Journal of Engineering, Science and Technology
ISSN: 2714-2396, Volume 9 Issue 1

FEDERATED LEARNING AND BLOCKCHAIN-ENHANCED SMART METERING FOR ELECTRICITY THEFT DETECTION IN RESIDENTIAL AREAS: A CASE STUDY OF GWAGWALADA DISTRIBUTION SYSTEM

Tijani Monsuru Abiodun[1], Eronu Emmanuel[1], Sadiq Abubakar U[1]

[1]Department of Electrical and Electronics Engineering, University of Abuja.

Corresponding email: tijani.monsuru@yahoo.com

Received: 9 April 2026 **Accepted**: Date 23 May 2026 **Published**: 16 June 2026

Abstract. Electricity theft has remained one of the major causes of non-technical losses in electric power distribution systems, particularly in developing nations where monitoring infrastructure and regulatory enforcement are often inadequate. In recent years, the growing use of artificial intelligence in smart energy systems has introduced new opportunities for electricity consumption analysis and theft detection. However, concerns relating to data privacy, confidentiality, and cybersecurity have continued to limit the adoption of centralized intelligent systems in residential power networks. Therefore, this study presents a secure smart metering framework that integrates federated learning and blockchain technology for effective electricity theft detection within residential distribution systems. The study make use of and evaluated four machine learning models consisting of hybrid CNN and GRU, CNN only, GRU only, and XGBoost using smart meter data obtained from fifty households between 2015 and 2024. It was revealed from the comparative analysis that the XGBoost model achieved the best performance under centralized training conditions with a Root Mean Square Error value of 1.52 kWh. Statistical evaluation further confirmed the superiority of the model with a significance value of p equals 0.031. Nevertheless, when the same model was implemented within a federated learning environment involving fifty communication rounds and partial client participation, the Root Mean Square Error increased to 6.77 kWh. So, this result then reflects the practical compromise that exists between maintaining user privacy and achieving maximum prediction accuracy in distributed intelligent systems. The framework proposed in this study contributes to addressing important limitations associated with existing artificial intelligence based smart metering systems, especially in the areas of data privacy, system security, and practical implementation within semi urban residential communities. The findings of this study demonstrate that the integration of federated learning with blockchain technology can provide a reliable, secure, and scalable approach for electricity theft detection and electricity consumption monitoring in modern residential power distribution networks.
.

Keywords: Electricity theft detection; Federated learning; Blockchain technology; Smart metering; XGBoost (Extreme Gradient Boost).

1. Introduction

The continuous advancement of smart grid technology has brought remarkable changes to the operation and management of modern electric power distribution systems. Over the years, power utility companies across the world have increasingly adopted intelligent monitoring and communication technologies to improve energy efficiency, strengthen system reliability, and support effective real time decision making within distribution networks. One of the most important developments in this transition is the deployment of smart meters through Advanced Metering Infrastructure systems, which enable electricity providers to continuously track energy consumption patterns and manage electricity distribution more efficiently. The integration of these technologies has played a major role in the modernization of power systems by supporting automated meter reading, rapid fault detection, demand response management, and improved communication between electricity consumers and utility providers.

Despite these improvements in smart grid, the growing digitalization of electricity networks has also introduced serious threats relating to data privacy, cybersecurity, and system security. Smart meters generate large volumes of consumer data which contain detailed information about the electricity usage patterns within households and commercial facilities. While such information is valuable for intelligent energy management and system optimization, unauthorized access to these data may expose sensitive consumer information and compromise user privacy. In addition, modern smart grid systems are increasingly vulnerable to cyber threats, malicious attacks, and unauthorized manipulation of energy data, particularly in decentralized electricity environments where large numbers of connected devices exchange information continuously. Consequently, there is an increasing demand for secure and reliable frameworks capable of protecting user data while maintaining transparency, accountability, and operational trust within modern power systems (Aitzhan and Svetinovic, 2018).

Electricity theft has continued to represent one of the most critical challenges confronting power distribution companies, especially in developing countries which a good example is Nigeria. The problem contributes substantially to non-technical losses within electricity distribution networks through activities such as illegal electricity connections, meter bypassing, meter tampering, and manipulation of billing records. These practices not only reduce the revenue generated by electricity providers but also weaken the stability, efficiency, and reliability of power supply systems (Aderibigbe et al., 2021; Momoh and Onime, 2020). In many developing regions, the persistent occurrence of electricity theft has negatively affected investment in power infrastructure and has further complicated efforts aimed at improving electricity access and service delivery. Conventional approaches used for electricity theft detection, including manual inspections and rule-based monitoring systems, have proven to be inefficient, time consuming, expensive, and often unreliable when deployed across large residential distribution networks.

