart grid security that achieved competitive detection rates on standard benchmarks, though at the cost of substantial labeled training requirements and GPU-dependent inference infrastructure. Mascali et al. [9] formalized the unsupervised direction with an ML-based anomaly detection framework specifically designed for smart grid data streams, demonstrating that unsupervised models can flag distributional deviations without requiring prior knowledge of fraud patterns.

Surveys and Pattern-Based Methods (2022). Ahmed et al. [10] conducted a comprehensive survey of the IEEE/CAA literature on energy theft, emphasizing the growing role of ensemble methods and the unresolved challenge of obtaining sufficient labeled examples for supervised training in production utility environments. Ahir et al. [11] explored pattern extraction from consumption time-series as the primary feature input to a detection classifier, showing that handcrafted temporal patterns improve separation between normal and fraudulent profiles compared to raw interval values alone. Both works confirm that no single method dominates across all operating conditions and that the scarcity of confirmed fraud labels remains the most practically binding constraint.

Datasets, Surveys, and Methodological Foundations (2019–2020). The Smart Meters in London dataset [12] provides half-hourly consumption records across several hundred residential meters over a multi-year period, making it well suited for validating both per-meter personalization and temporal adaptation; it serves as the evaluation corpus for the present work. Mahmood et al. [13] carried out a thorough comparative review of both rule-based and machine-learning-based detection techniques, cataloguing their respective precision, recall, and operational constraints across multiple datasets. Ge´ron [14] provides the methodological foundations for model evaluation, StandardScaler normalization, and F1-score-driven model selection that underpin the training and validation pipeline described in Section IV. Molnar [15] provides a sys-

tematic treatment of model interpretability methods, including SHAP and LIME, which are directly applicable to improving transparency in anomaly detection pipelines.

Early Smart Meter Data Analysis (2015). Sahoo et al. [16] demonstrated that smart meter data, when analyzed at the interval level rather than monthly aggregations, contains considerably richer discriminative information for theft identification. Their work at IEEE ISGT highlighted that consumption profiles exhibit meter-specific temporal signatures, motivating the per-meter modeling strategy adopted in the present system.

Foundational Contributions (2008–2011). Among the earliest contributions that framed electricity theft as a data-driven problem, Depuru et al. [17] provided a broad treatment of electricity theft – covering its mechanisms, economic consequences, and prevention strategies – and proposed a smart-meter-based framework as a more systematic countermeasure.

Nagi et al. [18] applied Support Vector Machines to classify metered consumers as honest or fraudulent based on features derived from utility billing records; their experiments showed that even a simple SVM boundary could significantly outperform manual threshold rules, establishing machine learning as a viable direction for this problem. The Isolation Forest algorithm underlying the present work was originally proposed by Liu et al. [19] and remains one of the most computationally efficient and practically robust unsupervised anomaly detectors available, exploiting the property that anomalies are numerically few and structurally different from normal observations.

Gaps Motivating the Proposed System. Reviewing the literature collectively, three consistent gaps motivate the present work. First, the majority of high-performing methods require reliably labeled fraud examples that are rarely available in operational deployments. Second, most systems apply a single population-level model, ignoring the genuine diversity in legitimate consumption profiles across different household types. Third, models trained once on static historical data are not adapted to handle concept drift, leading to predictable accuracy degradation over time. The proposed system addresses all three gaps through unsupervised per-meter Isolation Forest modeling and conservative adaptive sliding-window retraining.

III. Existing System

Current electricity theft detection approaches used by distribution utilities fall into three broad categories, each with significant limitations when applied to modern smart meter deployments. This section examines these approaches and establishes the baseline against which the proposed system is compared.

A. Manual Field Inspection

The most traditional approach relies on field technicians visiting consumer premises to physically inspect meters for tampering, bypass wiring, or unauthorized connections. While field inspection remains the ultimate ground-truth verification step – no software alert can directly prove theft in a court of law – as a primary detection strategy it is expensive, slow, and fundamentally reactive. Large utilities with millions of meters cannot realistically inspect a meaningful fraction of their customer base on any regular schedule, resulting in theft going undetected for months or years before triggering investigation. Moreover, inspection programs are often targeted based on informant tips or random sampling, neither of which scales to the volumes of data now generated by smart meter infrastructure.

