Privacy-Utility-Efficiency Trade-Offs in Personalized Federated Learning for Edge Computing

Clipping-Pressure Diagnostics and Pareto Operating Points for Privacy-Preserving Edge Learning

Mahavir Teraiya
Department of Computer Engineering, Marwadi University, Rajkot, GJ, India
mahavir.teraiya120510@marwadiuniversity.ac.in (corresponding author)

Madhu Shukla
Department of CSE–AI, ML & DS, Marwadi University, Rajkot, GJ, India
madhu.shukla@marwadieducation.edu.in

Received: 28 December 2025 | Revised: 21 January 2026 and 7 February 2026 | Accepted: 8 February 2026

 DOI: https://doi.org/10.48084/etasr.17222

ABSTRACT

For the edge applications of Industrial Internet of Things (IIoT), the task of learning from distributed and privacy-conscious data needs to be conducted under constrained communication resources and in the presence of highly heterogeneous clients. This paper offers a personalized Federated Learning (FL) solution where the global backbone can be decoupled from the per-client heads, enabling the concurrent execution of global representation learning and local adaptation, even with non-Independent and Identically Distributed (non-IID) data. Differential Privacy (DP) is used to ensure privacy protection, where the updates of the backbone are differentially privatized via clipping and Gaussian noise, with the total privacy budget (ε, δ) measured over the number of communication rounds via a Rényi Differential Privacy (RDP) accountant. This study also provides a set of training diagnostics related to clipping pressure, including the proportion of clipped clients and the dynamics of the norms of the updates, which can identify the points at which privacy noise begins to dominate the learning process, along with privacy-friendly operating configurations. Experimental results on the LEAF and FEMNIST datasets reveal higher average accuracy and a narrower per-client accuracy distribution than those of the FL counterparts, whereas the privacy–utility–efficiency analysis identifies Pareto points at which privacy can be improved with a constant per-round communication cost, since only the backbone updates are communicated. In practice, this means more reliable performance at the client level for controlled privacy loss and predictable communication overheads, which makes this approach appropriate for edge sites with varying quality and availability of data.

Keywords-edge computing; Federated Learning (FL); Differential Privacy (DP); personalized learning; non-IID data

I. INTRODUCTION

Edge computing has seen increased use in the area of Industrial Internet of Things (IIoT) to allow low-latency analytics and intelligence to be pushed closer to where sensors, equipment, and the manufacturing environment exist. However, learning accurate models in these scenarios is challenging because the data are distributed, privacy-sensitive, and highly heterogeneous. Each device experiences different conditions, workloads, and noise distributions, causing the local data to be non-Independent and Identically Distributed (non-IID). A natural fit for this type of learning problem is Federated Learning (FL), which enables collaborative learning without requiring the actual data to be shared. Nevertheless, conventional FL approaches, such as FedAvg, face significant challenges with non-IID data, which can worsen when privacy mechanisms are introduced.

This paper tackles a fundamental reality of edge FL: the three competing goals of personalization, privacy, and edge efficiency must be balanced. Figure 1 illustrates this triangle of competing goals. Personalization, exemplified by client-specific heads or adaptation, enhances robustness in the presence of heterogeneity but alters what is shared and how much is revealed. Differential Privacy (DP) formally protects shared elements by adding noise, yet the noise and clipping can

disproportionately affect clients, often further reducing the performance of the worst-performing clients even in IID settings. Budget constraints on edge efficiency, including the number of communication rounds, clients per round, and local steps, further limit resources available for other factors.

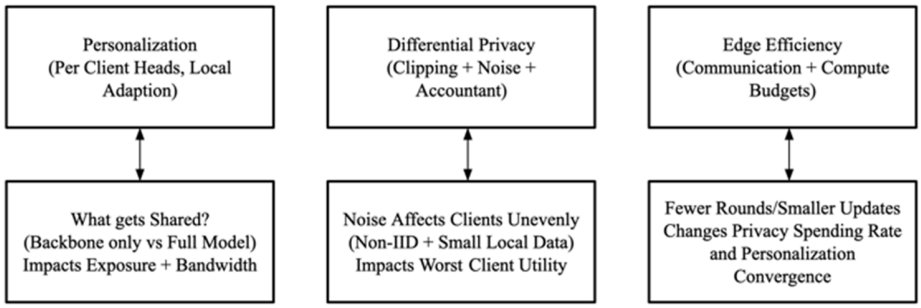

Fig. 1. Interaction of personalization, privacy, and edge efficiency in non-IID edge FL.

Figure 2 further illustrates the causal couplings underlying our design: sharing less (for example, sharing just the representation rather than the full model) reduces bandwidth use and preserves privacy, but edge resources are constrained by training dynamics, including the number of rounds, local steps, and client participation. These constraints directly affect the overall privacy budget (ε, δ) via the accountant and influence the model's personalization rate. Thus, a trade-off between privacy and accuracy is often necessary, as stronger privacy comes at the cost of increased noise and potentially more training steps.

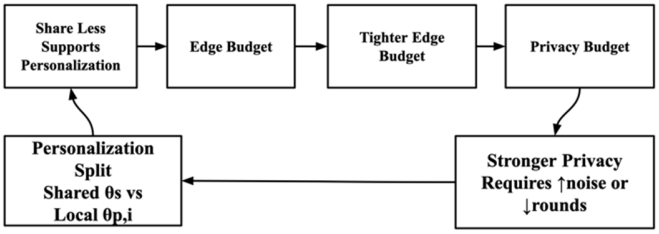

Fig. 2. Edge budget impact on privacy spending and personalization convergence.

To address this challenge, we build a personalized, privacy-preserving FL system for edge computing using a split model architecture, in which a shared backbone is learned while client-side models are protected. The system enforces client-level DP for shared updates and incorporates a privacy auditor for reporting privacy. To better assess FL in this context, we analyze not only average metrics such as accuracy but also individual client performance and Pareto optima.