In response to these challenges, researchers have increasingly explored the application of artificial intelligence techniques for electricity theft detection within smart grid systems. Machine learning and deep learning approaches have demonstrated considerable success in identifying abnormal electricity consumption patterns and detecting fraudulent activities from large smart meter datasets (Nwosu, Aneke and Okafor, 2022; Omagbemi and Adewole, 2021). Several studies have applied methods such as support vector machines, ensemble learning techniques, neural networks, and hybrid optimization models to improve the accuracy of theft detection systems (Omagbemi, Adewole and Akande, 2020; Akinlabi, Ighravwe and Adaramola, 2021; Akinlabi, Ighravwe and Adaramola, 2022). Although these methods have achieved promising results, most existing approaches depend heavily on centralized learning architectures in which consumer data from multiple users are collected and processed within a central server. This centralized approach raises major concerns regarding data privacy, confidentiality, and cybersecurity because sensitive consumer information must be transmitted and stored in a single location. Furthermore, centralized systems remain highly vulnerable to cyberattacks, unauthorized access, and single point system failures that may compromise the integrity of the entire network.

Recent developments in federated learning has introduced a more secure and privacy conscious approach to intelligent model training within distributed environments. Federated learning enables multiple devices or local nodes to collaboratively train machine learning models without transferring raw consumer data to a central server. Instead, only model parameters and updates are exchanged during the training process, thereby reducing the exposure of sensitive user information and improving data privacy. This decentralized learning approach has attracted growing attention in smart grid applications because of its ability to support intelligent analytics while maintaining confidentiality of consumer data (Zafar et al., 2023). Nevertheless, federated learning alone does not completely eliminate concerns relating to transparency, trust, and protection against malicious model manipulation during collaborative training processes.

To strengthen the security and reliability of federated learning systems, blockchain technology has emerged as a suitable complementary solution. Blockchain provides a decentralized and immutable digital ledger capable of securely recording transactions and validating information exchanges among participating nodes. Through its transparent and tamper resistant structure, blockchain technology can improve accountability, strengthen trust among distributed participants, and protect federated learning systems against unauthorized modifications and malicious updates (Okafor, Aneke and Okafor, 2020; Okafor, Aneke and Okafor, 2021). The integration of blockchain technology with federated learning therefore presents a promising approach for developing secure and privacy preserving electricity theft detection systems within modern smart grid environments.

Against this background, this study proposes a secure framework that integrates federated learning and blockchain technology for electricity theft detection within residential electricity distribution systems, using the Gwagwalada distribution network as a case study. The proposed framework combines the predictive capability of advanced machine learning models with the privacy preservation advantages of federated learning and the security features provided by blockchain technology. By addressing critical issues relating to prediction accuracy, consumer privacy, cybersecurity, and system trust, the study seeks to contribute toward the development of reliable, scalable, and secure intelligent energy management systems suitable for semi urban and developing environments (Shahzadi, Javaid and Akbar, 2024; Sharma and Tiwari, 2024; Nawaz et al., 2023; Ezeji, Okoro and Aginam, 2024).

2. Materials and Method

2.1 Study Area and Data Collection

The study area is the case of Gwagwalada distribution system in Abuja, Nigeria. The smart meter data was obtained from 50 residential households under the Abuja Electricity Distribution Company (AEDC), which cover the period of 2015 to 2024. The dataset attributes include electricity consumption records, timestamps, and derived features for theft detection.

2.2 Data Pre-processing and Feature Engineering

For this study, the dataset was preprocessed thereby cleaning missing values, normalization, and transformation of the dataset into structured time-series format. The feature engineering techniques used include Lag-based consumption features, Rolling mean and seasonal indicators, Consumption deviation metrics, Clustering-based household grouping. These features enabled effective learning of consumption behavior patterns and anomaly detection.

