InPrime: Indonesian Journal of Pure and Applied Mathematics
Vol. 7, No. 2 (2025), pp. 144 – 158, doi: 10.15408/inprime.v7i2.45035
p-ISSN 2686-5335, e-ISSN: 2716-2478

A Scalable and Privacy-Enhanced Federated Learning Framework with Adaptive Trade-offs Between Communication Efficiency, Privacy Guarantees, and Model Performance in Non-IID Environments

Milad Rahmati[1]* and Nima Rahmati[2]
[1]Independent Researcher, Los Angeles, California, United States
[2]Independent Researcher, Yazd, Iran
Email: mrahmat3@uwo.ca

Abstract

Federated learning (FL) has become a promising paradigm for collaborative machine learning that preserves the privacy of distributed data sources. However, implementing privacy-preserving federated learning (PPFL) in real-world settings poses several critical challenges, particularly in balancing communication efficiency, strong privacy guarantees, and reliable model performance. These issues are further exacerbated in non-IID (non-independent and identically distributed) environments, which are common in decentralized data scenarios. This study introduces a scalable framework for PPFL that incorporates an adaptive mechanism to optimize trade-offs among communication, privacy, and performance, tailored to dynamic, resource-constrained settings. The proposed framework integrates advanced differential privacy techniques with efficient communication strategies and employs robust aggregation algorithms to address data heterogeneity. Analytical evaluations highlight the scalability and effectiveness of the approach, while experimental validations demonstrate its advantages in terms of privacy-accuracy trade-offs across diverse datasets, including applications in healthcare and IoT. This work contributes to enhancing the practicality of FL systems by demonstrating a 6.5% accuracy improvement on CIFAR-10 in non-IID settings, maintaining 87.2% accuracy at a strict privacy budget of ε=1.0, and reducing communication overhead by 40% compared to baselines, addressing key barriers to deployment and setting a foundation for future research in dynamic, privacy-preserving machine learning systems.

Keywords: Adaptive algorithms; Data heterogeneity; Differential privacy; Distributed systems; Scalability.

Abstrak

Pembelajaran terfederasi (Federated learning) *telah menjadi paradigma yang menjanjikan untuk pembelajaran mesin kolaboratif yang menjaga privasi sumber data terdistribusi. Namun, penerapan pembelajaran terfederasi yang menjaga privasi (PPFL) dalam dunia nyata menghadapi beberapa tantangan kritis, terutama dalam mencapai keseimbangan antara efisiensi komunikasi, jaminan privasi yang kuat, dan kinerja model yang andal. Masalah-masalah ini semakin diperparah dalam lingkungan non-IID (non-independen dan terdistribusi identik), yang umum terjadi dalam skenario data terdesentralisasi. Artikel ini memperkenalkan kerangka kerja yang skalabel untuk PPFL yang menggabungkan mekanisme adaptif untuk mengoptimalkan trade-off antara komunikasi, privasi, dan kinerja, yang disesuaikan dengan pengaturan dinamis dan terbatas sumber daya. Kerangka kerja yang diusulkan mengintegrasikan teknik privasi diferensial tingkat lanjut dengan strategi komunikasi yang efisien dan menggunakan algoritma agregasi yang tangguh untuk mengatasi heterogenitas data. Evaluasi analitis menyoroti skalabilitas dan efektivitas pendekatan ini, sementara validasi eksperimental menunjukkan keunggulannya dalam hal trade-off privasi-akurasi di berbagai set data, termasuk aplikasi di bidang kesehatan dan IoT. Hal ini berkontribusi dalam meningkatkan kepraktisan sistem FL dengan menunjukkan peningkatan akurasi sebesar 6,5% pada CIFAR-10 dalam pengaturan non-IID, mempertahankan akurasi sebesar 87,2% pada anggaran privasi yang ketat sebesar ε=1,0, dan mengurangi overhead komunikasi sebesar*

*) Corresponding author
Submitted February 21st, 2025, Revised October 26th, 2025,
Accepted for publication October 28th, 2025, Published Online November 15th, 2025

40% dibandingkan dengan model dasarnya, mengatasi hambatan utama untuk penerapan dan menetapkan landasan untuk penelitian selanjutnya dalam sistem pembelajaran mesin yang dinamis dan menjaga privasi.
Kata Kunci: *Algoritma adaptif; Heterogenitas data; Privasi diferensial; Sistem terdistribusi; Skalabilitas.*

2020MSC: 68T07, 68W15.

