Sensor Fault Impact on Federated Learning for Electricity Demand Forecasting

Jaewon Jeoung
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
jjw0127@yonsei.ac.kr

Eunseong Song
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
en_s11@yonsei.ac.kr

Juwon Hong
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
juwonae@yonsei.ac.kr

Jongbaek An
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
ajb2577@yonsei.ac.kr

Hyuna Kang
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
hyuna_kang@yonsei.ac.kr

Taehoon Hong
Department of Architecture and
Architectural Engineering
Yonsei University
Seoul, Republic of Korea
hong7@yonsei.ac.kr

***Abstract*—Buildings contribute significantly to global energy consumption, necessitating advanced electricity forecasting methods to enable effective energy management and preserve data privacy. Federated learning (FL) has demonstrated promising performance in energy forecasting by protecting data privacy, yet its resilience against sensor faults remains underexplored. To address this gap, this study proposes an FL framework that integrates a lightweight model with an anomaly detection module to improve electricity forecasting accuracy in the presence of sensor faults. Within this framework, individual buildings preprocess hourly electricity data locally by detecting sensor faults (such as missing values and outliers) using statistical methods (e.g., Z-score and Median Absolute Deviation) and imputing missing data via seasonal-trend decomposition. The Time-Series Mixer (TSMixer) model, which employs multiple perceptron layers, is used to support resource-constrained environments (e.g., smart meters). The proposed framework is evaluated using data from 30 educational buildings in the Building Genome Project 2. The results demonstrate that the framework consistently outperforms traditional approaches, where each building independently trains its own model using only local data. Moreover, the integration of the anomaly detection module yields a 38.2% improvement in prediction accuracy compared to absence of the anomaly detection module. The framework maintains robust generalization with an 8.1% performance degradation when applied to previously unseen buildings. These findings highlight the potential of FL-based forecasting systems to enhance energy management in real-world smart grid applications while ensuring data privacy and fault tolerance.**

Keywords*—*Federated learning, electricity forecasting, anomaly detection, sensor fault tolerance

I. INTRODUCTION

With the urgent global challenge of achieving carbon neutrality, energy management in the building sector, accounting for 30 % of global energy use and one-third of related greenhouse gas emissions, has become a critical issue [1]. In particular, the increasing uncertainty in electricity supply and demand, driven by the expansion of renewable energy adoption and the growing frequency of extreme weather events, has intensified the need for accurate predictions of building energy consumption. As the frequency and intensity of extreme weather events rise, the volatility of building electricity patterns increases, revealing the limitations of traditional rule-based energy management methods. In this context, time-series forecasting of data-driven electricity is gaining attention [2], [3]. A data-driven method not only facilitates demand-side flexibility and proactive demand management, but also significantly contributes to peak load reduction and cost savings through the optimal operation of heating and cooling systems [4].

High-quality data collected from IoT sensors and smart meters are essential for accurate electricity forecasting. However, data privacy has emerged as a major challenge in this field, as high-resolution energy consumption data contain sensitive information that can reveal individuals' daily routines and corporate operational details. For example, analyzing 15-minute interval electricity consumption data has been shown to identify occupant presence and even appliance usage patterns [5]. In response to these privacy concerns, major countries have introduced stringent data protection regulations [6], [7]. For example, the European Union (EU)'s General Data Protection Regulation (GDPR) enforces strict regulations on the collection, processing, and storage of personal data, including energy consumption data [6].

Against this backdrop, Federated Learning (FL) has emerged as a promising solution [8]. FL is an innovative decentralized learning paradigm that enables multiple distributed data sources to collaboratively train machine learning or deep learning models without sharing their raw data. In this approach, each client (e.g., individual buildings or users) trains a local model using its own data and shares only the learned model parameters with a central server. The central server aggregates these local model parameters to update a global model, and this process is repeated to iteratively improve the model's performance. Since raw data never leaves the local environment, data privacy is inherently preserved, while the model benefits from diverse training sources, improving its generalization ability.