The proposed work focuses on a FL scenario where the data are non-IID, privacy-aware, and communicated under a limited budget. The aim is to develop an FL approach that ensures accurate performance for individual clients, as well as satisfactory performance for tail clients, under the privacy budget of (ε, δ)-DP within a limited number of communication rounds.

This paper proposes a customized FL architecture in which the global shared backbone is distinguished from the per-client head to enable non-IID edge diversity without sharing parameters for personalization. It enforces client-level DP only on the global shared updates via norm clipping and Gaussian noise addition and measures privacy through Rényi Differential Privacy (RDP) composition over training iterations. The study also introduces clipped-client fraction and update-norm dynamics for diagnosing training stress induced by DP. Lastly, it conducts an analysis on privacy–utility–efficiency on the LEAF and FEMNIST datasets through accuracy, per-client accuracy distribution, tail performance, and communication overhead to determine the deployable operating points.

Most previous approaches consider edge efficiency and privacy budgeting independently. IDMO advances edge efficiency while preserving 91.43% accuracy on the CIFAR-10 dataset by integrating structured pruning, quantization, and selective updates. However, the method targets IID clients, is mostly simulated, and, interestingly, recommends secure aggregation or DP as future work [1]. Neuro Fed-Light TCN is another system combining edge FL with a light Temporal Convolutional Network (TCN) to achieve 97.11% accuracy with 56 ms latency for Electroencephalography (EEG) seizure detection. Nevertheless, privacy considerations mainly focus on the decentralized architecture, and the method is highly application-dependent [2]. ATCP enhances edge privacy budgets for split learning using adversarial training and channel pruning to defend against reconstruction attacks. Yet, the approach is locked to split learning threat models and needs additional attack coverage and real-device cost assessment [3]. Rep Sys with asynchronous FL and blockchain improves communication costs and trust but also introduces verification and incentive system complexity, with "lightweight and scalable proof mechanisms and enhanced security" suggested for future work [4]. Cloud-ELastic selection for FedAdaSS improves parameter-server selection for edge clouds and convergence efficiency in simulations (~12–20% reduction in rounds-to-accuracy), although it is mostly simulated and focuses on heterogeneity and engagement rather than formal privacy budgets [5]. In contrast, our work integrates edge budgets, personalization with shared and local parameters, and formal DP budgets using RDP composition to deliver a controlled (ε, δ) privacy–cost–utility Pareto-optimal point for edge scenarios.

Across the next set of studies, privacy is often ensured via encryption, decentralization, authentication, or perturbation. However, most approaches do not provide a rigorous end-to-end privacy budget that can be composed over multiple rounds under edge constraints. Authors in [6] combine homomorphic encryption with FL and demonstrate privacy protection against inversion/gradient attacks, but the encryption/decryption latency may be prohibitive for edge devices. Authors in [7] propose differentially private federated learning with flexible privacy budget allocation, showing that controlled privacy spending improves accuracy but still faces leakage risks and edge-related trade-offs. Authors in [8] propose a vehicular edge authentication scheme that improves security by adding noise but incurs additional computation and communication, focusing on secure access rather than learning-time privacy accounting. Authors in [9] introduce the tMK-CKKS scheme that minimizes homomorphic encryption overheads and enables client drop-in/drop-out but cannot fully protect the global model. Finally, authors in [10] present the DISTPAB family of tasks, which defend against inference and model inversion attacks via distributed perturbations, yet face efficiency and attack-surface challenges in Distributed Machine Learning (DML). In contrast, the focus of our paper is an edge-aware FL framework with RDP composition, producing (ε, δ) and configurable control knobs (q, σ, S, T) that manage the trade-offs between privacy, utility, and edge budget constraints in a manner that is analytically expressible and tractable.

Authors in [11] focus on vehicular edge FL and perform selective aggregation based on image quality (motion blur) and computing ability, using a contract-theory approach to identify "fine" clients and improve accuracy and efficiency on MNIST and BelgiumTSC in FedAvg. The main limitation of this work is that privacy is addressed implicitly (the server is unaware of certain client attributes) and it lacks explicit (ε, δ) privacy budgeting. In contrast, our work presents an edge-conscious FL framework with explicit privacy budgeting, where composition is applied per round to achieve (ε, δ)-DP.

In other studies, privacy on the edge is treated differently. Authors in [12] propose an asynchronous aerial-aided FL framework incorporating an adaptive DP process (gradient clipping and Gaussian noise with a gradually increasing noise scale) to mitigate gradient leakage. However**,** this approach heavily depends on hyperparameters and does not provide a strong, turn-by-turn privacy accountant. Authors in [13] formulate edge learning as cost-effective DP distributed learning and design a double-layer auction mechanism to decide the number of iterations and corresponding privacy budgets within an overall budget, although the budgeting is tied to market/auction models and not generalizable for deployable FL rounds. Authors in [14] integrate blockchain-based FL, WGAN, and DP to enhance security and privacy. Yet, the additional blockchain and GAN overheads make the privacy–utility trade-off less directly manageable from the accountant's perspective. Conversely, authors in [15] focus on IIoT intrusion detection, and M16 addresses vehicular offloading in MEC using federated deep RL/MADDPG, largely relying on "keeping data local" to preserve privacy. However, these approaches remain vulnerable to potential data leakage in model updates, as no (ε, δ) budgeting is provided. Authors in [16] specifically address bandwidth-limited edge over-the-air FL (OTA-FL) with privacy-informed design, including dimensionality reduction and aggregation/perturbation decisions. While theoretically grounded, this work is tailored for band-limited OTA scenarios and is not intended for general-purpose privacy auditing across multiple rounds or arbitrary subsampling ratios. Relative to these works, the core contribution of our paper is an edge-aware DP pipeline that integrates per-round privacy into RDP and translates it into (ε, δ) -DP, linking privacy guarantees to real-world edge conditions, including communication and computational resources, personalization decisions, and number of training steps.