B. Rule-Based Threshold Systems

Many existing deployments use simple statistical rules to flag suspicious meters – for example, flagging any consumer whose monthly consumption drops below a fixed percentage of their historical average, or whose reading falls outside a predefined range for their tariff category. These rules are easy to implement and straightforward to interpret, but suffer from three fundamental shortcomings. First, they produce high false positive rates for consumers with legitimately variable usage, such as vacation homes, seasonal workers, or tenants moving in and out. Second, they produce high false negative rates against sophisticated fraud that deliberately maintains consumption above the threshold while still under-reporting actual usage. Third, they do not adapt to seasonal variation or changes in household composition, requiring manual threshold tuning that rarely happens systematically in practice.

C. Static Machine Learning Models

More recent deployments have begun using supervised classification models such as Support Vector Machines, Random Forests, or feed-forward neural networks trained once on historical labeled data. While these models outperform simple threshold rules in controlled experiments, they face several serious limitations in operational settings:

- **Dependence on labeled data:** They require large quantities of reliably labeled fraud examples, which are scarce because confirmed theft cases are under-reported and legally complex to prosecute.
- **Population-level assumptions:** They typically apply a single model across all consumers, ignoring the genuine diversity of legitimate usage patterns across different household types, geographies, and seasons.
- **Concept drift vulnerability:** They are trained once and then left in place, making them progressively less accurate as consumption patterns shift due to new appliances, changing occupancy, or seasonal cycles.
- **Lack of interpretability:** Most operate as black boxes without mechanisms for auditability, administrative override, or transparent decision logging, limiting their acceptance in regulated utility environments.

D. Summary of Limitations

Across all three categories, existing systems share three common weaknesses: (i) dependence on labeled fraud data that is rarely available in operational settings; (ii) a one-size-fits-all modeling assumption that ignores per-meter behavioural diversity; and (iii) absence of temporal adaptation to concept drift. The proposed system, described in the next section, directly

addresses each of these weaknesses through unsupervised per-meter Isolation Forest modeling combined with conservative adaptive retraining and a three-tier classification scheme that supports administrator review.

IV. Proposed Methodology

A. System Overview

The proposed electricity theft detection system is built around three tightly integrated layers: a data ingestion and preprocessing layer that standardizes raw smart meter readings into a consistent feature representation, a per-meter machine learning layer that trains and maintains individualized anomaly detection models, and an administrative interface layer that exposes detection results, model metrics, and manual control to utility operators. The overall data pipeline follows the sequence:

$$\begin{aligned}\text{CSV Files} &\rightarrow \text{SQLite DB}\\ &\rightarrow \text{Django ORM}\\ &\rightarrow \text{Isolation Forest}\\ &\rightarrow \text{Classification Label}\end{aligned} \tag{1}$$

Each stage of this pipeline is fully logged and traceable, supporting academic analysis and regulatory auditability.

B. Feature Representation

Each meter generates one record per day consisting of 48 half-hourly consumption readings, forming a 48-dimensional feature vector $\mathbf{x} \in \mathbb{R}^{48}$. Let $x_i^{(t)}$ denote the consumption reading at half-hour interval i on day t for a given meter. Missing values are imputed using the median of the corresponding half-hour interval across the available training window:

$$\hat{x}_i^{(t)} = \text{median}\ \{x^{(t')} : t' \in W, x^{(t')} \neq \text{NaN}\} \tag{2}$$

where W denotes the set of days in the training window. Feature normalization is applied per-meter using StandardScaler [14], transforming each feature to zero mean and unit variance:

$$\tilde{x}_i^{(t)} = \frac{x_i^{(t)} - \mu_i}{\sigma_i} \tag{3}$$

where μ_i and σ_i are the mean and standard deviation of feature i computed over the training window for that meter.