Despite the advantages of FL, real-world applications face challenges due to sensor failures, which are inherent in distributed systems. For instance, according to the UK government, as of June 2023, approximately 2.7 million smart meters reported faults [9]. Additionally, by the end of 2023, the proportion of smart meters operating normally was reported to be 88.6% [10]. Faulty sensor data, such as missing values, noise, or incorrect readings, can propagate across the federated model, significantly degrading performance. This underscores the need for a robust fault-tolerant framework that ensures reliable predictions.

Therefore, this study systematically analyzes the impact of sensor faults on FL-based time-series electricity forecasting and proposes an FL framework that integrates an anomaly detection module. The key contributions of this research are threefold. First, it quantitatively evaluates the effect of sensor faults on the prediction performance of FL-based time-series electricity forecasting. Second, it introduces an FL framework that incorporates anomaly detection and fault correction mechanisms, enabling stable model training even in the presence of faulty data. Third, through experiments conducted on real-world energy datasets, this study demonstrates the robustness of the proposed framework across different fault rates (0%, 5%, and 10%), highlighting its effectiveness in improving the reliability of electricity forecasting in practical applications.

II. RELATED WORKS

Electricity consumption profiles of buildings vary significantly due to factors such as occupancy patterns occupancy patterns (e.g., office- versus residential-use) and local climate. When integrating models trained on non-independent and identically distributed (non-IID) data, client drift may occur, causing local models to diverge from the global optimum. To mitigate this issue, several studies have proposed personalized FL, where clients with similar data distributions are grouped into clusters, and a separate global model is trained for each cluster [11], [12], [13]. For instance, Tang et al. dynamically clustered clients in each round based on validation loss [12]. However, a major challenge in this approach lies in determining the optimal number of clusters. An alternative strategy is fine-tuning local models in personalized FL, allowing clients to adjust the model to their unique characteristics [14], [15]. In personalized FL, all clients fine-tune a subset of global model parameters using their local datasets, capturing distinctive electricity patterns across buildings. A common approach involves keeping the final layer independent, so that client-specific patterns are retained while sharing lower layers (feature extractors) across all clients to learn general representations. This technique reduces client drift while maintaining personalized model performance.

Existing studies on FL for electricity forecasting have demonstrated advancements in privacy-preserving techniques. However, they exhibit limitations in addressing deployment challenges on real-world edge devices, particularly in scenarios with missing data. First, while previous studies propose imputation methods (e.g., autoencoders [16], moving average [17]) to handle missing values, there are research gaps on how increasing missing data rates caused by sensor failures affect the FL framework for electricity forecasting. Moreover, many studies in FL either discard outliers or do not mention how they handle missing values [12], [18], [19]. Second, though lightweight models, personalized FL, and missing data imputation are studied individually, no work integrates these components cohesively in FL for edge devices (e.g., smart meters) [14], [20], [21]. Specifically, state-of-the-art (SOTA) lightweight models for time-series forecasting are rarely tested in FL settings. To bridge these gaps, the proposed framework includes the following:

- A seasonal-trend decomposition using Locally Estimated Scatterplot Smoothing (LOESS) to handle missing values and outliers was introduced for resource-constrained edge devices, preserving temporal patterns without heavy computation.
- An empirical study was conducted to quantify how missing data rates (0%, 5%, 10%) degrade FL convergence and prediction accuracy of electricity forecasting model.
- A lightweight time-series electricity forecasting model was introduced within a personalized FL framework.

III. PROPOSED FRAMEWORK

A. Federated-learning framework

FL enables multiple clients to collaboratively train a global model while preserving the privacy of their local data (i.e., electricity data). This process involves iterative communication between a central server and distributed clients, ensuring both data privacy and scalability. The process for the FL framework applied to electricity forecasting consists of four steps (see Figure 1).

1) Initialization: The FL process begins with the server initializing the global model for electricity forecasting. The server then distributes this initialized model to all participating clients, providing each client with a uniform starting point.

