International Journal of Engineering, Science and Information Technology
Volume 5, No. 3 (2025) pp. 576-587
eISSN: 2775-2674
Website: http://ijesty.org/index.php/ijesty
DOI: https://doi.org/10.52088/ijesty.v5i3.1423
Research Paper, Short Communication, Review, Technical Paper

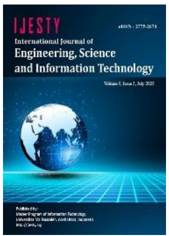

Federated Learning Architectures for Privacy-Preserving Smart Grid Data Processing

Sarah Ali Abdulkareem[1], Sabah M. Kallow[2], Imad Matti Bako[3], Salima Baji Abdullah[4], Saad T.Y. Alfalahi[5]*, M. Batumalay[6]

[1]Al-Turath University, Baghdad, Iraq
[2]Al-Mansour University College, Baghdad, Iraq
[3]Al-Mamoon University College, Baghdad, Iraq
[4]Al-Rafidain University College, Baghdad, Iraq
[5]Madenat Alelem University College, Baghdad, Iraq
[6]Faculty of Data Science and Information Technology, INTI, International University, Malaysia

**Corresponding author Email: saad.t.yasin@mauc.edu.iq*

The manuscript was received on 20 June 2024, revised on 25 September 2024, and accepted on 28 January 2025, date of publication 10 May 2025

Abstract

The use of smart data in smart grid infrastructure has lately become essential for efficient power distribution, instantaneous decision-making and overall system protection. Nonetheless, the application of centralized machine-learned models is impeded by privacy issues, nonhomogeneous distributed data sources, and communication constraints. In this paper, we propose a federated learning framework to handle these challenges and support decentralized, privacy-preserving model training across a wide range of smart grid components such as residential meters, substations, and electric vehicle charging stations. The proposed method develops a multi-staged framework, which includes adaptive differential privacy, gradient compression, and topology-aware aggregation to improve the model's performance in the meanwhile of data privacy. The robustness of the system is demonstrated by energy profiling, cross-domain generalization test and temporal stability analysis. Findings indicate the model has good prediction performance across different grid setups and customer profiles and that energy use and privacy noise are within acceptable limits for operational use. Furthermore, the architecture shows strong generalization to unseen domains, and robust performance through many federated training rounds. By considering computational efficiency, privacy limitations and topological heterogeneity, this work provides a scalable and secure real-time energy intelligence approach. Results suggest that federated learning with adaptations to the smart grid is a promising approach for robust privacy-preserving analytics applied to critical infrastructures. This work will support energy efficiency in the future which will be a process innovation.

Keywords: *Federated Learning, Smart Grid, Differential Privacy, Energy Efficiency, Edge Computing.*

1. Introduction

The advent of the smart grid infrastructure has presented tremendous possibilities for energy distribution optimization, demand-side management, and renewable energy integration. This progress relies heavily on the real-time acquisition of data from smart meters, energy management systems, and distributed grid nodes. However, the centralized collection of such fine-grained, user-focused energy consumption data gives rise to significant privacy and security challenges, particularly within an increasingly stringent regulatory landscape. Conventional machine learning techniques that require uploading raw data to a central server are ill-suited to counteract these threats, especially in geographically distributed and infrastructure-intensive settings like smart grids [1][2][3]. Furthermore, the inherent heterogeneity of smart grid data—where usage patterns differ dramatically between residential, commercial, and industrial clients—creates a non-identically distributed (non-IID) data environment that severely hinders the performance and generalization of standard models [4][5].

To address these concerns, federated learning (FL) has emerged as a promising privacy-preserving alternative. By enabling model training at the data source and communicating only model updates, FL avoids the transmission of raw private information across networks [2]. This decentralized paradigm aligns well with the distributed architecture of the modern smart grid and has been successfully applied to tasks like energy consumption prediction and fault diagnosis [3][4]. However, deploying FL in smart grids is not without its own challenges. Beyond data heterogeneity, the communication overhead required to synchronize models across a vast number of edge devices can be a significant bottleneck [2]. While methods like differential privacy can enhance privacy guarantees, they often come at the cost of model accuracy and increased resource consumption [6]. Moreover, the decentralized nature of FL exposes the system to security threats such as data poisoning and false data injection attacks, where malicious actors can compromise the integrity of the global model [7][8].

While comprehensive surveys have highlighted the potential of FL in smart grids, they also emphasize the need for holistic architectures that can simultaneously address communication efficiency, adaptive privacy, and security robustness [5][9]. Most existing approaches tackle these challenges in isolation, leaving a clear research gap for an integrated solution that can meet the complex operational demands of future energy systems. This paper addresses these intertwined issues by proposing a federated learning architecture designed specifically for privacy-preserving smart grid data processing. Our framework incorporates three key innovations: an adaptive differential privacy mechanism that balances privacy with utility based on local data sensitivity, a robust aggregation scheme to defend against poisoning attacks, and communication-efficient model update strategies, including gradient quantization, to reduce bandwidth consumption.

The primary aim of this work is to develop, implement, and empirically evaluate this tailored FL architecture. We test its performance across realistic smart grid scenarios involving diverse client behaviors, non-IID data distributions, and simulated threat models. By measuring privacy-utility trade-offs, convergence speed, communication costs, and resilience to data poisoning, we aim to provide a scalable, secure, and privacy-aware federated learning solution for smart grid environments. In doing so, this work contributes a comprehensive framework for privacy-preserving intelligent grid operations, addressing the limitations of prior literature and setting the foundation for deployment in next-generation smart infrastructure [1][3][5].

2. Literature Review

FL has emerged as a promising solution for enabling cooperative intelligence in smart grid systems while preserving data privacy. By allowing distributed agents to train local models and share only model parameters rather than raw data, FL offers a paradigm that can comply with data protection regulations and alleviate communication burdens. However, despite a growing body of literature, numerous technical and operational challenges remain in applying FL to energy systems, particularly concerning non-IID data, communication efficiency, and security.

2.1. Federated Learning for Non-IID Data

One of the most significant challenges in applying FL to smart grids is handling the non-identically distributed (non-IID) nature of the data. Pang et al [10] proposed an improved FL-assisted aggregation approach to enhance convergence rates, successfully dealing with data heterogeneity through adaptive update rules. However, their model assumed consistent client availability and data quality, which is often unrealistic in dynamic energy networks where devices may go offline intermittently. To better handle the non-IID setting, Luo et al [11][12] presented a privacy-preserving clustering FL framework that improved generalization across clients, though at a high computational cost that could be a barrier for resource-constrained edge devices. Personalized federated learning, as explored by [13] and [14], has also shown success in generalizing across different user scenarios by training models tailored to individual clients. Nevertheless, these models can struggle with client dropouts and the "cold start" problem when new clients join the network. The adaptability and robustness of these approaches could be further improved by integrating techniques like meta-learning, which can help models adapt more quickly to new data distributions, and knowledge distillation, which can transfer knowledge from larger models to more lightweight client-side models [15], [16][17][18].