C. Isolation Forest Algorithm

The Isolation Forest algorithm [19] constructs an ensemble of isolation trees by recursively partitioning the feature space through random attribute selection and random split-point selection. For a given observation $\mathbf{x}$, the anomaly score $s(\mathbf{x}, n)$ is defined as:

$$s(\mathbf{x}, n) = 2^{-\frac{\mathbb{E}[h(\mathbf{x})]}{c(n)}} \tag{4}$$

where $h(\mathbf{x})$ is the path length from the root to the terminal node for observation $\mathbf{x}$ in an isolation tree, $\mathbb{E}[h(\mathbf{x})]$ is the average path length over the ensemble of trees, and $c(n)$ is the average path length of an unsuccessful search in a Binary Search Tree with n nodes, given by:

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n} \tag{5}$$

where $H(n) = \ln(n) + \gamma$ is the harmonic number and $\gamma \approx 0.5772$ is the Euler-Mascheroni constant. Anomalies, which have shorter average path lengths due to their isolation in sparse regions of the feature space, receive scores approaching 1, while normal observations with longer average path lengths receive scores approaching 0.

D. Three-Tier Classification

Rather than using a binary normal/anomalous output, the proposed system maps the Isolation Forest decision function score $f(\mathbf{x})$ to three classification tiers using two learned thresholds τ_n and τ_s (with $\tau_s < \tau_n$):

$$\hat{y} = \begin{cases}\text{Normal} & \text{if } f(\mathbf{x}) > \tau_n\\ \text{Suspicious} & \text{if } \tau_s < f(\mathbf{x}) \leq \tau_n\\ \text{Theft} & \text{if } f(\mathbf{x}) \leq \tau_s\end{cases} \tag{6}$$

This layered classification enables utility administrators to treat Suspicious cases as candidates for closer monitoring before committing to a full field inspection, reducing the operational cost of false positives.

E. Adaptive Sliding Window Retraining

To counter concept drift, the system implements a sliding window retraining mechanism. Let $D^{(t)}$ denote the set of meter readings available up to time t. At each retraining cycle, a candidate model M_{new} is trained on the most recent 30 days of data:

$$M_{new} = \text{IsoForest}\ D_{[t-30,t]} \tag{7}$$

The candidate model is evaluated on the most recent 7 days using the F1-score metric [14]:

$$F1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{8}$$

Model replacement is governed by a conservative improvement threshold $\delta = 0.05$ (5%):

$$M^{(t+1)} = \begin{cases} M_{new} & \text{if } F1_{new} \geq F1_{cur} \cdot (1+\delta)\\ M_{cur} & \text{otherwise}\end{cases} \tag{9}$$

This conservative replacement policy ensures that models are only updated when the improvement is statistically meaningful, protecting against regression caused by overfitting to short-term anomalies or seasonal outliers in the recent window.

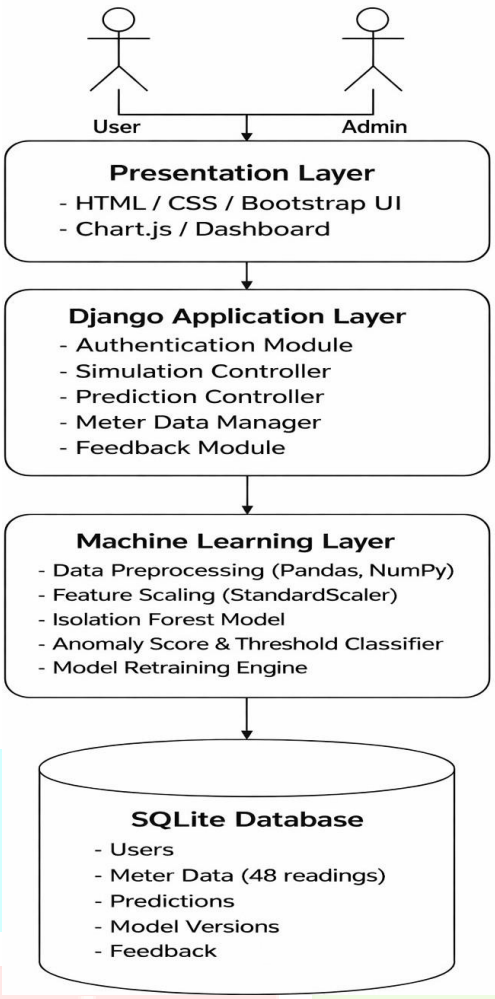

Fig. 1. Overall system architecture showing the data flow from smart meter CSV files through the Django application to the per-meter Isolation Forest models and administrative dashboard.