2) Local training on clients: Each client independently trains the received model using its private local dataset. This step comprises three substeps. First, to address sensor failures, each client employs an anomaly detection module to preprocess its local data. This module identifies missing values and outliers, ensuring that the training data are free of errors. By preprocessing data locally, clients provide high-quality updates to the server while adapting the global model to their unique data characteristics. Second, each client trains the received model using its preprocessed local dataset. These models are particularly suitable for resource-constrained edge devices, providing an efficient trade-off between computational requirements and predictive accuracy. Through local training, each client captures region-specific electricity patterns, which are then reflected in the parameters sent to the server. Third, upon completion of training, clients compute their local parameters (i.e., gradients). As these parameters do not include raw electricity data, privacy is preserved.

3) Server-side aggregation: At the central server, parameters received from all clients are aggregated to update

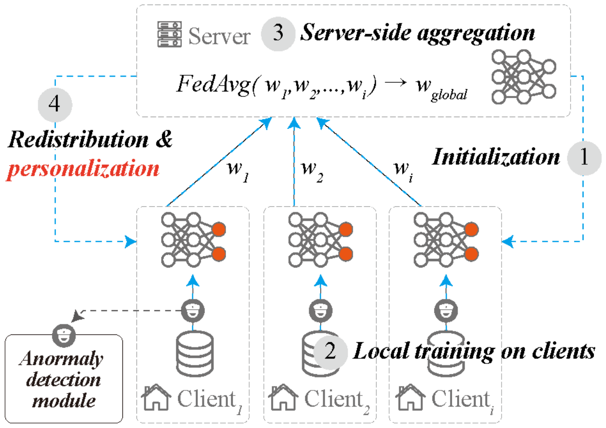

Fig. 1. Overview of the proposed federated learning framework for electricity forecasting.

the global model. The aggregation is performed using the Federated Averaging (FedAvg) algorithm, the most commonly used method in FL [8], which calculates the weighted average of client updates (see Equation 1):

$$w_{global} = 1/N * \sum w_i \qquad (1)$$

where w_i denotes the local update from client i, w_{global} denotes the aggregated weights of the global model, and N is the total number of participating clients.

4) Redistribution and personalization: The global model is updated based on the aggregated parameters and then redistributed to all participating clients for the next round of training. Upon receiving the updated global model's weights, each client reloads its previously trained personalized layer (i.e., the final layer of the electricity forecasting model). This approach ensures that client-specific patterns, including building-specific electricity characteristics, are retained. Since the final layer is trained using the local dataset, preserving it prevents the loss of personalized adaptations while still benefiting from the improvements of the global model. This iterative process continues until the final round is reached.

B. Lightweight model for electricity forecasting

The proposed framework employs the Time-Series Mixer (TSMixer) model on both client and server [22]. TSMixer is a lightweight architecture designed to efficiently capture complex time-series patterns using Multi-Layer Perceptrons (MLPs), making it well-suited for resource-constrained edge devices. In this study, the TSMixer model comprises 0.214 million parameters and has a total size of 0.819MB. Its architecture comprises three main components (see Figure 2).

- Time mixing block: The time mixing block focuses on aggregating information across the time dimension. It transforms the input sequence by applying a series of linear mappings, normalization, and non-linear activations (i.e., ReLU). A residual connection adds the original temporal information back to the transformed output, ensuring effective learning of temporal dependencies.

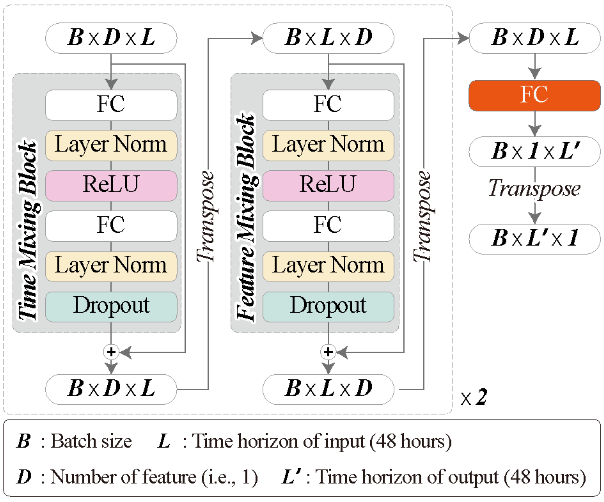

Fig. 2. Architecture of TSMixer for time-series electrcitiy forecasting.