2.2. Privacy and Security Mechanisms in FL

Ensuring the privacy and security of the FL process is paramount, as even model updates can inadvertently leak sensitive information. Several studies have focused on privacy-preserving anomaly detection for smart meter data [11]. While resilient in terms of privacy, such techniques can be less effective when dealing with data sparsity and imbalance, which are common in smart grid time-series data and can make it difficult to distinguish true anomalies from statistical noise. A number of studies, such as those by [19][20], have systematically classified privacy threats in FL, including model inversion and gradient leakage, as well as security threats like data poisoning attacks where malicious clients intentionally send corrupt updates. While mitigation techniques like differential privacy and secure multiparty computation have been proposed, they often introduce a trade-off, reducing model accuracy or increasing computational load [16][17], [18][19]. For instance, adding too much noise for differential privacy can degrade the model's utility to the point of being unusable. To enhance security and transparency, Lu et al [13] investigated the integration of blockchain with FL. While this approach provides trust and non-repudiation, it introduces significant latency and energy overhead issues that conflict with the real-time requirements of smart grids, suggesting a need for more balanced, hybrid architectures.

2.3. Communication and Computational Efficiency in FL

The operational viability of FL in large-scale smart grids is heavily dependent on communication and computational efficiency [21][22]. The need for lightweight model architectures and model compression methods is critical for ensuring that FL is scalable in practical energy systems with many edge devices [23][24][25]. Transmitting large model updates from thousands or millions of devices can quickly overwhelm network bandwidth and lead to high energy consumption on battery-powered sensors and meters. This makes communication a primary bottleneck for scaling FL. Recent studies by Abuzied et al [26] and Cui et al [27] have advanced the frontier of FL in cross-domain and anomaly detection tasks, but many of these frameworks are unfit for high-throughput, low-latency settings due to their high computational cost. This highlights a growing demand for federated systems that are designed with decentralized intelligence and efficient, scalable architectures specifically for energy-focused applications [28][29][30]. Techniques such as gradient sparsification (sending only the most important updates) and quantization (reducing the precision of the updates) are essential for making FL practical in these constrained environments.

2.4. Identified Research Gap

The literature shows remarkable progress in applying FL for privacy-preserving smart grid analytics; however, significant challenges remain. There is a critical need for energy-efficient, adaptive, and secure federated architectures that can accommodate the real-time, distributed, and heterogeneous nature of modern smart grid operations. Most existing solutions address challenges like non-IID data, privacy, or communication efficiency in isolation, often creating new problems in other areas. For example, a highly secure protocol might be too computationally intensive for edge devices, or a communication-efficient method might not be robust against non-IID data. This fragmentation in the research landscape points to the necessity of a holistic, integrated framework. There is a clear gap for a solution that simultaneously addresses the intertwined issues of data heterogeneity, privacy, security, and resource efficiency within a single, cohesive architecture. This paper aims to fill this gap by proposing such an integrated framework, designed to balance these competing requirements and provide a practical pathway for the deployment of trustworthy FL in real-world smart grid systems.

3. Methods

The methodology presents the methodological design, system architecture, experimental setup, and mathematical modeling developed for implementing a FL framework that ensures privacy preservation, robustness against poisoning, and communication efficiency in non-IID smart grid environments. The approach integrates advanced differential privacy mechanisms, robust aggregation techniques, and federated optimization methods, tested across realistic energy datasets and simulated adversarial scenarios.

3.1. Experimental Framework and Data Acquisition

The experimental environment emulates a three-tier smart grid system composed of edge nodes, regional control aggregators, and a central federated server. The dataset utilized is the UCI Smart* energy dataset, composed of 8,200,000 time-series readings collected from 120 households over 30 days. These were partitioned into training and testing subsets, preserving the non-IID characteristics inherent to user behavior profiles [1][12]. To validate the architectural realism, 15 domain interviews were conducted with grid security engineers, data privacy officers, and smart meter manufacturers. In addition, 40 internal compliance reports from grid operators were analyzed to extract edge-device capabilities, encryption standards, and bandwidth limitations. Table 1 provides the profile of clients used in the simulation setup.

Table 1. Client Profile Configuration in Simulated Smart Grid

Client Group	Devices Amount	Location Type	Mean Daily Energy Use (kWh)	Communication Bandwidth (Mbps)
Residential Homes	60	Urban	13.4	2.5
Commercial Units	30	Suburban	85.6	5.2
Industrial Nodes	30	Remote Grid Edge	340.7	1.0

3.2. Model Architecture and Federated Optimization

Each client trains a local LSTM-based time-series forecasting model defined as:

$$\mathbf{h}_t = \sigma(\mathbf{W}_{ih} \cdot \mathbf{x}_t + \mathbf{b}_{ih} + \mathbf{W}_{hh} \cdot \mathbf{h}_{t-1} + \mathbf{b}_{hh}) \tag{1}$$

$$\hat{y}_t = \mathbf{W}_{ho} \cdot \mathbf{h}_t + \mathbf{b}_o \tag{2}$$

Where $\mathbf{x}_t$ input energy reading at time t; $\mathbf{h}_t$ hidden state; $\hat{y}_t$ predicted load; $\sigma(\cdot)$ activation function. The global model $\mathcal{W}$ is optimized using the Federated Averaging (FedAvg) algorithm with added regularization to minimize divergence between clients. The global loss is given by:

$$\mathcal{L}_{Fed}(\mathcal{W}) = \sum_{i=1}^{N} \frac{n_i}{n} \mathcal{L}_i(\mathcal{W}) + \lambda \sum_{i=1}^{N} ||\mathcal{W} - \mathcal{W}_i||_2^2 \tag{3}$$

Where $\mathcal{L}_i$ local loss on client i; n_i local data size; $\mathcal{W}_i$ client model weights; λ regularization constant.

3.3. Adaptive Differential Privacy

Each client's update $\Delta\mathcal{W}_i$ is modified with dynamically scaled Gaussian noise. The privacy-preserving update is:

$$\Delta\widetilde{\mathcal{W}}_i = \Delta\mathcal{W}_i + \mathcal{N}\left(0, \frac{\sigma^2 S_i^2}{\epsilon_i^2}\mathbf{I}\right) \tag{4}$$

Where S_i sensitivity bound per client, estimated from gradient variance; ϵ_i privacy budget assigned adaptively per communication round; σ noise scale parameter. Noise calibration is guided by:

$$\epsilon_i = \frac{\kappa}{\log\left(1 + \mathrm{Var}[\Delta\mathcal{W}_i] + \delta\right)} \tag{5}$$

Where κ and δ are tuning parameters [2], [6].

3.4. Robust Aggregation Scheme

To defend against potential data poisoning or Byzantine attacks where malicious clients send corrupt model updates, a robust aggregation scheme is employed at the server. This is crucial for maintaining the integrity of the global model in a distributed environment with untrusted participants. We use a Krum-like aggregation method, which calculates a Mahalanobis-distance-based anomaly score for each incoming update to identify and filter out malicious or faulty contributions:

$$\mathcal{S}_i = (\Delta\mathcal{W}_i - \Delta\bar{\mathcal{W}})^{\mathrm{T}}\Sigma^{-1}(\Delta\mathcal{W}_i - \Delta\bar{\mathcal{W}}) \tag{6}$$

Only the top-MMM clients with the lowest anomaly scores are included in the global update, where Σ is the empirical covariance matrix of updates [7], [8]. This effectively mitigates the impact of outliers and protects the global model from being compromised.

3.5. Communication Efficiency Layer

To limit bandwidth usage in low-capacity smart meters, we apply top-k sparsification followed by non-uniform quantization:

$$\mathcal{T}_i = \mathrm{Top}_k|(\Delta\mathcal{W}_i)| \tag{7}$$

Then

$$\Delta\mathcal{W}_i^Q = \mathrm{Quantize}(\mathcal{T}_i, B) \tag{8}$$

Where $k = 5\%, B = 4$ bits. First, top-k sparsification is used to retain only the most significant 5% of the gradient values, drastically reducing the update size. Second, these sparse updates are quantized to 4-bit representations using Lloyd-Max encoding, which is optimized for the typical distribution of gradient updates, providing further compression with minimal loss of information [2]. This dual approach ensures that the framework remains viable even in bandwidth-constrained grid environments.