2.3 Model Development

Four models were developed and evaluated for this study:

(i) XGBoost Model

Extreme Gradient Boosting (XGBoost) was used due to its efficiency in handling structured data and its robustness to overfitting. The eXtreme Gradient Boosting (XGBoost) model was implemented as the traditional machine learning baseline. XGBoost is an optimised distributed gradient boosting library designed to be highly efficient, flexible, and portable. Its foundation lies in iteratively adding new regression trees that optimally minimise a regularised objective function.

(ii) CNN-GRU Hybrid Model

The hybrid model combines CNN for spatial feature extraction and GRU for temporal dependency learning. The Convolutional Neural Network (CNN) - Gated Recurrent Unit (GRU) Hybrid Model was developed to leverage the strengths of both architectural types. The CNN layer was positioned first to perform local feature extraction across the time steps in the input window. This allowed the model to automatically learn optimal local patterns, such as short-term fluctuations or intra-day seasonality. The CNN model extracts localized patterns in consumption data.

The outputs of the CNN layer were then passed to the GRU layer. The GRU component, which is a highly efficient variant of the Recurrent Neural Network (RNN), processed these extracted features sequentially, capturing the long-term temporal dependencies and memory effects necessary for accurate forecasting. The GRU model captures sequential dependencies in electricity consumption data. A final dense layer with a single output unit and a linear activation function was used to produce the final prediction.

2.4 Model Evaluation Metrics

The model performance was evaluated using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Cross-validation (5-fold), and Wilcoxon signed-rank test for statistical significance. The Root Mean Squared Error (RMSE) was chosen as the primary metric because it strongly penalised large errors, which was critical for ensuring accurate forecasting of peak demand periods. The RMSE was calculated as the square root of the average of the squared errors:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y_i})^2} \tag{1}$$

The Mean Absolute Error (MAE) provided a measure of the average magnitude of the errors, giving a more interpretable view of typical prediction accuracy without the squaring effect of RMSE. This metric offered a linear score of the prediction errors [12].

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y_i}| \tag{1}$$

The Mean Absolute Percentage Error (MAPE) was included to give a relative measure of accuracy, expressing error as a percentage of the actual value. This was useful for comparing performance across different datasets or time series with varied scales [13].

$$\text{MAPE} = \frac{100\backslash\%}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y_i}}{y_i}\right| \tag{2}$$

Finally, the Coefficient of Determination (R^2) provided a statistical measure of how well the model's predictions approximated the real data points. It quantified the proportion of the variance in the dependent variable that was predictable from the independent variables, with values closer to 1.0 indicating superior fit [14].

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \quad (3)$$

The statistical evaluation of the developed models was conducted using a household level cross validation approach in order to ensure reliable and generalized performance assessment. In this procedure, the complete dataset was divided into several folds, with each fold containing a unique group of households. The models were repeatedly trained and tested across the different folds, allowing the study to obtain not only the average performance values but also the corresponding standard deviation values for each evaluation metric. This validation strategy enhanced the robustness of the analysis by ensuring that model selection was based on consistent performance across varying household electricity consumption behaviours rather than dependence on a single dataset partition.

2.5 Federated Learning Framework

The Federated Averaging algorithm was adopted to coordinate the decentralized model training process within the federated learning environment. The system architecture consisted of a central server and multiple client nodes, where each client node represented either an individual household or a localized data center. Following the comparative performance evaluation of the developed models, the XGBoost model, which demonstrated the best predictive capability, was selected for deployment within the federated learning framework. The training process was conducted over fifty communication rounds with twenty percent client participation in each round, while model aggregation was achieved using the Federated Averaging technique. This framework enabled decentralized model training without requiring the transfer of raw consumer data to a central location, thereby improving data privacy and reducing potential security risks associated with centralized learning systems.

2.6 Blockchain Integration

Blockchain technology was incorporated into the proposed framework to strengthen system security, transparency, and data integrity throughout the collaborative training process. A Ganache based Ethereum blockchain environment was implemented to securely record model updates, preserve cryptographic hashes of the global models, and maintain tamper resistant audit records. In addition, smart contracts were utilized to regulate the validation and aggregation of model updates submitted by participating client nodes. This integration provided an additional layer of trust and accountability within the federated learning framework by ensuring that all training activities and model transactions remained secure, traceable, and resistant to unauthorized modification.