- Feature mixing block: The feature mixing block operates along the feature dimension. It performs a similar series of linear transformations, normalization, and activation functions. Additionally, a residual connection is incorporated to preserve feature-specific details.
- Projection layer: After processing by the mixing blocks, the projection layer maps the refined internal representation to the desired forecasting horizon. In the FL framework, this layer serves as a personalized layer for each client.

C. Anomaly detection module

The anomaly detection module is implemented on the client side to ensure that data preprocessing occurs before local model training. This module is responsible for detecting and addressing two major types of anomalies: missing values and outliers. By identifying these anomalies and applying appropriate imputation strategies, the module enhances data quality, thereby improving model performance.

To identify missing values and outliers, two complementary statistical techniques are employed. Missing values are directly identified as data points lacking recorded values, whereas outliers are determined using a combination of statistical measures. The first approach utilizes the Z-score method, which evaluates the deviation of each observation from the mean based on the standard deviation of the dataset. Data points with Z-scores greater than 3.0 are considered outliers. The second approach leverages the Median Absolute Deviation (MAD) method, where values deviating beyond 3.5 times the MAD from the median are flagged as outliers. While the Z-score method captures deviations relative to the overall data distribution, the MAD method is particularly effective in handling extreme values, ensuring a more comprehensive anomaly detection process.

For imputation, seasonal-trend decomposition using LOESS (STL) is applied [23]. The process begins with an initial estimation of missing values using linear interpolation. The interpolated time series is then decomposed into trend, seasonal, and residual components, considering a reference period of 24 hours. Each component is estimated iteratively using LOESS, ensuring smooth and reliable reconstruction of the time series. Finally, these components are recombined to generate refined values that preserve the underlying data patterns, minimizing the impact of missing values and outliers on subsequent analyses.

IV. EXPERIMENTS

A. Dataset

The building electricity dataset from the Building Genome Project 2 [24] was used to evaluate the proposed framework, which comprises hourly electricity consumption data collected from 30 educational buildings in the Panther region during the period from 1 January to 31 December 2017. To evaluate the forecasting framework, the dataset was partitioned into a training set, an in-building test set, and an out-of-building test set (see Figure 3).

- Training set: Data from 28 buildings covering the period from 1 January to 25 November 2017 (90% of the total period) were used for training the model.

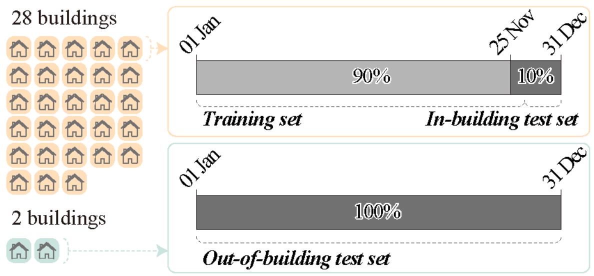

Fig. 3. Composition of the training set and two test sets.

- In-building test set: To evaluate the performance of the FL framework on participating clients, data from the same 28 buildings used for training, from 26 November to 31 December 2017 were reserved as the in-building test set.
- Out-of-building test set: To assess the adaptability of the global model to new environments, data from two buildings not included in the training phase were used. These buildings contributed data for the entire year (1 January to 31 December 2017), enabling the evaluation of the FL framework's performance on previously unseen clients (i.e., new buildings).
- Task: The time-series forecasting task involves predicting electricity for the next 48-hour period using the preceding 48 hours of data. This design is intended to capture both short-term fluctuations and long-term usage trends inherent in the hourly measurements.