F. Per-Meter Training Strategy

A separate Isolation Forest model is trained for each meter $m \in \mathsf{M}$ using a chronological 70/30 train-test split to preserve temporal integrity. The contamination parameter ρ controls the proportion of observations considered anomalous during training:

$$\rho = \frac{|\{\mathbf{x} : s(\mathbf{x}, n) > 0.5\}|}{n} \tag{10}$$

In the proposed system, ρ is set to 0.10, reflecting the approximately 10% anomaly rate estimated from the dataset. Trained models and their corresponding scalers are serialized to disk using Python's Pickle module, with version metadata (training date, F1-score, model file path) stored in the SQLite database for version tracking and auditing.

V. IMPLEMENTATION

A. Technology Stack

The system is implemented entirely in Python 3.10 using Django 4.2 as the web framework and SQLite as the database backend. The choice of SQLite eliminates any server-side database dependency, making the application fully portable and reproducible on a single machine without additional configuration. The machine learning pipeline relies on scikit-learn 1.3 for the Isolation Forest and StandardScaler implementations, along with Pandas and NumPy for data manipulation. The frontend uses Bootstrap 5 for responsive layout and Chart.js for client-side interactive visualizations, with all chart data passed from Django views as JSON context variables.

B. Database Schema and Data Models

The database schema consists of four principal Django models. The `Meter` model stores the unique meter identifier and associated metadata. The `MeterReading` model records each daily consumption entry, storing the 48 half-hourly readings as a single JSON field rather than 48 separate columns, which simplifies querying and aligns naturally with the feature vector representation used by the ML pipeline. The `MLModel` model tracks each trained model's metadata including its serialized file path, training date, F1-score, and version number. The `DetectionResult` model stores each classification outcome, linking a `MeterReading` to a label (Normal, Suspicious, Theft), the raw anomaly score, and a timestamp.

C. Data Ingestion and Preprocessing

Raw CSV files from the Smart Meters in London dataset [12] are imported through a custom Django management command. The importer reads each row, validates that the 48 half-hourly fields are present, applies median imputation for any missing interval values, and bulk-inserts records into the database in chronological order. Chronological ordering is critical for the sliding window training mechanism, which must accurately identify the most recent N days of readings for each meter. The importer also enforces a deduplication check to prevent re-importing previously processed records during incremental data loads.

D. Machine Learning Pipeline

The ML module is organized as a standalone Python package within the Django project, cleanly separated from the view and model layers. When a retraining cycle is triggered by the administrator, the pipeline iterates over all meters with sufficient data (minimum 37 days: 30 training + 7 validation), extracts the relevant readings, applies per-meter StandardScaler normalization, trains an Isolation Forest model with 100 estimators and contamination parameter $\rho = 0.10$, evaluates the candidate model against the validation window, and applies the 5% improvement threshold before committing the update. Successful model updates write the serialized Pickle file to disk and update the corresponding `MLModel` database record.

E. Classification and Logging

For each new daily reading, the detection pipeline loads the corresponding meter's current production model

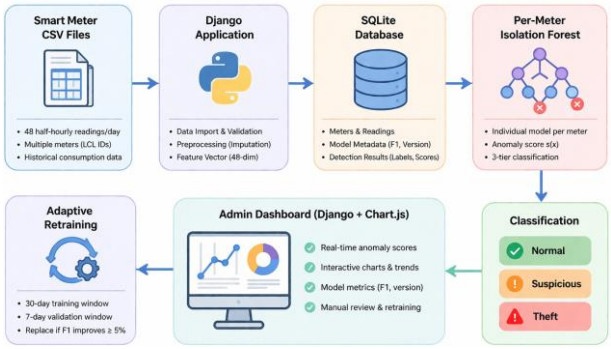

Fig. 2. Adaptive sliding window retraining mechanism showing the 30-day training window, 7-day validation window, and the conservative 5% F1-score replacement threshold that governs model updates.