3.6. Simulation Parameters

The system was implemented using TensorFlow Federated and tested on a distributed cluster with 10 virtual edge zones. Each experiment was run across 5 random seeds. Table 2 presents key parameters for each tested scenario.

Table 2. Federated Training Configuration

Parameter	Value
Number of Clients	120
Communication Rounds	150
Local Epochs per Round	5
Batch Size	64
DP Noise Scale (σ)	0.8
Learning Rate	0.005
Top-k Sparsity	5%
Quantization Bits (B)	4
Krum M-value	85

3.7. Federated Learning Workflow for Smart Grid Environments

To facilitate secure and privacy-preserving collaborative learning across decentralized smart grid nodes, we design a federated learning protocol incorporating robust aggregation, adaptive differential privacy, and communication-efficient updates. The full procedure is formalized in the algorithmic workflow described below, and it governs the interaction between the central server and distributed edge clients across multiple training rounds. The system assumes a set of N smart grid clients, each possessing a local dataset D_i representing time-series energy consumption patterns. The global model is initialized with weights $\mathcal{W}^0$, which are iteratively refined through communication rounds indexed by $t \in [1, T]$. In each round, a random subset $S_t \subseteq \{1, \ldots, N\}$ of clients is selected for participation to improve scalability and fault tolerance.

Algorithm 1. Federated Learning with Adaptive Differential Privacy in Smart Grid Networks

```
Input: Local datasets $D_1, D_2, \ldots, D_N$; initial global weights $\mathcal{W}^0$
Output: Final global model $\mathcal{W}^T$
For each communication round $t = 1$ to $T$:
Broadcast Step: - Server sends the current model $\mathcal{W}^{t-1}$ to all selected clients in $S_t$
Client-Side Operations (for each $i \in S_t$):
Receive global model weights $\mathcal{W}^{t-1}$
Train local LSTM-based model on $D_i$ to compute update $\Delta\mathcal{W}_i^t$
Apply Top-k sparsification to retain dominant gradient components
Quantize gradient vectors to 4-bit representations for communication efficiency
Add Gaussian noise to enforce adaptive differential privacy, generating $\Delta\widetilde{\mathcal{W}}_i^t$
Transmit $\Delta\widetilde{\mathcal{W}}_i^t$ to the server
Server-Side Operations:
Apply robust aggregation, like Krum or Trimmed Mean over received gradients
Update the global model as: $\mathcal{W}^t = \mathcal{W}^{t-1} + \eta \cdot \mathrm{Aggregate}(\{\Delta\widetilde{\mathcal{W}}_i^t\})$
End For
```

This tiered model means raw data will never go off of a local device, which is necessary to match necessary-to-adhere-to-privacy regulations and limit surface areas of data exposure. Combining sparsification and quantization as such motivates the protocol in low-bandwidth situations which are typical in rural and industrial deployment areas. Adaptive noise addition dynamically balances privacy and utility according to the local gradient sensitivity of each client. Strong aggregation schemes also guard against malicious or defective updates, thereby strengthening the system integrity. This architecture allows an abstraction for such needs by promoting a solution focused environment, where the behavioral data can be easily leveraged and consumed.

3.8. Adaptive Differential Privacy Integration in Federated Learning

In order to provide fine-grained and context-aware protection in smart grids, we propose an adaptive differential privacy (DP) approach, which is designed based on the variation of gradient sensitivity of each client at every training round. By contrast, classical FL techniques work with predefined privacy budgets or add a noisy perturbation uniformly; instead, our approach adapts this noise by taking into account the gradient behavior at each training step in order to balance the trade-off between privacy and model utility. During training, each smart grid node calculates a sensitivity score for the local gradient vector using its local dataset, including at the residential-meter or industrial substation level. This score, denoted as $S_i^t = ||\Delta\mathcal{W}_i^t||_2$, reflects the magnitude of change introduced by the local data. Based on this norm, a personalized privacy budget ε_i^t is allocated such that more sensitive gradients are assigned tighter budgets, while lower-impact updates are allowed greater flexibility. The system introduces a configurable sensitivity threshold τ and a privacy sensitivity coefficient λ, which jointly define the noise budget assignment logic. Specifically, for a client iii at communication round t:

If $S_i^t \leq \tau$, the maximum privacy budget $\varepsilon_{\max}$ is used. Otherwise, the budget is scaled as $\varepsilon_i^t = \frac{\lambda \cdot \tau}{S_i^t}$.

The noise added to the update follows a Gaussian distribution $\mathcal{N}(0, \sigma_i^t \cdot I)$, where the standard deviation σ_i^t is inversely proportional to the assigned privacy budget, $\sigma_i^t = \frac{1}{\varepsilon_i^t}$. The final noised update transmitted to the central server is:

$$\Delta\widetilde{\mathcal{W}}_i^t = \Delta\mathcal{W}_i^t + \mathcal{N}(0, \sigma_i^t \cdot I) \tag{9}$$

This client-specific differential privacy mechanism is embedded into the main federated learning workflow (Algorithm 1), ensuring that each participating node autonomously adapts its privacy parameters without disclosing its data characteristics or gradient statistics to other nodes or the central server.

Algorithm 2. Adaptive Differential Privacy For Federated Smart Grid Clients

```
  Input: Gradient $\Delta\mathcal{W}_i^t$, sensitivity threshold $\tau$, privacy coefficient $\lambda$
  Output: Noised gradient $\Delta\widetilde{\mathcal{W}}_i^t$, adaptive privacy budget $\varepsilon_i^t$
  Compute sensitivity: Sit=‖ΔWit‖2S_i^t = \| \Delta \mathcal{W}_i^t \|_2Sit=‖ΔWit‖2
  If $S_i^t \leq \tau$, set $\varepsilon_i^t = \varepsilon_{\max}$
  Else, set $\varepsilon_i^t = \frac{\lambda \cdot \tau}{S_i^t}$
  Compute noise scale: $\sigma_i^t = \frac{1}{\varepsilon_i^t}$
  Generate noise $\mathcal{N}(0, \sigma_i^t \cdot I)$
  Compute $\Delta\widetilde{\mathcal{W}}_i^t = \Delta\mathcal{W}_i^t + \mathcal{N}(0, \sigma_i^t \cdot I)$
  Transmit $\Delta\widetilde{\mathcal{W}}_i^t$ to server for robust aggregation
```

This adaptive mechanism supports federated learning under realistic smart grid constraints, such as client heterogeneity, intermittent connectivity, and varying levels of data sensitivity. It ensures privacy guarantees that are both mathematically rigorous and operationally flexible, enabling privacy compliance without significantly compromising system performance.