To evaluate the robustness of the proposed framework and mimic real-world sensor fault conditions, missing values (NaN) were artificially introduced into the training set at three fault levels: 0%, 5%, and 10%. For the 5% and 10% scenarios, the missing values were randomly distributed across different time intervals for each building. Any pre-existing outliers and missing values in the original dataset were left unmodified.

B. *Evaluation metrics*

The performance of the model was evaluated using standard time-series forecasting metrics, specifically Mean Absolute Error (MAE). MAE measures the average absolute difference between actual and predicted values, providing a direct and interpretable measure of prediction accuracy. A lower MAE indicates that the model's predictions closely follow the actual electricity consumption patterns.

C. *Implementation details*

The experiments were conducted on a local machine equipped with an Intel Core i7-12700KF CPU and an NVIDIA GeForce RTX 3060 GPU. The Flower Python library was utilized to implement the FL framework, while PyTorch 2.2.1 was used for model training and optimization.

Regarding the FL-related parameters, the number of server rounds was set to 60, with each model trained for one epoch before transmitting parameters to the central server. The learning rate followed a cosine annealing schedule, with a minimum learning rate of 1e-6, a peak learning rate of 1e-4, and a warmup phase of four epochs. The model was optimized using the Adam optimizer, and the Mean Squared Error (MSE) was employed as the loss function. The batch size was set to 32.

V. RESULTS AND DISCUSSIONS

This study evaluates the proposed FL framework through two experiments. The first experiment focuses on the in-building test set, comparing the FL framework against a traditional approach. In the traditional approach, each building independently trains its model using only local data. The second experiment tests the framework's generalizability using an out-of-building test set, which comprises data from buildings not involved in the training phase. The experiment is designed to assess how well the global model adapts to entirely new environments, highlighting the robustness and scalability of the FL framework.

A. *Performance comparison on In-building test set*

Figure 4 compares the performance of the FL framework with that of the traditional single-building training approach. On the in-building test set, the FL framework consistently outperformed the traditional approach across all fault rates, showing improved robustness to sensor failures. At 0% fault rate (baseline scenario), the the FL framework achieved an MAE of 19.33, compared to 20.49 for the traditional approach, representing a 6.0% improvement. This improvement in performance can be attributed to the collaborative learning nature of the FL framework, where the model benefits from diverse building patterns while maintaining privacy. The advantage of the FL framework