3. Results and Discussion

3.1 Centralized Model Performance

Four different machine learning models were developed, trained, and evaluated in order to determine the most suitable approach for electricity theft detection and electricity consumption prediction within the residential smart metering network in Gwagwalada. The models were trained using electricity consumption data obtained from fifty households grouped into five distinct clusters. To ensure a comprehensive assessment of model performance, several evaluation metrics were employed, including Root Mean Squared Error, Mean Absolute Error, Coefficient of Determination, and Mean Absolute Percentage Error. These performance indicators provided a detailed basis for comparing the predictive accuracy, reliability, and generalization capability of the developed models across varying household electricity consumption patterns.

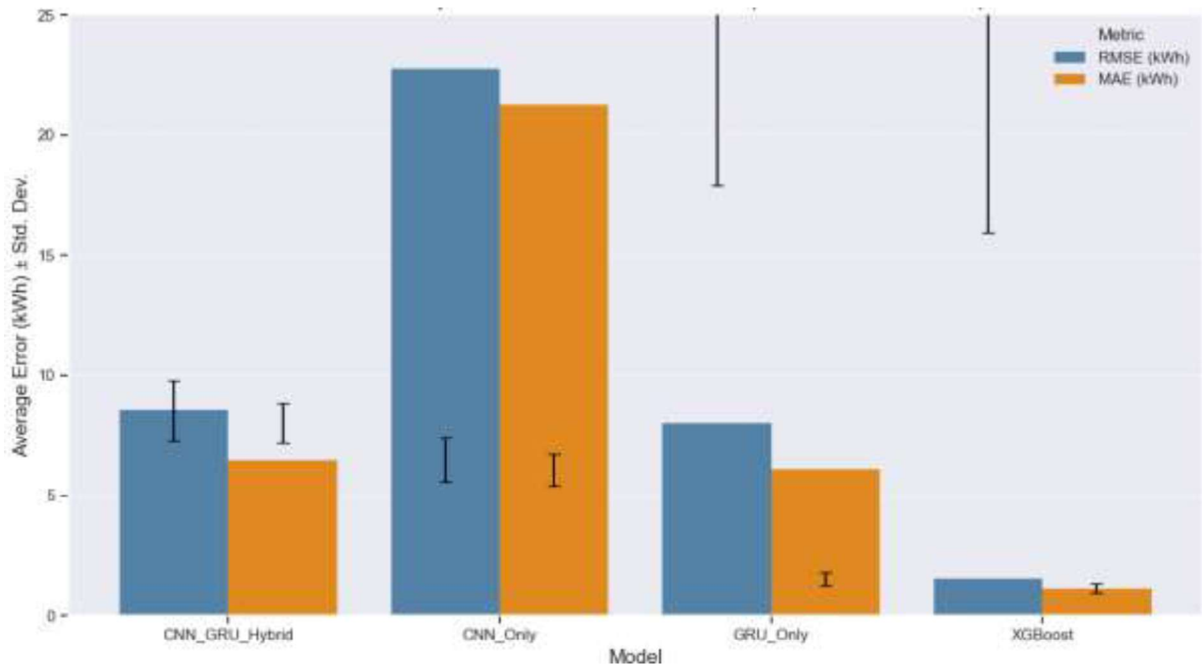

Figure 1 **Statistical Comparison of Model Performance (Error Bars = Std. Dev.)**

Among the evaluated models, XGBoost delivered the strongest overall performance, achieving a Root Mean Squared Error of 1.52 kWh and outperforming the CNN, GRU, and CNN GRU architectures. Further statistical validation using the Wilcoxon signed rank test confirmed that the superiority of XGBoost was statistically significant, with a p value of 0.031, thereby indicating meaningful differences in predictive accuracy among the compared models.

Although the deep learning-based models demonstrated an improved ability to capture complex temporal patterns in the data, their performance was comparatively lower. This can be attributed to their higher computational demands as well as their increased sensitivity to variations within the dataset, which affected their overall stability and predictive consistency.