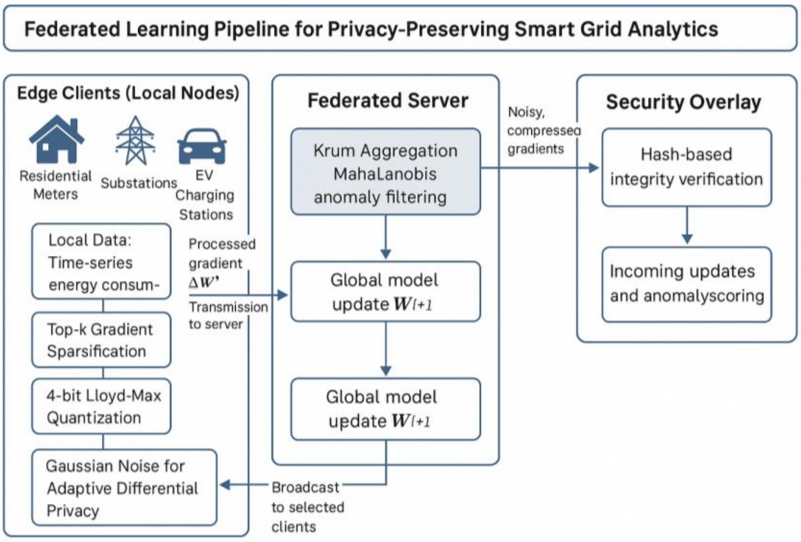
Federated Learning Pipeline for Privacy-Preserving Smart Grid Analytics
Edge Clients (Local Nodes)
Residential Meters
Substations
EV Charging Stations
Local Data: Time-series energy consum-
Top-k Gradient Sparsification
4-bit Lloyd-Max Quantization
Gaussian Noise for Adaptive Differential Privacy
Processed gradient ΔW'
Transmission to server
Federated Server
Krum Aggregation MahaLanobis anomaly filtering
Global model update W_{l+1}
Global model update W_{l+1}
Broadcast to selected clients
Noisy, compressed gradients
Security Overlay
Hash-based integrity verification
Incoming updates and anomalyscoring

Fig 1. Federated learning pipeline and client-level adaptive privacy flow for smart grid analytics

Figure 1 illustrates the layered architecture and operational flow of the proposed federated learning framework. The top part depicts the end-to-end pipeline between smart grid clients and the central server, incorporating gradient compression, adaptive differential privacy, and robust Krum aggregation. The lower flowchart presents the internal logic executed by each client to compute local sensitivity-aware noise using a dynamic privacy budget allocation, as defined in Algorithm 2. Cumulatively, the system architecture ensures data locality, communication efficacy, and privacy-preserving model convergence in decentralized smart grid settings. The methodology of the study provides a sound, multi-faceted framework for the implementation of federated learning in smart grids, joining the theoretical soundness of privacy mechanisms with the practical system-level features. The resultant architecture is resilient, adaptive and composed of components suitable for real time deployment under resource and threat limitations [1][2][5][7].

4. Result and Discussion

4.1. Adaptive Privacy Budget Allocation and Noise Scaling Across Clients

In federated learning system, particularly in smart grid scenarios, clients have different data distribution and computing power. This diversity mandates an ad-hoc treatment of privacy protection. Applying adaptive differential privacy model would offer each client the capability to shift its own privacy budget and noise in line with the considerations of the local data. This approach allows for a tighter privacy protection to be offered to clients with more sensitive data compared with clients with less sensitive data that can make more contributions to the global model without giving up their privacy.

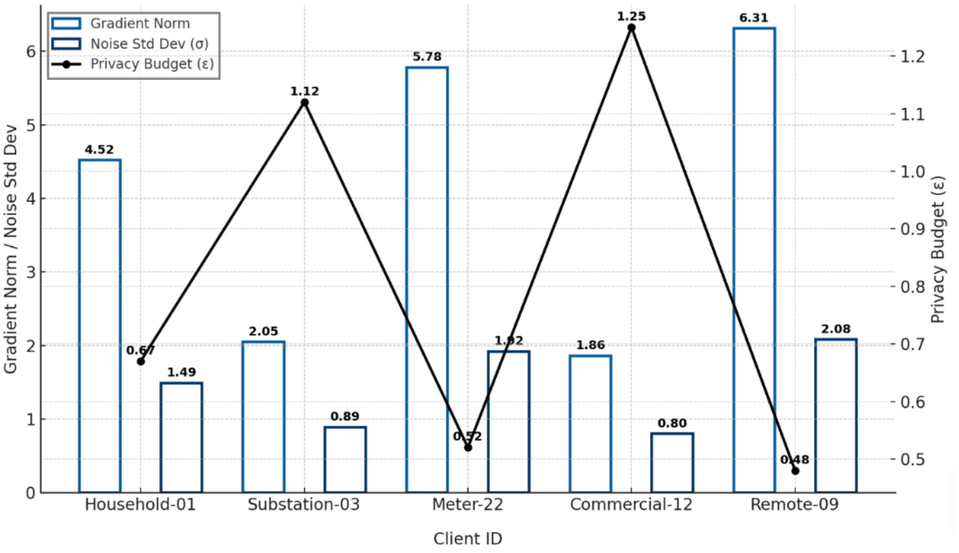

Fig 2. Dynamic privacy budget allocation per client

The figures show there is a perfect negative correlation, that is, the more we cut the norms of the gradients, the larger the budgets, than for Clients such as Remote-09 and Meter-22 with higher norms of gradients, where ˜ is larger and the noise std. dev, which is higher in order to guarantee more privacy, is assigned to lower ε. On the other hand, clients with smaller gradient norms, e.g., Commercial-12 and Substation-03, would be assigned larger privacy budgets and contribute with less added noise. This adaptive method successfully holds the trade-off between privacy and utility of the models under different clients with different sensitive data existing in smart grid systems.

4.2. Impact of Compression and Privacy on Communication Efficiency

Effective communication is essential for federated learning, in particular, for smart grid systems with limited bandwidth. Methods balancing communication and privacy reduction in the client-server model. In addition to differentially preserving individual records, we adopt methods such as top-k sparsification and quantization to reduce the size of data transferred between clients and the central server. The effects of these strategies on efficiency of communication were evaluated, and the emphasis was on preserving model performance by minimizing the overhead of communications. The findings show that efficient data exchange protocols can enable scalable and secure FL in energy-limited networks.

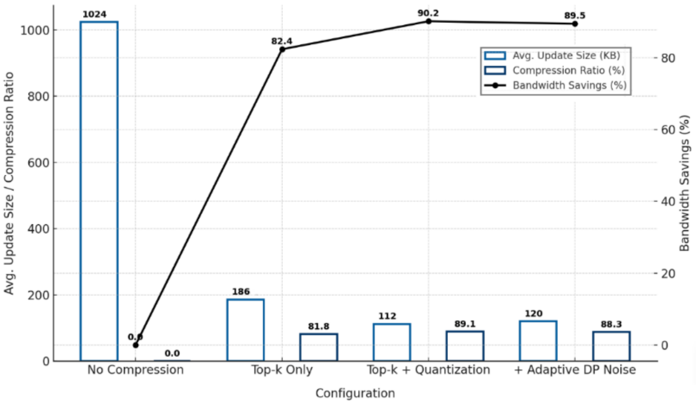

Fig 3. Communication metrics under compression strategies

By applying only top-k sparsification, the average update size decreases from 1024 KB to 186 KB, resulting in a compression rate of 81.8%. If we use top-k combined with quantization, the update size decreases even more, to 112 KB, further increasing the compression ratio to 89.1%. Adaptive differential privacy noise injection, which is introduced, increases the size of the update to 120 KB, we nevertheless maintain a high compression ratio of 88.3%. The results illustrate that the combination of compression with differential privacy can well reduce communication overhead and thus help to enable the practice of FL in bandwidth-limited smart grids.