Recent research on FL for edge/IIoT emphasizes that edge FL is limited not only by non-IID data distributions but also by device heterogeneity, availability, incentives, and privacy-efficiency trade-offs. Authors in [17] propose NebulaFL, a multilayer, hierarchical FL framework for edge environments, using hybrid synchronous-asynchronous training with adaptive load tuning to improve latency and training efficiency. Studies focusing on on-demand client deployment argue that edge FL must support runtime availability for suitable clients, and propose containerized or orchestrated client formation via Docker and Kubernetes to dynamically mobilize additional clients when local ones are insufficient [18]. Research addressing client quality and contribution incentives develops pricing mechanisms using yardstick pricing and Stackelberg game theory to better align data transmission power and corresponding incentives [19]. Finally, studies on privacy-focused edge FL demonstrate improved utility under strict privacy constraints by integrating lottery ticket pruning and DP variants, reducing overall computation and communication costs while maintaining acceptable accuracy [20].

II. MOTIVATION

In the context of IIoT, edge computing platforms produce large amounts of privacy-sensitive, device-specific data (e.g., machine health data, event logs, camera feeds) that cannot be centrally aggregated for analysis without the permission of their owners. However, applying FL, which aggregates model updates rather than raw data, to improve models on these platforms is challenging for two main reasons: (i) non-IID heterogeneity, where device conditions can differ substantially, including operating conditions, workloads, sensor characteristics, and other device-specific factors; and (ii) resource constraints. In the case of heterogeneity, a learning method may improve the average accuracy, yet variability among devices can remain high.

Additionally, FL at the edge must provide formal privacy guarantees, as model updates may reveal sensitive information about local data. DP is an effective approach, but clipping and noise addition can disproportionately affect clients in the tail of the performance distribution, exacerbating the worst-case utility, which ironically needs privacy protection in the first place. These challenges create an intertwined optimization problem involving three objectives: personalization for handling data heterogeneity, privacy to limit leakage, and edge efficiency to satisfy deployment constraints. Therefore, this work explores the privacy–utility–efficiency trade-offs for

personalized FL architectures with client-level DP applied only to the shared backbone updates.

III. METHODOLOGY

We focus on an edge computing federation with K clients (edge devices) that are managed by a centralized server. Each client $k \in \{1, \dots, K\}$ holds its own local dataset $D_k = \{(x^j, y^j)\}_{j=1}^{n_k}$, which is not shared. The goal is to learn a model that generalizes well over these non-IID data points while providing privacy assurances for the information contained in the updates shared by the clients. We consider the LEAF/FEMNIST dataset [21], where each client represents a different writer, providing a natural non-IID distribution. In addition, we use the Lorenz curve and Gini coefficient to quantify inequality in the number of samples per client.

A. Split Personalized Model for Edge Federated Learning

To cope with heterogeneity, we adopt a split-personalized architecture consisting of a common backbone θ_s and client-specific heads $\theta_{p,k}$. The prediction function on client k is given by:

$$\hat{y} = f_k(x; \theta_s, \theta_{p,k}) = h_k(g(x; \theta_s);\ \theta_{p,k}) \quad (1)$$

where $g(\cdot)$ is the common representation learned on the federation, and $h_k(\cdot)$ is the lightweight local head maintained on the device. This architecture corresponds to the process illustrated in Figures 3 and 4, where θ_s is broadcast by the server. Each client updates θ_s and $\theta_{p,k}$ locally, and only the update of the common part is qualified for aggregation on the server, whereas $\theta_{p,k}$ is kept on the client for personalized inference.

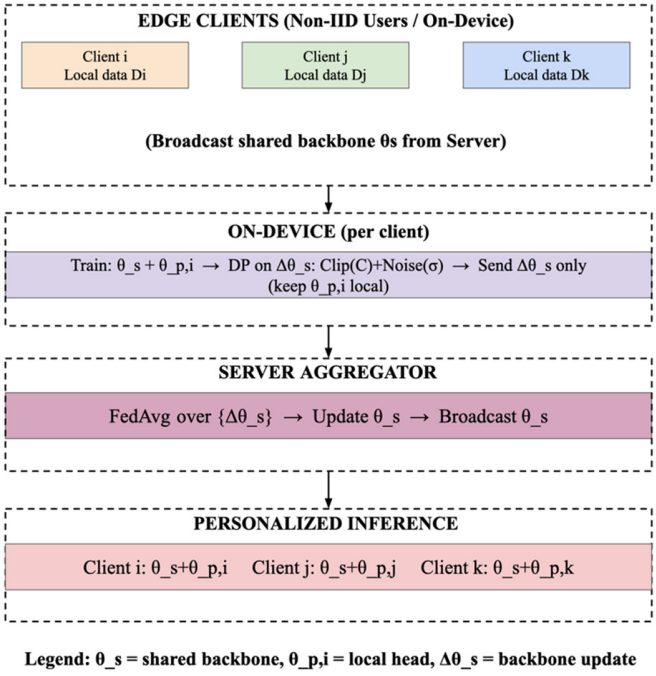

Fig. 3. Split personalized FL with client-level DP on shared backbone updates.