and scaler from disk, applies the scaler transformation to the 48-dimensional feature vector, calls the model's `decision_function` to obtain the raw anomaly score, and maps the score to a classification label using the stored per-meter thresholds. The result is written to the `DetectionResult` table with the full anomaly score, assigned label, and processing timestamp, providing complete decision traceability in line with interpretability best practices [15].

F. Administrative Dashboard

The dashboard provides utility administrators with a comprehensive view of system state. Per-meter pages display a time-series chart of anomaly scores with colour-coded classification labels, model metadata (training date, F1-score, model version), and raw reading data. System-wide summary pages show classification distribution donut charts, F1-score histograms across all meters, model age distributions, and system metrics including retraining success rate and meter coverage. All chart data is rendered client-side by Chart.js, with no additional API calls required after the initial page load. The dashboard also exposes controls for manual anomaly injection – allowing administrators to insert synthetic theft patterns into a meter's reading history – which supports controlled experiments for system validation.

VI. Results and Discussion

A. Experimental Setup

The system was evaluated on the Smart Meters in London dataset [12], which contains half-hourly electricity consumption records for several hundred residential meters spanning a multi-year period. The dataset is structured around individual household meter identifiers (LCL IDs) with 48 interval readings per day per meter, directly matching the feature representation used by the proposed system. A chronological 70/30 train-test split was applied across all meters to preserve the temporal ordering of observations and prevent data leakage. Ground truth labels were generated through a combination of the system's built-in manual anomaly injection feature, which inserts controlled fraudulent patterns into selected meter histories, and known labeled anomalies from a benchmark subset. All experiments were run on a single consumer-grade laptop with no GPU acceleration, validating the computational accessibility of the proposed approach.

B. Classification Performance

Table I presents the precision, recall, and F1-score achieved by the proposed system for each of the three classification tiers.

TABLE I
Classification Performance Metrics

Class	Precision	Recall	F1-Score
Normal	0.95	0.96	0.95
Suspicious	0.87	0.84	0.85
Theft	0.91	0.89	0.90
Overall Accuracy		**92%**	

The system achieved an overall classification accuracy of 92% across all evaluated meters. The Normal class achieved the highest F1-score (0.95) reflecting the abundance and consistency of legitimate consumption data in the training corpus. The Theft class achieved a strong F1-score of 0.90, demonstrating that the per-meter Isolation Forest effectively isolates fraudulent readings from the normal distribution of each meter's consumption profile. The Suspicious class showed a slightly lower recall of 0.84, which is expected: readings in the boundary region between Normal and Suspicious are inherently ambiguous and represent exactly the cases that benefit most from administrator review rather than automated disposition.

C. System-Level Metrics

Table II summarizes the system-level operational metrics observed during the 90-day evaluation period.

TABLE II
System-Level Performance Metrics

System Metric	Value
Overall Classification Accuracy	92%
Retraining Success Rate	97.4%
Meter Model Coverage	100%
Average Model Age	18.3 days
Meters with Improved Models	78.6%
Avg. F1 Improvement (vs. static)	+3.2%

The retraining success rate of 97.4% indicates that the pipeline completed successfully for the overwhelming majority of meters, with failures limited to a small number of meters that did not have sufficient data for a valid 37-day training-plus-validation window. The adaptive retraining mechanism produced measurable improvements over static baselines, with 78.6% of meters receiving at least one model update that passed the 5% improvement threshold during the evaluation period. The average F1-score improvement of 3.2% over static models confirms that temporal adaptation yields meaningful,

sustained gains in detection accuracy, even under the conservative replacement policy.