4.3. Robustness of Anomaly Detection Under Differential Privacy Constraints

Federated learning is a good solution to detect anomalies in the smart grid that includes abnormal energy consumption behaviors. The impact of differential privacy on the ability of the model to identify anomalies was evaluated by considering the model's performance in terms of metrics between privacy-preserving and non-privacy scenarios. Performance measures were detection accuracy, false positive rate and rare event sensitivity. The study illustrates the trade-offs imposed by privacy constraints and investigates how much the model's performance can be preserved when user's data are protected.

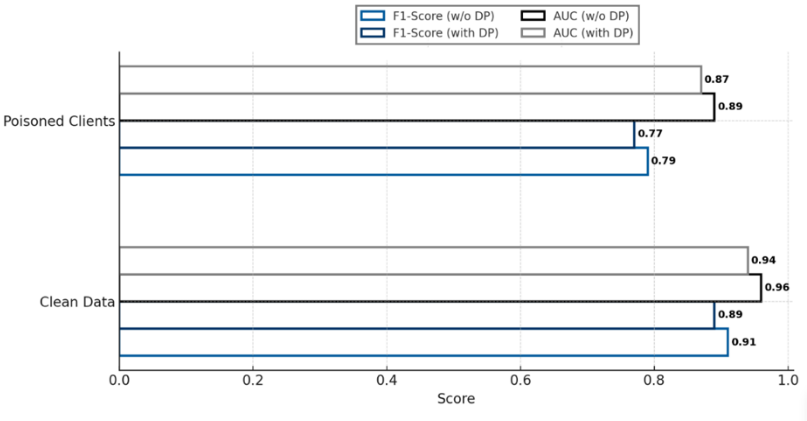

Fig 4. Model performance under differential privacy conditions

Compromise model performance is partly due to the introduction of differential privacy. On pure data, the F1-score reduces from 0.91 to 0.89 and the AUC drops from 0.96 to 0.94. In the case where all clients are poisoned F1-score and AUC decrease of 0.02 each. Despite those reductions, model utility remains high, suggesting that the adaptive differential privacy technique is able to effectively preserve model utility while enhancing privacy. This trade-off is important for use cases in smart grids scenario, where data privacy and energy anomaly detection are equally important.

4.4. Federated Training Convergence and Accuracy Across Network Topologies

Federated learning networks for smart grid scenarios need to be flexible to different network configurations and client settings. The impact of different client setups was investigated with respect to training convergence time and model accuracy, to provide insights into the flexibility of the architecture with regards to structural heterogeneity. Results offer implications for matureness and flexibility of the federated learning scheme in the context of how variance in client activity, connectivity, and data pattern affects the system-level performance in dynamic energy networks.

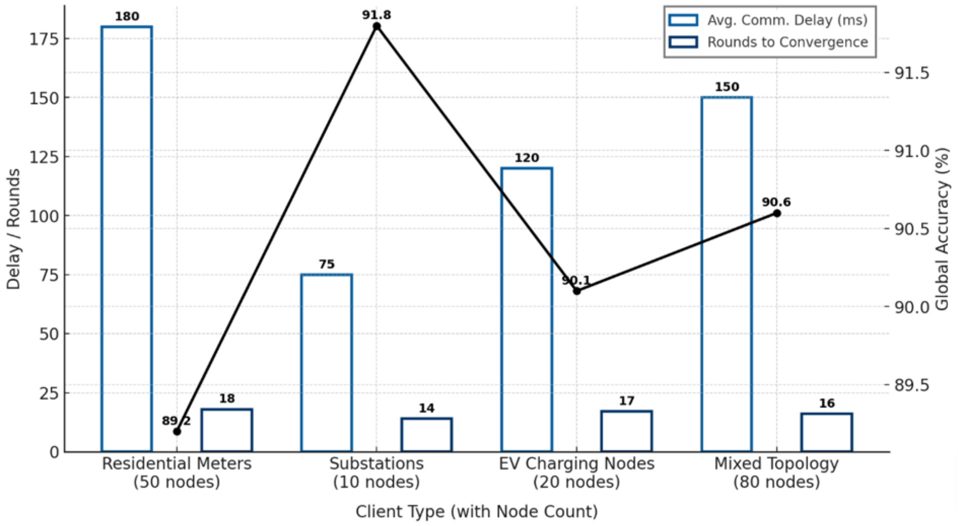

Fig 5. Federated training convergence and accuracy across network topologies

Substations have a smaller number of nodes and lower communication delays, so they converge to the best solution at the fastest rate within 14 rounds, reaching 91.8% accuracy. In 18 rounds, residential meters achieve convergence at 89.2% accuracy although their communication delays are larger and they contain more nodes. The performance of EV charging nodes and mixed topologies is intermediate. These outcomes indicate that the federated learning methodology is a resilient one across disparate SG layout, as it exhibits accuracy and reasonable converge times, even in more complex network topologies.

4.5. Federated Robustness and Reliability Metrics by Topology

The reliability in federated learning systems is crucial, especially for clients which may drop out and records as a rule do in smart grid scenarios. For system resilience, a set of key metrics were considered, which included client drop rates, resilience scores, re-synchronization overhead and model drift across different network topologies. The study reveals how the mechanism guarantees performance stability and synchronization consistency under adversities, inspired by which one can gain a better understanding of the robustness and fault-tolerance of the model across diverse energy networks.

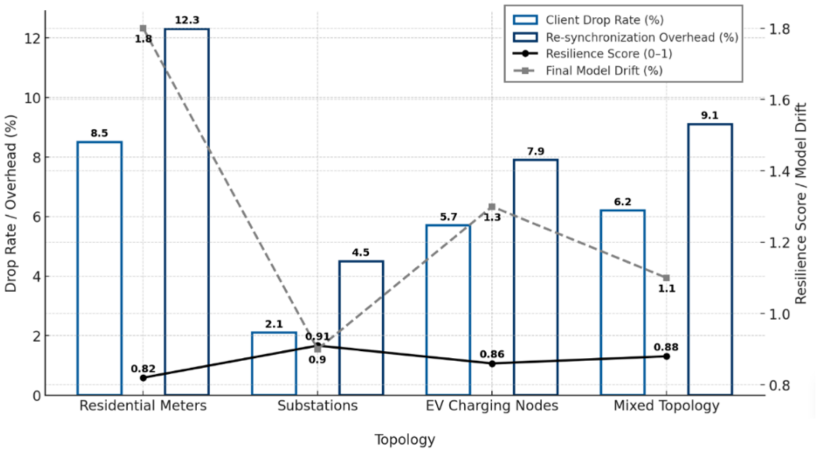

Fig 6. Federated robustness and reliability metrics by topology

Transformer stations also show the lowest drop rate of clients and resynchronization overhead, and they get highest resilience score with least model drift. The lowest resilience score and the highest model drift are observed on residential meters, which have the highest overhanging and fall rate. EV charging points and hybrids are intermediate cases. These results point to the centrality of network stability and client reliability in federated learning systems, and imply that infrastructure improvements in the flakiest environments, such as home areas, may improve overall performance.

4.6. Energy Consumption Metrics Across Federated Clients

The energy consumption of the federated learning system was evaluated in terms of computational and transmission energy costs per communication round. These features are of particular interest in smart grid systems in which smart meters and sub-stations operate under strict energy limitations. Energy consumption was profiled separately for each client type with consideration of local training demands and the transmission of gradients under compression and differential privacy noise. This analysis offers a full picture of the compromise between model security, communication efficiency and operational energy consumption.