Table 1: Comprehensive Model Performance Comparison across Multiple Metrics

Model	Avg RMSE (kWh)	Std RMSE (kWh)	Avg MAE (kWh)	Std MAE (kWh)	Avg R^2	Std R^2	Avg MAPE (%)	Std MAPE (%)	Avg Epochs	RMSE Rank	Performance Index
XGBoost	1.52	0.29	1.13	0.21	0.962	0.009	4.65	0.82	499.0	1	1.80
GRU_Only	8.02	0.84	6.09	0.66	-0.066	0.071	23.70	4.38	15.0	2	8.86
CNN_GRU Hybrid	8.53	1.24	6.47	0.91	-0.211	0.215	24.39	4.48	16.0	3	9.77
CNN_Only	22.76	4.86	21.28	5.33	-8.169	3.518	77.79	19.13	15.0	4	27.62

The results presented in Table 1 above demonstrate a clear and consistent pattern of model performance across all evaluation metrics, which is also reflected in the trends observed in Figure 1 above. Among all the evaluated approaches, XGBoost recorded the strongest overall performance as it achieved a Coefficient of Determination value of 0.962, indicating that the model was able to explain approximately 96.2 percent of the variability in electricity consumption. The very low standard deviation of 0.009 further confirms the stability and reliability of its predictions across different evaluation runs. In addition, the model attained a Mean Absolute Percentage Error of 4.65 percent with a variation of plus or minus 0.82 percent, which reflects a high level of predictive accuracy across different consumption ranges. With a performance index of 1.80, XGBoost clearly outperformed all other models and established itself as the most dependable approach for electricity consumption prediction in this study.

In contrast, the deep learning based models demonstrated relatively poor predictive performance across the same evaluation framework. The GRU only model produced a negative Coefficient of Determination value of minus 0.07, alongside a relatively high Mean Absolute Percentage Error of 23.70 percent, indicating limited predictive effectiveness. Although this model converged quickly after 15 training epochs, the results suggest that it likely suffered from underfitting and was unable to adequately capture the underlying data patterns. Similarly, the hybrid CNN GRU model performed even worse, with a Coefficient of Determination value of minus 0.21 and a Mean Absolute Percentage Error of 24.39 percent, showing that increasing architectural complexity did not translate into improved predictive capability for this dataset. The CNN only model recorded the weakest performance overall, with a highly negative Coefficient of Determination value of minus 8.17 and an extremely high Mean Absolute Percentage Error of 77.79 percent, confirming that it is not suitable for this specific application.

Overall, these findings indicate that traditional machine learning, particularly XGBoost, provides a more effective and reliable solution for electricity consumption prediction in this study compared to the deep learning architectures evaluated.

3.2 Federated Learning Performance

The final study architecture of the study prompted a strategic shift which resulted into a significant performance advantage of the centralized XGBoost model (RMSE = 1.52 kWh) over the best-performing deep learning alternative (GRU-only, RMSE = 8.02 kWh). In other to achieve optimal predictive accuracy, the research adopt a non-standard federated learning approach known as Federated Gradient Boosting (FGB), which enabled the integration of the superior tree-based XGBoost model into a distributed environment by employing a secure tree aggregation mechanism, thereby overcoming the limitations of traditional FedAvg. The FGB framework was implemented over 50 communication rounds, with 20% client participation (10 households) per round to balance computational efficiency and data heterogeneity. Key performance milestones showed an initial RMSE of 10.82 kWh in Round 1, a best performance of 3.96 kWh in Round 4, and a final converged RMSE of 6.77 kWh by Round 50.

Table 2: Federated Learning Training Performance Milestones

Training Round	Selected Clients	Evaluation Client ID	RMSE (kWh)	Notable Observation
Round 1	10	27	10.82	Initial baseline performance
Round 4	10	11	3.96	Best performance achieved
Round 30	10	31	8.17	Stabilization phase
Round 43	10	32	13.65	Maximum error observed
Round 50	10	31	6.77	Final converged performance

The training process has demonstrated effective collaborative learning across distributed clients, which results in a substantial reduction in prediction error of approximately 37.48% relative to the

initial baseline. Despite this overall improvement, the training history revealed noticeable fluctuations in model performance throughout the communication rounds. The Root Mean Squared Error values varied between 3.96 and 13.65 kWh, with a standard deviation of 2.11 kWh, which indicate a moderate variability in the learning process across different training iterations.
These oscillations can be attributed to several inherent characteristics of the federated learning framework. In particular, the random selection of participating clients in each round introduced variations in the underlying data distribution, as different subsets of households contributed to model updates at different times. In addition, the aggregation of tree based models, each with distinct decision structures, added further complexity to the learning process. The presence of non independent and identically distributed data across clients also contributed significantly to model instability, especially during the later stages of training when the model began to converge under heterogeneous data conditions.
Notwithstanding these challenges, the overall learning trend remained positive. The model exhibited rapid improvement during the early training rounds, followed by a period of relatively stable performance in the middle stages. Although increased variability was observed toward the final rounds, the general trajectory still reflected consistent progress in predictive accuracy across the federated learning process.