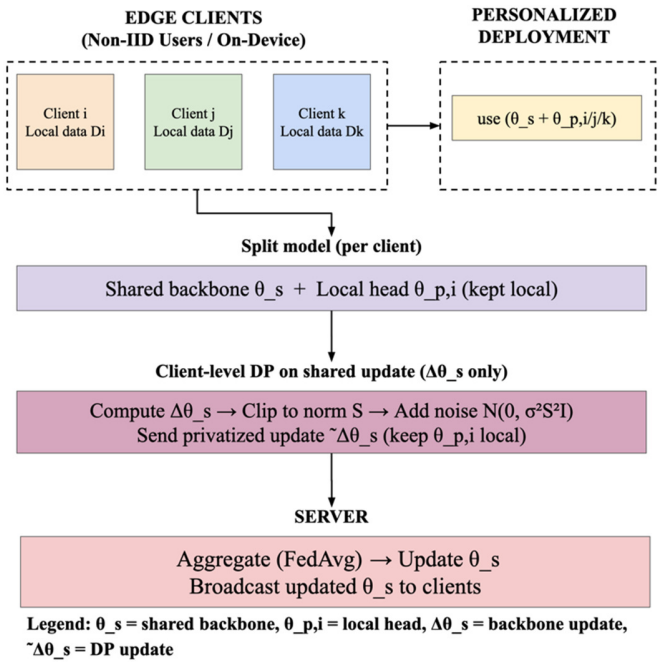

Fig. 4. Training workflow: local heads remain on-device, and privatized backbone updates are aggregated.

B. Local Training Objective

At each communication step t, a subset of clients C_t participates. Each client $k \in C_t$ solves a local optimization problem to minimize its empirical risk:

$$\min_{\theta_s, \theta_{p,k}} \mathcal{L}_k(\theta_s, \theta_{p,k}) = \frac{1}{n_k} \sum_{(x,y)\in D_k} \ell\big(f_k(x; \theta_s, \theta_{p,k}), y\big) \quad (2)$$

and computes an update for the shared backbone, either via delta or gradient update. The personalized head is updated locally and is never communicated.

C. Server Aggregation of the Shared Backbone

The server aggregates only the updates on the shared backbone from participating clients using a weighted average, where the weights correspond to the size of each client's local dataset, as in FedAvg:

$$\theta_s^{t+1} = \theta_s^t + \sum_{k\in C_t} \frac{n_k}{\sum_{j\in C_t} n_j} \Delta\theta_s^{k,t} \quad (3)$$

In this case, $\Delta\theta_s^{k,t}$ represents the update to the shared backbone from client k at iteration t. The formula maintains global generalization and enables specialization on the client side using $\theta_{p,k}$.

D. Client-Level Differential Privacy on Shared Updates

To protect the confidentiality of client data, client-level differential privacy (DP) is applied to $\Delta\theta_s^{k,t}$ before it is sent to the server. Each client update is first clipped to an L2 norm threshold S:

$$\bar{\Delta}\theta_s^{k,t} = \Delta\theta_s^{k,t} \cdot \min\left(1, \frac{S}{\left\|\Delta\theta_s^{k,t}\right\|_2}\right) \quad (4)$$

Next, Gaussian noise is added:

$$\bar{\Delta}\theta_s^{k,t} = \bar{\Delta}\theta_s^{k,t} + \mathcal{N}(0, \sigma^2 S^2 I) \quad (5)$$

Only the privatized $\bar{\Delta}\theta_s^{k,t}$ is sent to the server.

The "DP mechanism stress," including the rate of clipping and changes to update norms, is then analyzed. In this work, we provide a description of the end-to-end privacy achieved via the RDP accountant (Figure 5). Using the client sampling rate q, noise multiplier σ, and total rounds T, we compute per-round RDP starting from orders α, then convert the RDP into an (ε, δ) guarantee using the specified δ. This yields the final privacy budget ε used in the privacy–utility–efficiency trade-off.

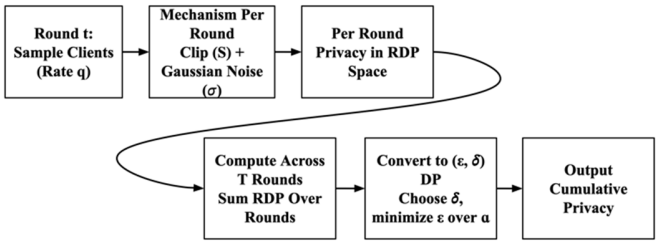

Fig. 5. RDP privacy accountant: per-round DP composition and conversion to (ε, δ).

E. Evaluation Metrics and Analysis

The performance of the models is evaluated on individual clients to reflect realistic deployment at the edge. Client accuracy and tail metrics, such as Bottom-10% accuracy of clients, are reported to enable multi-metric comparison. The points of intervention for personalization can be observed through personalization gains over the baseline and explored using heatmaps. The label entropy for client k is defined as:

$$H_k = -\sum_c p_k(c) \log p_k(c) \quad (6)$$

with $p_k(c)$ representing the empirical class distribution on client k. Edge efficiency is measured in terms of communication per client per round (MB) and the convergence rate in rounds. These metrics, along with privacy and utility, are summarized in method comparisons and the privacy–utility–efficiency trade-off analysis.

We assume an edge FL scenario in which raw data remain on-device, and only model updates are shared. The proposed scheme adopts a split personalized FL scheme with a shared backbone network and client-specific heads. Only the updated backbone network is transmitted. Client-level DP is enforced on shared updates via L2 clipping with threshold S and Gaussian noise with scale σ. Privacy is analyzed using the RDP accountant, with composition over T communication rounds to achieve (ε, δ)-DP. Potential adversaries include the server or any entity intercepting shared models. Local heads and raw data are kept private.

Training proceeds in synchronous rounds with partial participation. In each round, a fraction q of clients is selected. Each client receives the current backbone, locally optimizes it for a fixed number of steps/epochs, privatizes the backbone update, and sends it to the server for weighted aggregation.