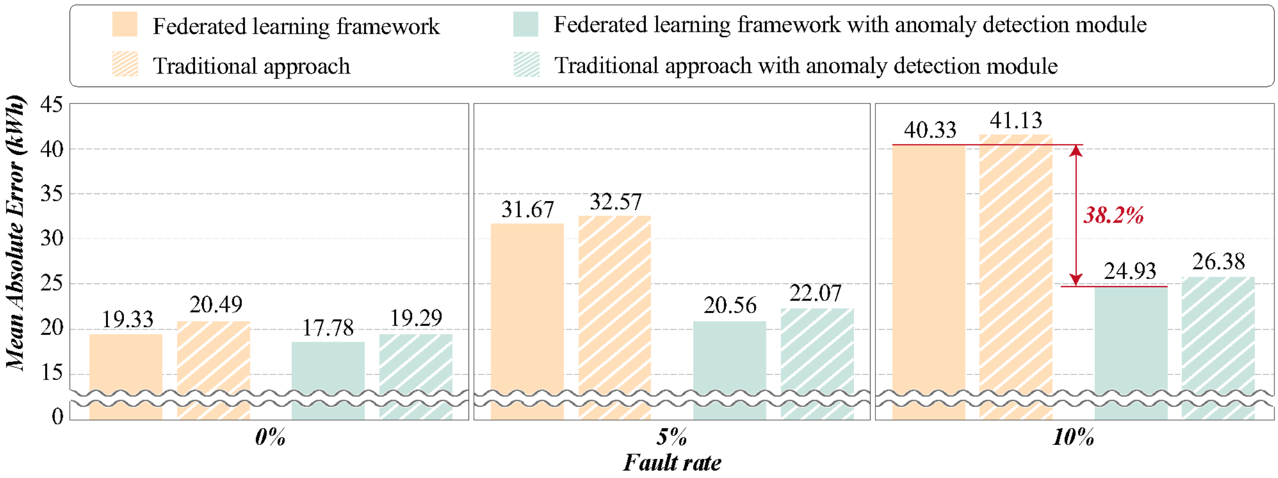

Fig. 4. Performance of the proposed federated learning framework with anomaly detection module against traditional single-building training approaches under three different levels of artificially introduced missing values.

became more pronounced as fault rates increased. At 5% fault rate, the FL framework maintained an MAE of 20.56, while the traditional approach its MAE increased to 22.07. The most significant difference was observed at 10% fault rate, where the FL framework achieved an MAE of 26.38 compared to 24.93 for the traditional approach, demonstrating superior resilience to sensor failures. Notably, the incorporation of the anomaly detection module significantly improved performance across all scenarios. With the module enabled at a 10% fault rate, the MAE was reduced from 40.33 to 24.93, representing a 38.2% improvement. This substantial improvement validates the effectiveness of imputation strategies in handling missing values and outliers.

B. *Generalization to new buildings*

Figure 5 presents the ability of the FL framework to generalize to unseen buildings using the out-of-building test set. In addition, Figure 6 presents a representative example of the electricity forecasting results. The results show stability across different fault rates, with MAE values ranging from 10.84 (0% fault rate) to 11.79 (10% fault rate), representing an 8.1% degradation in performance despite the increase in fault rate. The anomaly detection module proved equally effective for new buildings, maintaining consistent performance even at higher fault rates. With the anomaly detection module enabled, the MAE increased marginally from 10.84 (0% fault rate) to 10.93 (10% fault rate), demonstrating robust generalization capabilities.

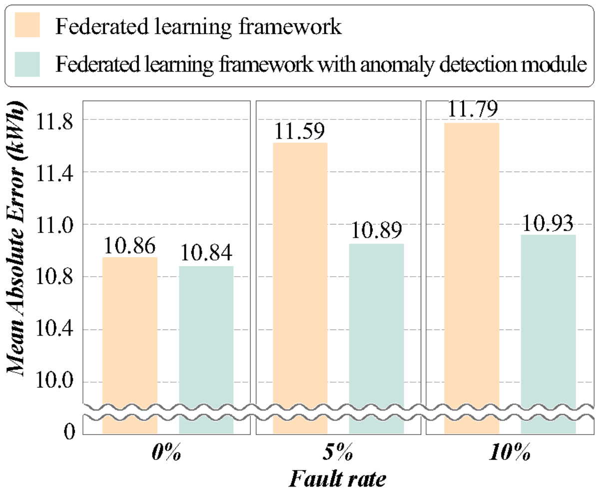

Fig. 5. Mean absolute error across different fault rates for unseen buildings.

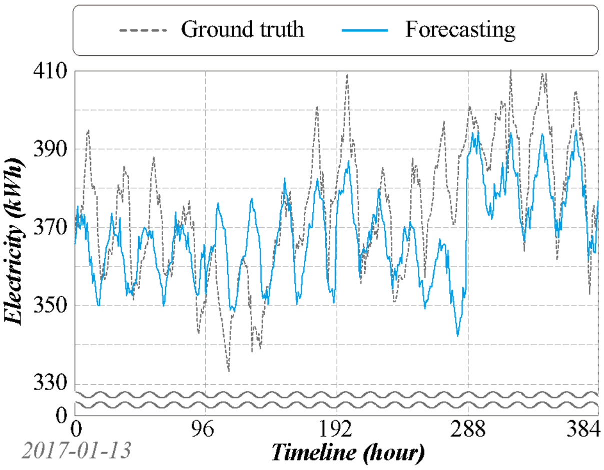

Fig. 6. Example of the electricity forecasting results for unseen buildings.