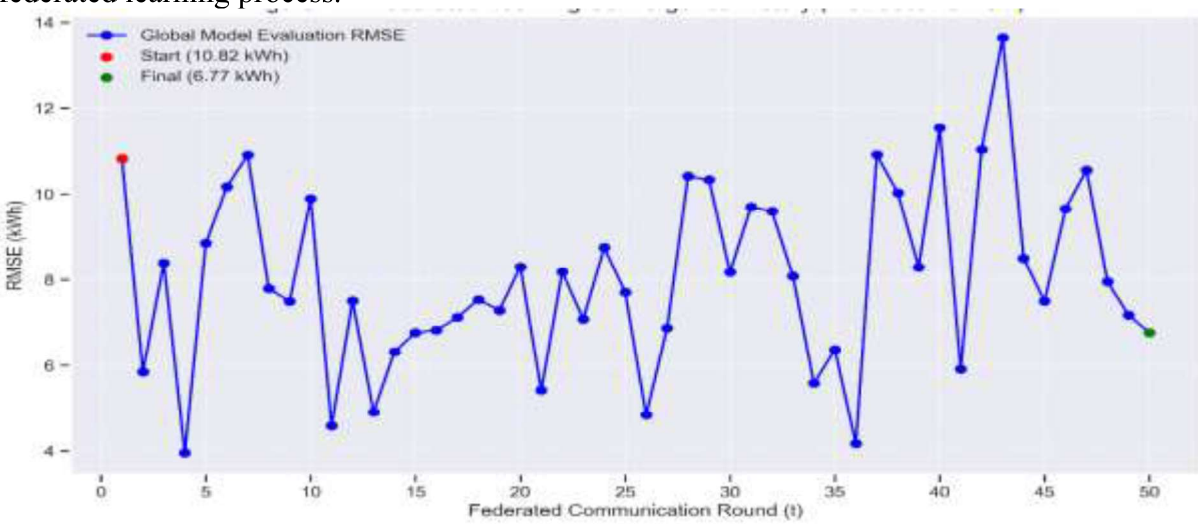

Figure 2: Federated Learning Training History - RMSE Across 50 Rounds

While there is still a performance difference between the federated implementation and the centralized XGBoost model, this difference reflects the trade-off necessary to provide decentralized and privacy-preserving learning. Crucially, the final FGB model (RMSE = 6.77 kWh) demonstrated its efficacy in a distributed system by still outperforming the centralized deep learning models, especially the GRU-only model. This outcome shows that excellent predictive performance can be maintained while addressing important data privacy and security concerns, and it confirms the strategic choice to implement XGBoost in a federated framework. Overall, the results verify that the study's goal of applying a precise and useful model within the limitations of federated learning was successfully accomplished.

3.3 Blockchain Integration

The Federated Gradient Boosting (FGB) framework was integrated with a simulated blockchain audit layer in the study's final phase to address security and accountability. This layer ensures transparency and tamper-proofing in the electricity theft detection system by creating an unchangeable record of all model updates, configurations, and performance metrics. Important information including timestamps, model identities, training settings, and evaluation outcomes are recorded by the blockchain audit. Interestingly, the model's cryptographic hashes at the start and finish of training function as distinct digital fingerprints, verifying that the model experienced real changes over the course of the 50 federated learning rounds and eliminating any chance of unnoticed changes.