We compare our approach to standard and personalized variants of FedAvg, FedEM, and FedPer/Ditto-type personalization, as well as private FedAvg and a DP-resistant version developed in our experiments, all within the same LEAF/FEMNIST federation. Performance is measured using average accuracy, tail performance (Bottom-10% accuracy), and edge efficiency metrics, including communication per client and rounds to target. These metrics are analyzed with respect to practical levels of ε.

IV. RESULTS AND DISCUSSION

All experiments are conducted on the LEAF/FEMNIST federation, where each client represents a different author, yielding a realistic non-IID distribution among edge clients. Before analyzing model performance, we quantify the degree of client data imbalance using the Lorenz curve in Figure 6. The cumulative proportion of clients is plotted on the x-axis and the cumulative proportion of client data on the y-axis, with clients sorted by local dataset size. The orange dashed line in Figure 6 represents the case where all clients contribute equally, whereas the blue line shows the actual Lorenz curve of the federation. The region between these two curves corresponds to the Gini coefficient, which is calculated as 0.220. This result indicates a non-negligible imbalance in client participation, where most clients contribute only a few samples. Such clients are more vulnerable to both non-IID effects and privacy noise in edge FL, and hence we concentrate on client and tail performance in the following analysis.

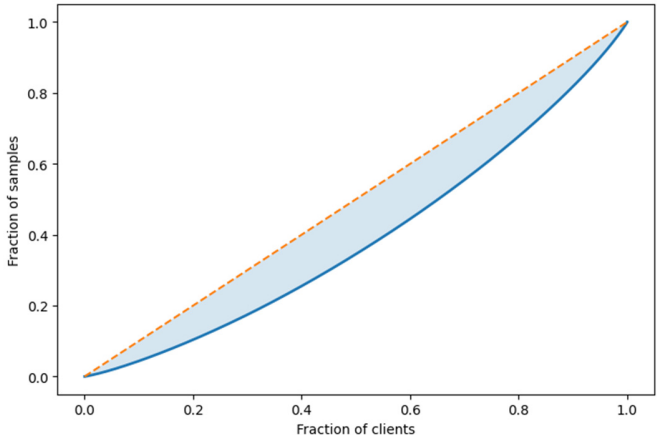

Fig. 6. Lorenz curve of client data imbalance.

Client-level performance is modeled through per-client accuracy distributions instead of being constrained to averaged summaries. As is clear from Figure 7, FedAvg exhibits greater variability and a longer tail of low-accuracy clients, reflecting heterogeneity-induced instability. FedEM improves upon FedAvg in terms of central tendency but still displays visible variation. On the other hand, there is a clear indication of a greater and more tightly distributed per-client accuracy through personalization, which is more desirable from a practical perspective since device reliability is often more important than accuracy improvements.

To identify which clients benefit most from personalization, we analyze gains across heterogeneity bins in Figure 8, where client difficulty is grouped according to local dataset size and label entropy. The heatmap shows that personalization yields substantially larger gains than FedAvg for clients with small datasets. These gains decrease as client data size and label entropy increase, indicating that personalization is less beneficial for clients with highly diverse label distributions. This behavior aligns with IIoT scenarios, where rare events are often associated with limited operating regimes or modes.

Next, we examine behaviors that can be attributed to DP itself in order to understand why different private training methodologies produce different utility levels under a fixed privacy regime. Figure 9 plots the stress of the DP mechanisms through clipping pressure (measured as the fraction of clipped client updates) and the dynamics of update norms. The DP-FedAvg method exhibits consistently higher clipping pressure throughout training, with peaks during the early and middle stages. This behavior indicates that a large fraction of client updates are clipped, thereby removing much of the learning signal before noise is added.

The joint trade-off among utility, tail robustness, privacy, and efficiency is illustrated in Figure 10 for both private and non-private personalization schemes. Non-private personalization baselines occupy the high-utility region, whereas privacy-aware methods illustrate how privacy constraints affect both accuracy and robustness. A strong privacy setting for DP-FedAvg leads to substantial utility degradation, particularly for worst-case (tail) clients. A scheme inspired by Fed-SAM improves this degradation by preserving utility and tail performance without violating privacy constraints.

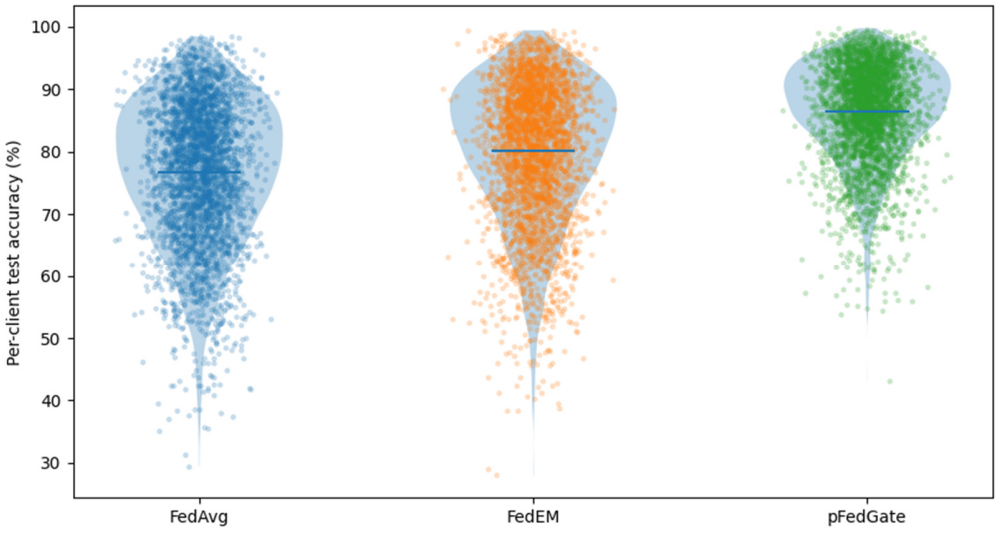

Fig. 7. Per-client accuracy distribution.