C. *Discussions*

The experimental results reveal several key insights about the relationship between fault rates and accuracy. First, both the FL framework and the traditional approach showed performance degradation with increasing fault rates. However, the FL framework exhibited lower MAE, indicating better fault tolerance. Second, the anomaly detection module's impact became more pronounced at higher fault rates. At a 10% fault rate, this module reduced MAE by 38.2% for in-building predictions, highlighting its crucial role in maintaining model reliability. Third, the consistent performance on the out-of-building test set, even at higher fault rates, suggests that the FL framework can effectively scale to new buildings without significant performance degradation. This is particularly important for real-world deployments where the system needs to accommodate new buildings over time. Moreover, the experimental results have several important implications for practical deployment. The lightweight nature of the TSMixer model, combined with local anomaly detection, makes it suitable for edge device deployment. The superior performance of the FL framework compared to traditional approaches demonstrates that effective electricity forecasting can be achieved while maintaining data privacy, addressing a key concern in building electricity management. The significant impact of the anomaly detection module suggests that investing in robust data preprocessing on edge devices can substantially improve prediction accuracy, potentially reducing the frequency of sensor maintenance requirements.

VI. CONCLUSIONS

This study presents the FL framework for building electricity forecasting that effectively handles sensor failures in real-world deployments. Experiments across 30 educational buildings demonstrated that the FL framework outperforms traditional approaches under normal conditions while maintaining robust performance with up to 10% fault rates. The integrated anomaly detection module significantly improved prediction accuracy, reducing MAE by 38.2% for in-building predictions at high fault rates. The framework also showed strong generalization capability to new buildings, with an 8.1% performance degradation at increased fault rates.

Future work could explore the framework's adaptability to different building types beyond educational facilities and investigate its performance under more diverse fault scenarios. Additionally, research into dynamic client clustering methods could further improve the framework's ability to handle non-IID data distributions across different buildings. The promising results of this study lay the groundwork for privacy-preserving, fault-tolerant electricity forecasting systems that can be practically deployed at scale.

ACKNOWLEDGMENT

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (Ministry of Science and ICT) (RS-2024-00357129). It was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (Ministry of Science and ICT) (NRF-2021R1A3B1076769).

REFERENCES

[1] IEA, "IEA (2022), Buildings, IEA," International Energy Agency.

[2] J. Bedi and D. Toshniwal, "Deep learning framework to forecast electricity demand," Appl Energy, vol. 238, 2019, doi: 10.1016/j.apenergy.2019.01.113.

[3] B. Yildiz, J. I. Bilbao, and A. B. Sproul, "A review and analysis of regression and machine learning models on commercial building electricity load forecasting," 2017. doi: 10.1016/j.rser.2017.02.023.

[4] O. Aponte and K. T. McConky, "Forecasting an electricity demand threshold to proactively trigger cost saving demand response actions," Energy Build, vol. 268, 2022, doi: 10.1016/j.enbuild.2022.112221.

[5] G. F. Angelis, C. Timplalexis, S. Krinidis, D. Ioannidis, and D. Tzovaras, "NILM applications: Literature review of learning approaches, recent developments and challenges," 2022. doi: 10.1016/j.enbuild.2022.111951.

[6] General Data Protection Regulation, "General Data Protection Regulation (GDPR) – Official Legal Text," General Data Protection Regulation.

[7] Korea Institute of Energy Technology, "Privacy Policy," Korea Institute of Energy Technology Official Website. Accessed: Feb. 06, 2025. [Online]. Available: https://www.kentech.ac.kr/

[8] H. Brendan McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 2017.