Table 3: Blockchain Audit Summary for Federated Learning Process

Audit Parameter	Value	Description
Audit Timestamp	2025-10-22 22:50:47	Time of blockchain audit generation
Project ID	Federated Learning Energy Consumption Prediction	System identifier
Model Name	XGBoost	Base algorithm used in federated training
Initial Model Hash (SHA-256)	9b0985fb82e72e5ca2460...caa81a04	Cryptographic hash of model at training start
Final Model Hash (SHA-256)	43ef06a675c64997508831...c9c775afb430	Cryptographic hash of model after round 50
FL Rounds Executed	50	Total number of federated training rounds
Total Clients	50	Complete client pool size
Clients per Round Fraction	0.2 (10 clients)	Proportion of clients participating per round
Aggregation Method	Secure Tree Aggregation (Simulated)	Federated boosting aggregation protocol
Initial Evaluation RMSE	10.82 kWh	Performance at training start
Final Evaluation RMSE	6.77 kWh	Performance after final round
RMSE Reduction	37.48%	Overall improvement metric
Minimum RMSE Achieved	3.96 kWh	Best performance across all rounds
Maximum RMSE Observed	13.65 kWh	Worst performance across all rounds
RMSE Standard Deviation	2.11 kWh	Measure of training stability

By recording crucial training configurations, such as XGBoost hyperparameters and federated learning parameters like client participation rate and aggregation technique, the audit further improves reproducibility and accountability. Comprehensive performance measures are also recorded, such as the reduction in RMSE from 10.82 kWh to 6.77 kWh (a 37.48% improvement) and variability indicators like minimum, maximum, and standard deviation values. Stakeholders can independently confirm the training process and its results thanks to this comprehensive and permanent record, which promotes trust in the system's dependability and efficacy.

All things considered, the study's goal of guaranteeing system security and trust is immediately fulfilled by the integration of blockchain with federated learning. The platform offers openness through easily accessible training metadata, tamper-proof verification via cryptographic hashing, and an auditable history for regulatory supervision. By protecting data privacy while still guaranteeing integrity and trust, this hybrid architecture provides a more reliable solution than conventional centralized methods. As a result, it creates a safe and expandable basis for implementing AI-based electricity theft detection systems in settings where stakeholder responsibility and data sensitivity are crucial.

The majority of current research (Akinlabi et al. 2021; Ezeji et al. 2024) uses centralized machine learning or deep learning methods, such as CNN-XGBoost ensembles, TLENET with LoRAS, Deep Federated Learning with ConvGRU, SVM-PSO, and PSM-SVM models, to detect electricity theft. Although these methods show impressive performance metrics with accuracy ranging from 85.7% to 98.9%, they mostly rely on centralized data processing and do not include blockchain. The current study, on the other hand, uses a Federated XGBoost approach and achieves an RMSE of 6.77 kWh in the federated setting and 1.52 kWh in the centralized setting. It also uniquely integrates security through blockchain (Ganache) and privacy preservation through federated learning within the context of Gwagwalada, Nigeria.

The comparison reveals clear trade-offs between privacy and accuracy. Because they have access to aggregated datasets, centralized models like those in research Nwosu et al. (2022), Momoh & Onime (2020) and Okafor et al. (2021) achieve higher accuracy levels (92–98.9%); however, they establish single points of failure and pose significant privacy concerns. The federated technique used in this work, on the other hand, ensures user privacy by avoiding the sharing of raw data, although at the expense of worse predictive performance. The quantifiable cost of upholding a privacy-preserving framework is demonstrated by the federated model's RMSE of 6.77 kWh, which is a significant decrease from the centralized XGBoost model's 1.52 kWh. However, in contemporary energy systems where data security and regulatory compliance are crucial, this trade-off is becoming more and more warranted.

The study using federated ConvGRU is the most closely related of the reviewed works since it also uses privacy-preserving learning. Its inability to integrate blockchain, however, restricts its capacity to guarantee openness and confidence. In order to close this gap, the current study introduces a blockchain-based audit layer that uses cryptographic hashing to preserve immutable training histories, allowing for model integrity verification. Furthermore, although research (Ezeji et al. 2024) shows how AI-based theft detection can be used in a Nigerian setting (Enugu), it is dependent on centralized techniques. By fusing blockchain technology with federated learning, the current study expands this regional relevance and provides a more open, safe, and privacy-conscious solution for local electricity distribution issues.

4 Conclusion

In order to detect electricity theft in home smart metering systems, this study offers a novel architecture that combines blockchain technology and federated learning. The findings show that federated learning guarantees privacy-preserving model training, while XGBoost offers the best predictive performance. Federated deployment results in some speed degradation, but improved data security and privacy make the trade-off worthwhile. By guaranteeing transparency, immutability, and confidence in model updates, blockchain integration further fortifies the system. In semi-urban areas like Gwagwalada, where infrastructural constraints and privacy issues impede the uptake of centralized AI systems, the suggested method is especially appropriate. Future studies should concentrate on increasing the effectiveness of federated learning and investigating practical deployment techniques.