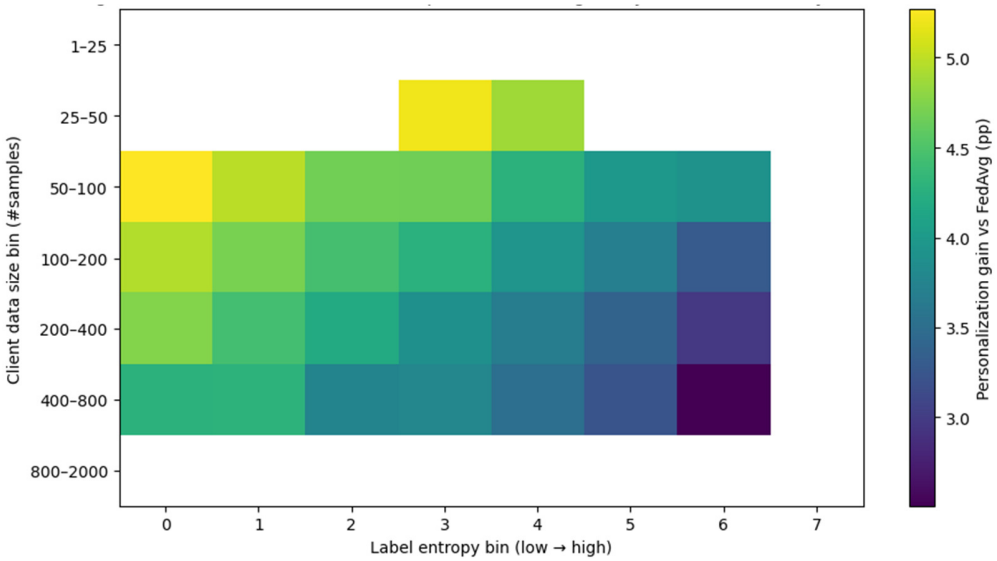

Fig. 8. Gain heatmap across heterogeneity bins.

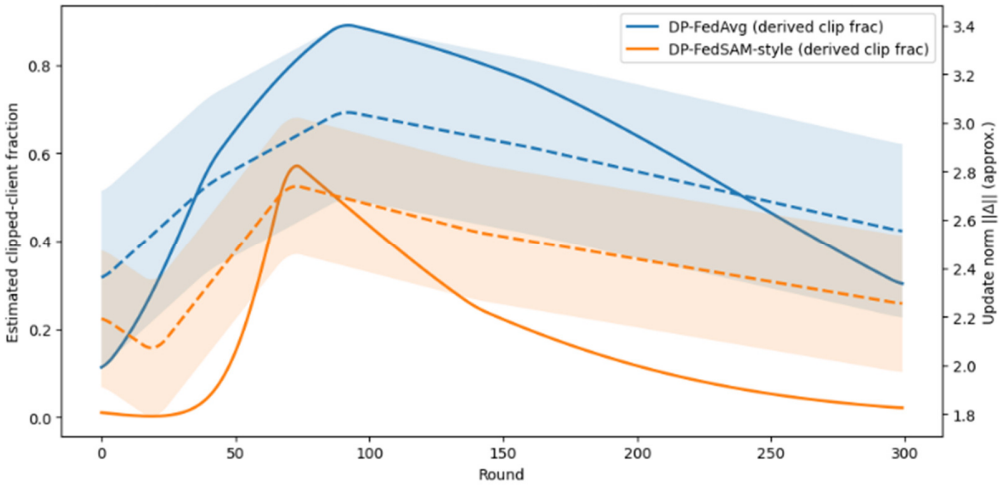

Fig. 9. DP mechanism stress during training.

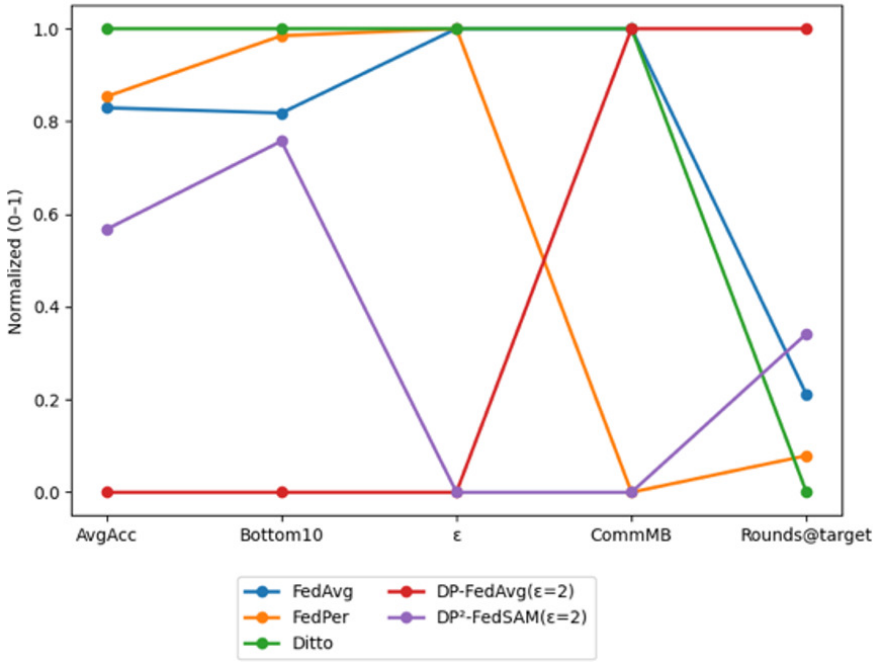

Fig. 10. Multi-metric method comparison.