[9] Jack Simpson, "New blow to British smart meter rollout as number of faulty machines leaps to 4m," The Guardian, Mar. 26, 2024. Accessed: Feb. 06, 2025. [Online]. Available: https://www.theguardian.com/business/2024/mar/26/smart-meter-rollout-number-faulty-machines-leaps-great-britain

[10] Adarsh Krishnan, "Smart electricity meter market 2024: Global adoption landscape," IoT Analytics. Accessed: Feb. 06, 2025. [Online]. Available: https://iot-analytics.com/smart-meter-adoption/

[11] Y. Wang, J. Ma, N. Gao, Q. Wen, L. Sun, and H. Guo, "Federated fuzzy k-means for privacy-preserving behavior analysis in smart grids," Appl Energy, vol. 331, 2023, doi: 10.1016/j.apenergy.2022.120396.

[12] L. Tang, H. Xie, X. Wang, and Z. Bie, "Privacy-preserving knowledge sharing for few-shot building energy prediction: A federated learning approach," Appl Energy, vol. 337, 2023, doi: 10.1016/j.apenergy.2023.120860.

[13] A. Dogra, A. Anand, and J. Bedi, "Consumers profiling based federated learning approach for energy load forecasting," Sustain Cities Soc, vol. 98, 2023, doi: 10.1016/j.scs.2023.104815.

[14] D. Qin, C. Wang, Q. Wen, W. Chen, L. Sun, and Y. Wang, "Personalized Federated DARTS for Electricity Load Forecasting of Individual Buildings," IEEE Trans Smart Grid, vol. 14, no. 6, 2023, doi: 10.1109/TSG.2023.3253855.

[15] M. G. Arivazhagan, V. Aggarwal, A. K. Singh, and S. Choudhary, "Federated Learning with Personalization Layers," CoRR, vol. abs/1912.00818, 2019, [Online]. Available: http://arxiv.org/abs/1912.00818

[16] R. Shrestha et al., "Anomaly detection based on LSTM and autoencoders using federated learning in smart electric grid," J Parallel Distrib Comput, vol. 193, p. 104951, 2024, doi: 10.1016/j.jpdc.2024.104951.

[17] P. Riedel, K. Belkilani, M. Reichert, G. Heilscher, and R. von Schwerin, "Enhancing PV feed-in power forecasting through federated learning with differential privacy using LSTM and GRU," Energy and AI, vol. 18, p. 100452, 2024, doi: 10.1016/j.egyai.2024.100452.

[18] L. Xu et al., "Swarm Learning for temporal and spatial series data in energy systems: A decentralized collaborative learning design based on blockchain," Appl Energy, vol. 381, p. 125053, 2025, doi: 10.1016/j.apenergy.2024.125053.

[19] Y. Lu, L. Cui, Y. Wang, J. Sun, and L. Liu, "Residential Energy Consumption Forecasting Based on Federated Reinforcement Learning with Data Privacy Protection," CMES - Computer Modeling in Engineering and Sciences, vol. 137, no. 1, 2023, doi: 10.32604/cmes.2023.027032.

[20] H. T. Truong et al., "Light-weight federated learning-based anomaly detection for time-series data in industrial control systems," Comput Ind, vol. 140, 2022, doi: 10.1016/j.compind.2022.103692.

[21] R. Wang et al., "An adaptive federated learning system for community building energy load forecasting and anomaly prediction," Energy Build, vol. 295, 2023, doi: 10.1016/j.enbuild.2023.113215.

[22] S.-A. Chen, C.-L. Li, N. Yoder, S. O. Arik, and T. Pfister, "TSMixer: An All-MLP Architecture for Time Series Forecasting," 2023. [Online]. Available: https://arxiv.org/abs/2303.06053

[23] R. B. Cleveland, W. S. Cleveland, J. E. McRae, and I. Terpenning, "STL: A Seasonal-Trend Decomposition Procedure Based on Loess (with Discussion)," J Off Stat, vol. 6, 1990.

[24] C. Miller et al., "The Building Data Genome Project 2, energy meter data from the ASHRAE Great Energy Predictor III competition," Sci Data, vol. 7, no. 1, 2020, doi: 10.1038/s41597-020-00712-x.