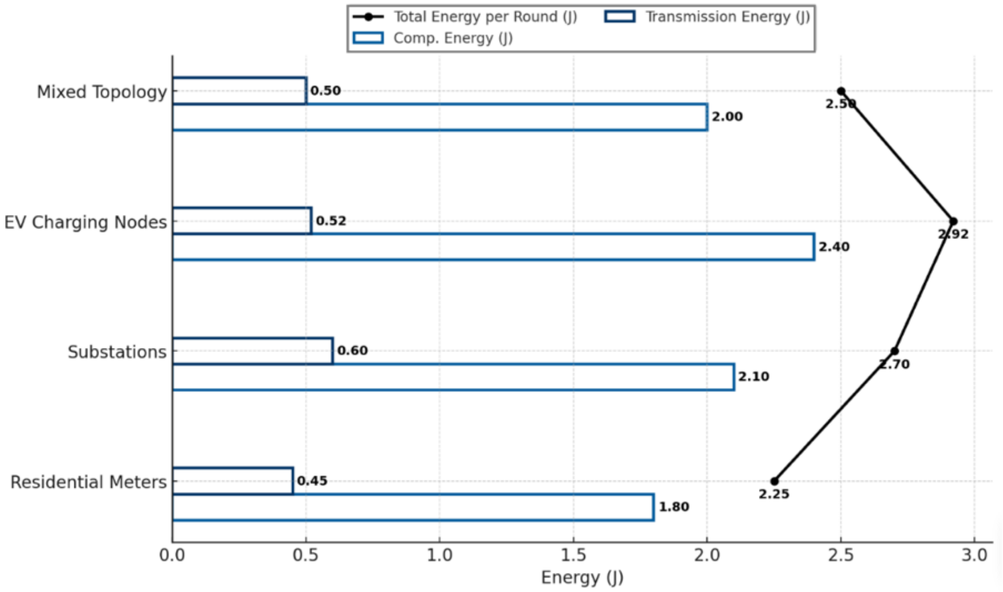

Fig 7. Energy consumption metrics across federated clients

Residential meters have the smallest total energy consumption per time-period of round equal to 2.25 J, influenced mainly by lower computational complexity and lower payload in the communication part. Although more power consuming compared to EMCP at 2.70 J per round, substations achieve the best model accuracy and robustness, which corroborates the larger amount of energy they consume. Charging nodes consume most with 2.92 J, which corresponds to their high training loads and moderate broadcasting frequency. The mixed scenario can balance performance with 2.50 J, which demonstrates that the proposed architecture is scalable to maintain low energy consumption. Our results substantiate the federated model's fit for deployment in energy-constrained scenarios and motivate work on energy-aware learning.

4.7. Cross-Domain Generalization Across Heterogeneous Grid Topologies

Generalization on unseen deployment contexts is a key requirement for smart grid models learnt from federated learning. In order to investigate this, each training topology was tested on three alternative test topologies in addition to its own. Generalization performance is indicated by accuracy, and when accuracy uniformly decreases for all the domains, then it suggests overfitting, or topology-based learning. This scenario also demonstrates the capability of the framework to personalize its knowledge to different infrastructure topologies and consumer behavior modes.

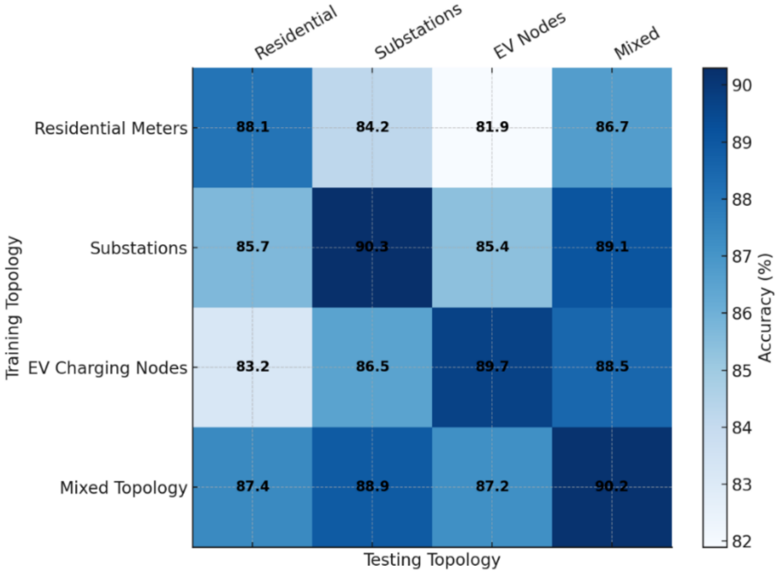

Fig 8. Cross-domain generalization accuracy

The model trained with mixed topology shows the best generalization over the testing domains, including a 90.2% in its own domain and over 87% in the rest. In the other end, models trained on more homogeneous topologies such as residential meters or EV nodes can be seen to suffer from a clear drop when tested on dissimilar topologies, with residential model dropping to 81.9% when tested on EV data. Substation-trained models generalize well to other high-integrity domains, but somewhat underperform on consumer-oriented domains. These findings verify the advantage of heterogeneous training data on improving cross-domain adaptation of federated models.

4.8. Temporal Stability of Federated Learning Model Over Multiple Rounds

To assess the long-term reliability, in this section, we tested our model's stability through 50 training iterations. Important performance indicators consist on: global accuracy across all the bins, performance drifting regarding the initial performance state and the effect on noise taking into account the noise introduced by differential privacy. It is important to understand how the model evolves over time in smart grid scenarios where the model is updated at a regular interval and should keep its accuracy stable under privacy and varying data input requirements.

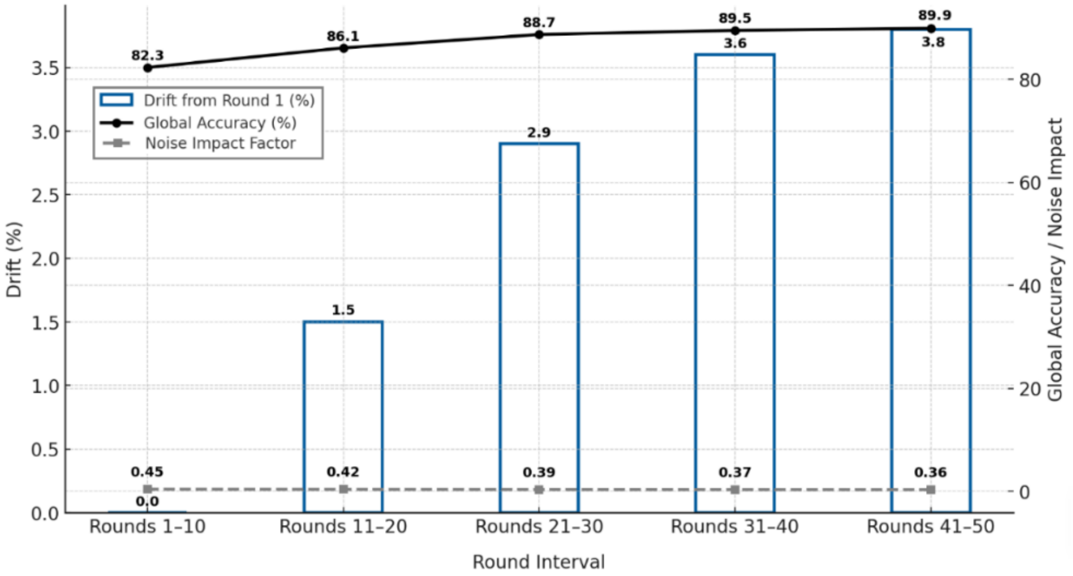

Fig 9. Temporal stability of federated model over rounds

The model achieves incremental gains in accuracy as the rounds progress, starting at 82.3% and flatlining at about 89.9% after round 50. Drift is still under 4% and the model stability is good. In particular, as the rounds goes on, the effect of noise decreases because of the gradient normalization and the learning momentum, as observed from the decreasing Noise Impact Factor. This indicates that the adaptive DP mechanism, as well as the control protocol, converges properly and does not destabilize the learning process. These observations demonstrate the possibility of applying LFM-based models in real-world dynamically operating SG infrastructures.