Acknowledgement

There is no funding support for this research.

References

Aitzhan, N.Z. and Svetinovic, D. (2018) 'Security and privacy in decentralized energy trading through multi signatures, blockchain and anonymous messaging streams', *IEEE Transactions on Dependable and Secure Computing*, 15(5), pp. 840–852.

Aderibigbe, A.A., Oladipo, F.O., Olufemi, O.I. and Adebayo, T.S. (2021) 'Smart meter data analytics for energy theft detection in the Nigerian power sector', *International Journal of Energy Economics and Policy*, 11(2), pp. 114–121.

Akinlabi, O.A., Ighravwe, D.E. and Adaramola, B. (2021) 'Electricity theft detection in a smart grid environment using a hybrid approach', *Energy Reports*, 7, pp. 5620–5628.

Akinlabi, O.A., Ighravwe, D.E. and Adaramola, B. (2022) 'Smart grid electricity theft detection using an ensemble learning model', *Sustainable Energy Technologies and Assessments*, 51, p. 101952.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. and Bengio, Y. (2014) 'Learning phrase representations using RNN encoder–decoder for statistical machine translation', *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734.

Chai, T. and Draxler, R.R. (2014) 'Root mean square error (RMSE) or mean absolute error (MAE) Arguments against avoiding RMSE normalization by standard deviation', *Geoscientific Model Development*, 7, pp. 1247–1250.

Ezeji, I., Okoro, O.I. and Aginam, C.H. (2024) 'Particle swarm-based SVM model for electricity theft detection: A practical implementation with Enugu Electricity Distribution Company data', *IEEE Transactions on Smart Grid*, 15(2), pp. 587–596.

Momoh, S. and Onime, C.E. (2020) 'Electricity theft detection in Nigeria using hybrid fuzzy logic–genetic algorithm model', *Journal of Engineering*, 14(2), pp. 44–55.

Nawaz, A., Ali, T., Mustafa, G. and Rehman, S.U. (2023) 'A novel technique for detecting electricity theft in secure smart grids using CNN and XGBoost', *Intelligent Systems with Applications*, 16, p. 200231.

Nwosu, C.E., Aneke, N.E. and Okafor, E.C. (2022) 'Deep learning-based electricity theft detection in smart grids: A review', *Sustainable Energy, Grids and Networks*, 29, p. 100560.

Omagbemi, O.O. and Adewole, K.S. (2020) 'Electricity theft detection in smart grids using support vector machines', *Journal of Electrical Engineering*, 71(4), pp. 266–274.

Omagbemi, O.O. and Adewole, K.S. (2021) 'Electricity theft detection in smart grids using a deep learning approach', *Heliyon*, 7(6), e07195.

Okafor, E.C., Aneke, N.E. and Okafor, C.J. (2020) 'Electricity theft detection in smart grids using machine learning and blockchain technology', *Heliyon*, 6(9), e04974.

Okafor, E.C., Aneke, N.E. and Okafor, C.J. (2021) 'A hybrid machine learning and blockchain-based electricity theft detection system for smart grids', *Sustainable Cities and Society*, 67, p. 102756.

Shahzadi, N., Javaid, N. and Akbar, M. (2024) 'A novel data-driven approach for combating energy theft in urbanized smart grids using artificial intelligence', *Expert Systems with Applications*, 242, p. 120074.

Sharma, A. and Tiwari, R. (2024) 'Anomaly detection in smart grid using optimized extreme gradient boosting with SCADA system', *Electric Power Systems Research*, 235, p. 110876.

Tofallis, C. (2012) 'A better measure of relative prediction accuracy for model comparisons', *Journal of the Royal Statistical Society Series A*, 175(4), pp. 379–396.

Willmott, C.J. and Matsuura, K. (2005) 'Advantages and limitations of the mean absolute error (MAE) and root mean square error (RMSE) in the evaluation of hydrological model performance', *Journal of Hydrology*, 30(1–4), pp. 79–82.

Zafar, M.H. et al. (2023) 'Step towards secure and reliable smart grids in Industry 5.0: A federated learning assisted hybrid deep learning model for electricity theft detection using smart meters', *Energy Reports*, 10, pp. 3001–3019.