Figure 11 presents the Pareto distribution with respect to privacy, utility, and communication from a deployment perspective. Accuracy improves as the privacy budget ε is relaxed, whereas communication per client per round remains approximately constant, consistent with the idea of sharing only backbone updates in this implementation. This indicates that, in this design, the dominant control variable for the trade-off is the privacy budget rather than bandwidth. Consequently, edge deployments can select an operating point based on privacy requirements while maintaining a stable communication profile.

Overall, the findings verify that split personalization substantially improves robustness under non-IID edge conditions by raising the entire distribution of per-client accuracies and enhancing the performance on the worst client (Figure 7). The largest gains occur in low-data regimes characteristic of edge environments (Figure 8). Analysis of the DP mechanism further shows that reducing clipping stress is critical for preserving the learning signal, which explains the improved utility observed under DP-friendly training behaviors (Figure 9). Lastly, the multi-metric and Pareto plots illustrate the trade-offs that can be practically explored during the deployment phase, where increasing ε improves accuracy whereas communication cost remains largely unchanged due to backbone-only sharing. This allows the system to be operated at various points along the ε distance depending on the edge system deployment requirements (Figures 10 and 11).

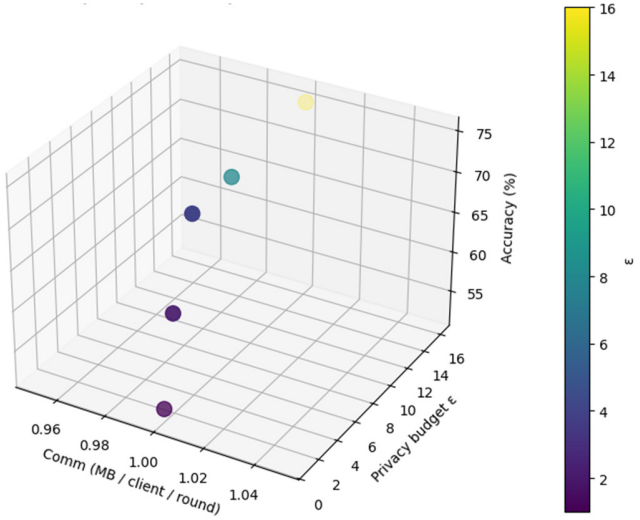

Fig. 11. Privacy–utility–efficiency trade-off.

V. CONCLUSION AND FUTURE SCOPE

This study proposed a privacy-preserving personalized Federated Learning (FL) solution for edge computing, addressing the challenges of non-Independent and Identically Distributed (non-IID) client data using a split-model approach. In this approach, a shared backbone is trained jointly across clients, whereas personalized heads remain localized on each edge device. Formal privacy guarantees are achieved through client-level Differential Privacy (DP) applied to split-model updates communicated across the network. Privacy is accounted for using a Rényi Differential Privacy (RDP) accountant, all while respecting edge resource constraints for practical deployment.

Experiments on the LEAF/FEMNIST datasets demonstrated that personalization improves reliability for individual clients by tightening the per-client accuracy distribution. The privacy–utility–efficiency analysis further highlighted operational regions in which privacy can be adjusted without significantly impacting communication overhead, since only the shared backbone is transmitted per round. Apart from increased accuracy, this personalized DP-FL model is well suited for real-world edge computing and Industrial Internet of Things (IIoT) scenarios, where clients have varying amounts of data. By distinguishing between the shared backbone and local heads, applying DP only to the backbone, client performance is optimized. Additionally, clipping pressure analysis serves as a valuable monitoring tool for evaluating DP-induced training stress.

The novelty of this work lies in the integration of: (i) a split-model personalized FL approach that preserves client specificity through local heads while learning a shared representation through a global backbone; (ii) client-level privacy guarantees via selective DP application on the backbone and end-to-end privacy analysis using RDP; and (iii) extensive experiments that go beyond average accuracy, exploring per-client accuracy distributions, privacy–utility–communication trade-offs, and clipping as a practical metric. Collectively, these contributions provide a practical and theoretically grounded solution for client-level, privacy-preserving personalization in heterogeneous edge and IIoT environments.

Future research may extend this framework to more realistic IIoT settings, incorporating system heterogeneity such as stragglers, connection availability, and energy-aware client selection. Adaptive strategies for clipping thresholds, noise scaling factors, and client participation rates could further improve performance. Other areas of research may include combining secure aggregation with DP to defend against more powerful attackers, personalization strategies for edge cases and concept drift, and larger-scale validation with sensor-driven data to investigate privacy and personalization in the presence of real fault distributions and rare events.

REFERENCES

[1] Ch.Ellaji, R. S. Ponmagal, and V. Saritha, "IDMO: A Multi-Stage Optimized Deep Learning Framework for Efficient and Scalable IoT Big Data Analytics," *International Journal of Electronics and Communication Engineering*, vol. 12, no. 11, pp. 8–20, Nov. 2025, https://doi.org/10.14445/23488549/IJECE-V12I11P102.

[2] Z. Y. Lim, Y. H. Pang, S. Y. Ooi, W. H. Khoh, and Y. J. Chew, "NeuroFed-LightTCN: Federated Lightweight Temporal Convolutional Networks for Privacy-Preserving Seizure Detection in EEG Data," *Applied Sciences*, vol. 15, no. 17, Sept. 2025, Art. no. 9660, https://doi.org/10.3390/app15179660.

[3] A. Alhindi, S. Al-Ahmadi, and M. Maher Ben Ismail, "Balancing Privacy and Utility in Split Learning: An Adversarial Channel Pruning-Based Approach," *IEEE Access*, vol. 13, pp. 10094–10110, 2025, https://doi.org/10.1109/ACCESS.2025.3528575.