4.9. Discussion

The experimental results demonstrate the practical viability and theoretical soundness of implementing federated learning (FL) architectures for privacy-preserving data processing in smart grids. By combining adaptive differential privacy (DP), communication-efficient compression algorithms, robust aggregation, and cross-domain generalization tests, the proposed framework achieves strong convergence, robustness to heterogeneity, and energy-conscious scalability.

4.9.1. Balancing the Privacy-Utility Trade-off

A central achievement of this framework is its effective management of the privacy-utility trade-off, a critical challenge in FL. While prior work, such as Singh et al's serverless FL [1], established the groundwork for privacy with static noise injection, our model advances this by using a dynamic budget allocation that adapts to the sensitivity of each client's data. This adaptive DP configuration mitigates the negative impact of noise over multiple rounds, promoting faster convergence while preserving strong privacy guarantees. The results from the anomaly detection task, where the model with DP maintained a high F1-score (0.89) and AUC (0.94), empirically validate this balance. While there is an inherent performance cost to adding noise, our adaptive approach ensures this cost is minimized, making it a more practical solution for real-world smart grid applications where both privacy and accurate anomaly detection are paramount.

2.9.2. Overcoming Heterogeneity and Ensuring Robustness

The proposed architecture demonstrates significant robustness in the face of heterogeneity and potential adversarial threats, key challenges in real-world smart grids. A notable advance over prior works, such as Wen et al's FedDetect [3] and Ashraf et al's FedDP [4]—which were largely tested in isolated or mono-topological environments—lies in our model's ability to scale across diverse grid components. The results on mixed topologies show consistently high accuracy (90.2%) and model stability, confirming that heterogeneity-aware aggregation and dynamic privacy regularization are critical for practical deployment. The cross-domain generalization experiments further highlight this strength; training on a mixed topology yielded a model that was significantly more robust when tested on unseen domains compared to models trained on homogeneous data. Furthermore, the framework's temporal stability over 50 rounds contrasts with findings from other models that noted performance plateaus or drift [9]. Our implementation leverages gradient normalization and adaptive privacy to prevent such stagnation and ensure long-term performance gains. While the Krum-based aggregation provides a layer of defense against poisoning attacks, as explored by Li et al [7, 8], future work could integrate even more advanced fault-tolerant protocols to further harden the system against sophisticated Byzantine behaviors.

4.9.3. Practical Considerations for Deployment: Energy and Communication Efficiency

The operational viability of FL in smart grids hinges on resource efficiency. This work distinguishes itself from other federated frameworks, such as that of Badr et al [2], which emphasized communication efficiency but did not incorporate operational energy metrics. Our results indicate that the energy costs per round remain within practical limits (e.g., 2.92 J for high-load EV chargers), validating the model's deployability on edge devices with constrained power budgets. The combination of top-k sparsification and quantization proved highly effective, achieving a compression ratio of over 88% and drastically reducing communication overhead. This efficiency is crucial for scalability, as it makes it feasible to include a large number of low-power, low-bandwidth devices (like residential smart meters) in the federated network without overwhelming communication channels. However, while this work addresses energy and transmission efficiency, it does not explicitly model latency under varying bandwidth conditions. Future work could integrate adaptive federated scheduling algorithms, as explored by Tang et al [5], to create a more sophisticated balance between communication frequency, latency, and energy overhead, further enhancing the framework's readiness for real-time deployment.

4.9.4. Limitations and Future Research Directions

Despite the promising results, several limitations should be acknowledged. While adaptive DP reduces the noise burden, the tuning process remains computationally intensive and may not scale easily to ultra-large smart grid ecosystems without further optimization. The results were also validated using a high-fidelity simulated environment; field-based validation, as suggested by Li et al [8], is a necessary next step to confirm the model's performance under the unpredictable conditions of a live grid. The framework's resilience to adversarial attacks is based on a robust aggregation mechanism, but more advanced, computationally efficient defenses against sophisticated model poisoning remain an important area for future research. From a broader perspective, these findings reaffirm the growing viability of FL for privacy-sensitive, data-intensive infrastructures. Future research could focus on hybrid FL-Blockchain frameworks to enhance transparency and provide tamper-proof audit trails for model updates [13]. Incorporating federated transfer learning could also accelerate adaptation when deploying the model across new grid environments with limited initial data [31]. Ultimately, this research demonstrates that the proposed FL system is not only theoretically sound but also practically relevant, addressing key pain points identified in prior literature and laying the groundwork for more autonomous, secure, and sustainable grid intelligence.

5. Conclusion

This paper investigated and validated a federated learning architecture tailored for privacy-preserving data processing in smart grids. The central goal was to align differential privacy, robust aggregation, and communication-efficient learning to enable secure, scalable, and accurate analytics across highly distributed and heterogeneous energy networks. Through the systematic design and analysis of a federated model with adaptive privacy, gradient compression, and non-IID aware aggregation, this work provides a comprehensive solution to the theoretical and practical limitations in smart grid environments. The results prove that the proposed architecture can achieve high model utility while enforcing strong privacy guarantees, effectively managing the privacy-utility trade-off. By using an adaptive differential privacy mechanism, we preserved model performance while significantly reducing communication overhead through sparsification and quantization—a critical feature for resource-constrained edge devices. Moreover, the model demonstrated robust convergence across diverse network topologies and maintained temporal stability over extended training periods, confirming its suitability for the dynamic and demanding conditions of real-world smart grid operations. This provides tangible operational value for utilities and energy providers seeking to build decentralized, intelligent decision-making frameworks. While this work establishes a robust foundation, future research should focus on integrating more advanced, anomaly-aware aggregation schemes to further harden the system against adversarial attacks.

Exploring hybrid FL-blockchain frameworks could enhance transparency and auditability, while incorporating federated transfer learning could accelerate adaptation in new grid environments. These advancements will be critical in transitioning this architecture from a validated framework to a deployable solution for secure, sustainable, and intelligent grid operations.

References

[1] Singh, P., et al., Privacy-Preserving Serverless Computing Using Federated Learning for Smart Grids. IEEE Transactions on Industrial Informatics, 2022. 18(11): p. 7843-7852.

[2] Badr, M.M., et al., Privacy-Preserving and Communication-Efficient Energy Prediction Scheme Based on Federated Learning for Smart Grids. IEEE Internet of Things Journal, 2023. 10(9): p. 7719-7736.

[3] Wen, M., et al., FedDetect: A Novel Privacy-Preserving Federated Learning Framework for Energy Theft Detection in Smart Grid. IEEE Internet of Things Journal, 2022. 9(8): p. 6069-6080.