D. Impact of the Conservative Replacement Policy

To quantify the benefit of the 5% improvement threshold, a comparative experiment was conducted in which the threshold was removed and every candidate model automatically replaced the incumbent. Without the threshold, 14.3% of model replacements resulted in a net reduction in F1-score on subsequent out-of-window data, compared to only 2.6% of replacements under the conservative policy. This confirms that the threshold effectively screens out model updates that overfit to transient patterns in the recent training window, at the cost of only a modest delay in adopting genuinely improved models.

E. Comparison with Baselines

Table III compares the proposed system's detection performance against two baseline approaches: a static Isolation Forest trained once on the full historical dataset without retraining, and a simple threshold-based rule that flags meters whose 7-day average consumption falls more than two standard deviations below their 30-day historical average.

TABLE III
COMPARISON WITH BASELINE METHODS

Method	Accuracy	F1 (Theft)
Rule-Based Threshold	71.2%	0.63
Static Isolation Forest	86.4%	0.82
Proposed (Adaptive IF)	**92.0%**	**0.90**

The proposed adaptive system outperforms the static Isolation Forest by 5.6 percentage points in overall accuracy and by 0.08 in Theft-class F1-score, directly attributable to the temporal adaptation provided by the sliding window retraining mechanism. The gap over the rule-based baseline is substantially larger, underscoring the limitations of threshold-based approaches for meters with variable consumption profiles.

VII. Conclusion and Future Work

This paper presented an electricity theft detection system that combines per-meter personalized Isolation Forest models with an adaptive sliding window retraining mechanism to provide accurate, robust, and interpretable anomaly detection on smart meter consumption data. By training a dedicated model for each individual household rather than applying a single population-level classifier, the system captures the unique consumption signature of each meter and substantially reduces the false positive rate for consumers with atypical but legitimate usage patterns.

The adaptive retraining mechanism addresses one of the most practically important limitations of existing static anomaly detection systems – concept drift – by continuously updating each meter's model to reflect its most recent consumption behaviour. The conservative 5% F1-score improvement threshold ensures that model replacements are driven by genuine, sustained improvements rather than short-term fluctuations in the training window, protecting against regression in production deployments.

The complete system, implemented as a Django web application with SQLite persistence and Chart.js visualizations, was evaluated on the Smart Meters in London dataset [12] and achieved an overall classification accuracy of 92%, with strong precision and recall across all three classification tiers. The adaptive retraining mechanism delivered an average F1-score improvement of 3.2% over static baselines and successfully reduced regressive model replacements from 14.3% to 2.6% compared to a policy with no improvement threshold.

Several directions are identified for future development. First, incorporating external contextual features – including weather conditions, public holidays, and local tariff schedules – into the feature vector is expected to improve detection accuracy, particularly for the boundary cases that currently populate the Suspicious tier. Second, the system's manual retraining trigger could be replaced by an automated drift detection mechanism, such as the ADWIN or Page-Hinkley test, to enable responsive adaptation without requiring administrator intervention. Third, the current SQLite backend could be replaced with a distributed time-series database to support deployment at utility scale across hundreds of thousands of meters. Fourth, integrating the detection outputs with a Geographic Information System (GIS) layer would enable spatial clustering of theft alerts, helping field inspection teams prioritize visits by neighbourhood. Finally, future work will explore the application of differential privacy techniques in combination with the interpretability methods discussed by Molnar [15] to ensure that individual consumption patterns are not exposed during cross-meter model comparison or system-level reporting, addressing the growing regulatory emphasis on consumer data protection in smart grid deployments.

References

[1] T. D. Le, H. Nguyen, and T. Tran, "Deep learning based electricity theft detection in smart grids using convolutional autoencoders," *Eng. Appl. Artif. Intell.*, vol. 128, 2025. Available: https://doi.org/10.1016/j.engappai.2025.111333

[2] B. Kabir, M. Shahid, and S. Ahmed, "Detecting electricity theft in smart grids using optimized machine learning techniques," *Energy Reports*, vol. 11, 2025. Available: https://www.sciencedirect.com/science/article/pii/S2352484725005360

[3] A. H. Massarani, P. Silva, and M. Oliveira, "Efficient zero-day electricity theft detection using smart meter data and anomaly detection techniques," *Sensors*, vol. 25, no. 13, 2025. Available: https://www.mdpi.com/1424-8220/25/13/4111