[4] X. Sheng, C. Yu, Y. Zhou, and X. Cui, "Reputation-Driven Asynchronous Federated Learning for Optimizing Communication Efficiency in Big Data Labeling Systems," *Mathematics*, vol. 12, no. 18, Sept. 2024, Art. no. 2932, https://doi.org/10.3390/math12182932.

[5] Y. Xu, B. Zhao, H. Zhou, and J. Su, "FedAdaSS: Federated Learning with Adaptive Parameter Server Selection Based on Elastic Cloud Resources," *Computer Modeling in Engineering & Sciences*, vol. 141, no. 1, pp. 609–629, Aug. 2024, https://doi.org/10.32604/cmes.2024.053462.

[6] M. G. Hegde, B. Ruthvika, R. B. Jain, P. D. Shenoy, K. R. Venugopal, and A. Canchi, "A Privacy-Preserving Federated Learning Method with Homomorphic Encryption for Chronic Kidney Disease Stage Prediction," *Engineering, Technology & Applied Science Research*, vol. 15, no. 4, pp. 26019–26026, Aug. 2025, https://doi.org/10.48084/etasr.11928.

[7] W. Qian, Q. Shen, X. Chen, C. Li, Y. Fang, and Z. Wu, "FDP-FL: differentially private federated learning with flexible privacy budget allocation," *The Computer Journal*, vol. 67, no. 12, pp. 3180–3195, Dec. 2024, https://doi.org/10.1093/comjnl/bxae081.

[8] V. S. Sadlapur and N. Hegde, "A Privacy-Preserving Reliable Authentication Scheme for 6G-Enabled Internet of Vehicular Edge Computing Networks," *Engineering, Technology & Applied Science Research*, vol. 15, no. 6, pp. 28810–28817, Dec. 2025, https://doi.org/10.48084/etasr.11534.

[9] W. Du, M. Li, L. Wu, Y. Han, T. Zhou, and X. Yang, "A efficient and robust privacy-preserving framework for cross-device federated learning," *Complex & Intelligent Systems*, vol. 9, no. 5, pp. 4923–4937, Oct. 2023, https://doi.org/10.1007/s40747-023-00978-9.

[10] M. A. P. Chamikara, P. Bertok, I. Khalil, D. Liu, and S. Camtepe, "Privacy preserving distributed machine learning with federated learning," *Computer Communications*, vol. 171, pp. 112–125, Apr. 2021, https://doi.org/10.1016/j.comcom.2021.02.014.

[11] D. Ye, R. Yu, M. Pan, and Z. Han, "Federated Learning in Vehicular Edge Computing: A Selective Model Aggregation Approach," *IEEE Access*, vol. 8, pp. 23920–23935, 2020, https://doi.org/10.1109/ACCESS.2020.2968399.

[12] Y. Zhang, H. Zhang, Y. Yang, W. Sun, H. Zhang, and Y. Fu, "Adaptive differential privacy in asynchronous federated learning for aerial-aided edge computing," *Journal of Network and Computer Applications*, vol. 235, Mar. 2025, Art. no. 104087, https://doi.org/10.1016/j.jnca.2024.104087.

[13] Y. Song, D. He, M. Dai, and M. Guizani, "Cost-Efficient and Privacy-Preserving Distributed Learning: A Double Layer-Based Auction Design," *IEEE Transactions on Mobile Computing*, vol. 24, no. 9, pp. 8824–8840, Sept. 2025, https://doi.org/10.1109/TMC.2025.3560550.

[14] Y. Wan, Y. Qu, L. Gao, and Y. Xiang, "Privacy-preserving blockchain-enabled federated learning for B5G-Driven edge computing," *Computer Networks*, vol. 204, Feb. 2022, Art. no. 108671, https://doi.org/10.1016/j.comnet.2021.108671.

[15] X. Zhao, Y. Wu, T. Zhao, F. Wang, and M. Li, "Federated deep reinforcement learning for task offloading and resource allocation in mobile edge computing-assisted vehicular networks," *Journal of Network and Computer Applications*, vol. 229, Sept. 2024, Art. no. 103941, https://doi.org/10.1016/j.jnca.2024.103941.

[16] Y. Tao *et al.*, "Private Over-the-Air Federated Learning at Band-Limited Edge," *IEEE Transactions on Mobile Computing*, vol. 23, no. 12, pp. 12444–12460, Dec. 2024, https://doi.org/10.1109/TMC.2024.3411295.

[17] Z. Lian, J. Cao, Q. Cao, W. Liu, Z. Zhu, and X. Zhou, "NebulaFL: Self-Organizing Efficient Multilayer Federated Learning Framework With Adaptive Load Tuning in Heterogeneous Edge Systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 11, pp. 3358–3369, Nov. 2024, https://doi.org/10.1109/TCAD.2024.3443715.

[18] M. Chahoud, S. Otoum, and A. Mourad, "On the feasibility of Federated Learning towards on-demand client deployment at the edge," *Information Processing & Management*, vol. 60, no. 1, Jan. 2023, Art. no. 103150, https://doi.org/10.1016/j.ipm.2022.103150.

[19] Q. Yu, H. Xue, C. Wu, Y. Liu, and W. Guo, "Yardstick-Stackelberg pricing-based incentive mechanism for Federated Learning in Edge Computing," *Computer Networks*, vol. 262, May 2025, Art. no. 111186, https://doi.org/10.1016/j.comnet.2025.111186.

[20] Y. Shi *et al.*, "Efficient Federated Learning With Enhanced Privacy via Lottery Ticket Pruning in Edge Computing," *IEEE Transactions on Mobile Computing*, vol. 23, no. 10, pp. 9946–9958, Oct. 2024, https://doi.org/10.1109/TMC.2024.3370967.

[21] S. Caldas *et al.*, "LEAF: A Benchmark for Federated Settings." arXiv, Dec. 09, 2019, https://doi.org/10.48550/arXiv.1812.01097.