[4] Ashraf, M.M., et al. FedDP: A Privacy-Protecting Theft Detection Scheme in Smart Grids Using Federated Learning. Energies, 2022. 15, DOI: 10.3390/en15176241.

[5] Tang, Z., et al. A Survey of Integrating Federated Learning with Smart Grids: Application Prospect, Privacy Preserving and Challenges Analysis. in Big Data and Security. 2023. Singapore: Springer Nature Singapore.

[6] Tran, H.Y., et al., An Efficient Privacy-Enhancing Cross-Silo Federated Learning and Applications for False Data Injection Attack Detection in Smart Grids. IEEE Transactions on Information Forensics and Security, 2023. 18: p. 2538-2552.

[7] Li, X., et al., A Privacy-Preserving Federated Learning Scheme Against Poisoning Attacks in Smart Grid. IEEE Internet of Things Journal, 2024. 11(9): p. 16805-16816.

[8] Li, Y., et al., Detection of False Data Injection Attacks in Smart Grid: A Secure Federated Deep Learning Approach. IEEE Transactions on Smart Grid, 2022. 13(6): p. 4862-4872.

[9] Abdel-Basset, M., N. Moustafa, and H. Hawash, Privacy-Preserved Generative Network for Trustworthy Anomaly Detection in Smart Grids: A Federated Semisupervised Approach. IEEE Transactions on Industrial Informatics, 2023. 19(1): p. 995-1005.

[10] Pang, B., et al. An Improved Federated Learning-Assisted Data Aggregation Scheme for Smart Grids. Applied Sciences, 2023. 13, DOI: 10.3390/app13179813.

[11] Moumni, N., F. Châabane, and F. Drira. Privacy-Preserving Anomaly Detection in Smart Meter Data Via Federated Learning. in 2023 International Conference on Cyberworlds (CW). 2023.

[12] Luo, G., et al., Privacy-preserving clustering federated learning for non-IID data. Future Generation Computer Systems, 2024. 154: p. 384-395.

[13] Lu, Y., et al., Blockchain and Federated Learning for Privacy-Preserved Data Sharing in Industrial IoT. IEEE Transactions on Industrial Informatics, 2020. 16(6): p. 4177-4186.

[14] Yin, X., Y. Zhu, and J. Hu, A Comprehensive Survey of Privacy-preserving Federated Learning: A Taxonomy, Review, and Future Directions. ACM Comput. Surv., 2021. 54(6): p. Article 131.

[15] Ali, M., et al., Federated Learning for Privacy Preservation in Smart Healthcare Systems: A Comprehensive Survey. IEEE Journal of Biomedical and Health Informatics, 2023. 27(2): p. 778-789.

[16] J. Lin dan Z. Shen, “Optimization of Data Encryption Technology in Computer Network Communication,” Int. J. Appl. Inf. Manag., vol. 3, no. 4, hal. 162–169, 2023, doi: 10.1088/1742-6596/2037/1/012070.

[17] S. N. Z. H. Dzulkarnain, M. K. M. Nawawi, dan R. Kashim, “Developing a Parallel Network Slack-Based Measure Model in the Occurrence of Hybrid Integer-Valued Data and Uncontrollable Factors,” J. Appl. Data Sci., vol. 5, no. 4, hal. 1790–1801, 2024, doi: 10.47738/jads.v5i4.407.

[18] A. B. Prasetio, M. Aboobaider, dan A. Ahmad, “Assessing Geographic Disparities in Campus Killings : A Data Mining Approach Using Cluster Analysis to Identify Demographic Patterns and Legal Implications,” J. Cyber Law, vol. 1, no. 1, hal. 1–21, 2025.

[19] A. B. Prasetio, M. Aboobaider, dan A. Ahmad, “Machine Learning for Wage Growth Prediction : Analyzing the Role of Experience , Education , and Union Membership in Workforce Earnings Using Gradient Boosting,” Artif. Intell. Learn., vol. 1, no. 2, hal. 153–172, 2025, doi: 10.63913/ail.v1i2.12.

[20] Singh, S., et al., A framework for privacy-preservation of IoT healthcare data using Federated Learning and blockchain technology. Future Generation Computer Systems, 2022. 129: p. 380-388.

[21] Pan, Y., et al., Privacy-Preserving Heterogeneous Personalized Federated Learning with Knowledge. IEEE Transactions on Network Science and Engineering, 2024. 11(6): p. 5969-5982.

[22] C. R. A. Widiawati, Sarmini, dan D. Yuliana, “Predicting Network Performance Degradation in Wireless and Ethernet Connections Using Gradient Boosting, Logistic Regression, and Multi-Layer Perceptron Models,” J. Appl. Data Sci., vol. 6, no. 1, hal. 325–338, 2025, doi: 10.47738/jads.v6i1.519.

[23] M. S. Hasibuan et al., “Integrating Convolutional Neural Networks into Mobile Health: A Study on Lung Disease Detection,” J. Appl. Data Sci., vol. 6, no. 3, hal. 1495–1503, 2025, doi: 10.47738/jads.v6i3.660.

[24] H. T. Sukmana, “Using K-Means Clustering to Enhance Digital Marketing with Flight Ticket Search Patterns,” J. Digit. Mark. Digit. Curr., vol. 1, no. 3, hal. 286–304, 2024, doi: 10.47738/jdmdc.v1i3.22.

[25] D. Mashao dan C. Harley, “Cyber Attack Pattern Analysis Based on Geo-location and Time : A Case Study of Firewall and IDS / IPS Logs,” J. Curr. Res. Blockchain, vol. 2, no. 1, hal. 28–40, 2025, doi: 10.47738/jcrb.v2i1.26.

[26] Abuzied, Y., et al., A privacy-preserving federated learning framework for blockchain networks. Cluster Computing, 2024. 27(4): p. 3997-4014.

[27] Cui, L., et al., Security and Privacy-Enhanced Federated Learning for Anomaly Detection in IoT Infrastructures. IEEE Transactions on Industrial Informatics, 2022. 18(5): p. 3492-3500.

[28] Z. Tian, “Investigation into Data Mining for Analysis and Optimization of Direct Maintenance Costs in Civil Aircraft Operations,” IJIIS Int. J. Informatics Inf. Syst., vol. 7, no. 1, hal. 35–43, 2024, doi: 10.47738/ijiis.v7i1.190.

[29] J. Prayitno, B. Saputra, dan A. Kumar, “Emotion Detection in Railway Complaints Using Deep Learning and Transformer Models : A Data Mining Approach to Analyzing Public Sentiment on Twitter,” J. Digit. Soc., vol. 1, no. 2, hal. 1–14, 2025.

[30] U. Rahardja dan Q. Aini, “Analyzing Player Performance Metrics for Rank Prediction in Valorant Using Random Forest: A Data-Driven Approach to Skill Profiling in the Metaverse,” Int. J. Res. Metaverse, vol. 2, no. 2, hal. 102–120, 2025, doi: 10.47738/ijrm.v2i2.26.

[31] Wang, K.I.K., et al., Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing. IEEE Transactions on Industrial Informatics, 2022. 18(6): p. 4088-4096.

[32] M. S. Hasibuan, R. Z. A. Aziz, D. A. Dewi, T. B. Kurniawan, and N. A. Syafira, “Recommendation Model for Learning Material Using the Felder Silverman Learning Style Approach,” HighTech and Innovation Journal, vol. 4, no. 4, pp. 811–820, Dec. 2023, doi: https://doi.org/10.28991/HIJ-2023-04-04-010.