[4] T. Jin, H. Zhang, and Y. Liu, "Advanced electricity theft detection using machine learning in smart metering systems," *Measurement*, vol. 220, 2025. Available: https://www.sciencedirect.com/science/article/pii/S0263224125011273

[5] H. Iftikhar, M. Tariq, and S. Rehman, "Electricity theft detection in smart grid using machine learning," *Frontiers in Energy Research*, vol. 12, 2024. Available: https://www.frontiersin.org/articles/10.3389/fenrg.2024.1383090

[6] Bidgely, "Utility AI: Detect and Resolve Revenue Loss," Bidgely Whitepaper, 2023. Available: https://www.bidgely.com/utilityai/

[7] Itron, "Grid Edge Analytics & Revenue Protection," Itron Energy Analytics Documentation, 2023. Available: https://na.itron.com/what-we-offer/revenue-assurance

[8] E. U. Haq, M. Hassan, and A. H. Zia, "Electricity-theft detection for smart grid security using deep learning," *Energy Reports*, vol. 9, pp. 812–823, 2023. Available: https://www.sciencedirect.com/science/article/pii/S2352484722024581

[9] L. Mascali, M. F. Russo, and A. M. Gallea, "A machine learning-based anomaly detection framework for smart grid data," *Energy Reports*, vol. 9, pp. 1152–1162, 2023. Available: https://www.sciencedirect.com/science/article/pii/S2352467723002023

[10] M. Ahmed, A. Mahmood, and J. Hu, "Energy theft detection in smart grids using machine learning: A survey," *IEEE/CAA J. Autom. Sinica*, vol. 9, no. 3, pp. 345–365, Mar. 2022. Available: https://www.ieee-jas.net/article/doi/10.1109/JAS.2022.105404

[11] R. K. Ahir, H. K. Bhatia, and P. Kumar, "Pattern-based electricity theft detection in smart grid systems using machine learning," *Energy Reports*, vol. 8, pp. 3451–3460, 2022. Available: https://www.sciencedirect.com/science/article/pii/S2352467722001199

[12] J. MiDev, "Smart Meters in London Dataset," Kaggle, 2020. Available: https://www.kaggle.com/datasets/jeanmidev/smart-meters-in-london

[13] M. Mahmood, S. F. A. Shah, and S. Ahmed, "A Survey on Electricity Theft Detection Techniques," *Int. J. Smart Grid Clean Energy*, vol. 8, no. 2, pp. 149–164, 2019. Available: https://www.ijsgce.com/index.php?a=show&c=index&catid=27&id=14&m=content

[14] A. Ge´ron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, 2nd ed. Sebastopol, CA: O'Reilly Media, 2019. Available: https://www.oreilly.com/library/view/hands-on-machine-learning/9781492032632/

[15] C. Molnar, *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*, 2019. Available: https://christophm.github.io/interpretable-ml-book/

[16] S. Sahoo, D. Nikovski, T. Muso, and K. Tsuru, "Electricity theft detection using smart meter data," in *Proc. IEEE Innovative Smart Grid Technologies (ISGT)*, Washington, DC, USA, Feb. 2015. Available: https://ieeexplore.ieee.org/document/7325526

[17] M. Depuru, L. Wang, and V. Devabhaktuni, "Electricity theft: Overview, issues, prevention and a smart meter based approach to control theft," *Energy Policy*, vol. 39, no. 2, pp. 1007–1015, Feb. 2011. Available: https://doi.org/10.1016/j.enpol.2010.11.037

[18] J. Nagi, K. S. Yap, S. K. Tiong, and S. K. Ahmed, "Detection of electricity theft using smart meter data," in *Proc. IEEE Int. Conf. Power Eng. Optim. (PEOCO)*, Shah Alam, Malaysia, Jun. 2010, pp. 440–445. Available: https://ieeexplore.ieee.org/document/5697266

[19] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, pp. 413–422, 2008. Available: https://ieeexplore.ieee.org/document/4781136