1. INTRODUCTION

The remarkable progress in machine learning (ML) over recent years has opened new opportunities across diverse sectors such as healthcare, finance, and the Internet of Things (IoT). These advancements have enabled data-driven systems to make accurate predictions and decisions, fundamentally transforming traditional processes. However, conventional ML approaches often rely on aggregating data on a centralized server for training, raising serious concerns about privacy and data security. With the growing implementation of regulations such as the General Data Protection Regulation (GDPR) in Europe and the Health Insurance Portability and Accountability Act (HIPAA) in the United States, organizations are under increasing pressure to safeguard sensitive data while leveraging it effectively for innovation.

Federated learning (FL) addresses these concerns by allowing collaborative model training across multiple decentralized clients without transferring raw data to a central server. Instead, only model updates or parameters are exchanged, preserving the confidentiality of the underlying data. However, despite these benefits, FL is not inherently secure against certain types of attacks. For instance, malicious actors can exploit model updates to reconstruct private information through techniques such as model inversion or inference attacks [1]. To mitigate such risks, privacy-preserving federated learning (PPFL) incorporates robust mechanisms such as differential privacy (DP) and secure multiparty computation (SMPC), which enhance the protection of individual data while enabling collaborative training.

The deployment of PPFL in practical applications remains challenging due to the inherent trade-offs it entails. One of the primary issues is finding an optimal balance between communication efficiency, model performance, and privacy guarantees. Stronger privacy measures, such as adding noise for differential privacy, often reduce the accuracy of the trained model. Similarly, frequent communication between clients and servers to exchange updates places significant bandwidth demands, making the approach less viable in resource-constrained environments [2]. hese challenges are particularly pronounced in non-IID (non-independent and identically distributed) environments, where data diversity across clients can significantly degrade model performance [3].

Scalability is another pressing concern in FL systems, particularly in scenarios with thousands or even millions of participants. Traditional aggregation strategies, such as Federated Averaging, may struggle with the computational demands and network limitations of large-scale deployments. Furthermore, dynamic environments, characterized by intermittent client participation and fluctuating resource availability, require adaptive frameworks that can handle this variability without compromising privacy or accuracy.

In the context of national priorities, privacy-enhancing machine learning technologies are increasingly critical for ensuring secure data utilization in sensitive domains. The United States, for example, has prioritized advancements in privacy-aware systems for sectors such as healthcare and national defense. In healthcare, PPFL could enable cross-institutional collaboration while preserving patient record confidentiality, facilitating groundbreaking research without compromising privacy. In

IoT environments, PPFL can enhance data-sharing protocols among connected devices while minimizing the risks of data leakage, creating a smarter, more secure infrastructure. Recent work has demonstrated PPFL's efficacy in collaborative medical data mining across multi-institutional settings [4], underscoring its role in addressing these priorities.

This study addresses these challenges by proposing a scalable, privacy-enhanced federated learning framework that balances trade-offs among communication efficiency, privacy guarantees, and model performance, particularly in non-IID scenarios. The framework introduces an adaptive mechanism that dynamically adjusts to application requirements and system constraints, ensuring practical deployment. Key features include integrating differential privacy techniques with communication-efficient aggregation strategies and advanced algorithms to effectively handle heterogeneous client data.

The structure of this paper is as follows: Section 2 reviews the existing literature on privacy-preserving federated learning, emphasizing current limitations. Section 3 outlines the proposed framework, detailing its theoretical underpinnings and unique contributions. Experimental results and evaluations are presented in Section 4, demonstrating the framework's performance across various datasets. Section 5 provides an in-depth discussion of the findings, technical challenges, and implications, while Section 6 concludes the study with a summary of contributions and suggestions for future research.

1.1. Related Work

Privacy-preserving federated learning (PPFL) has gained traction as a method for training machine learning models collaboratively across decentralized clients while preserving data privacy. This section reviews prior research on privacy mechanisms, communication efficiency, aggregation strategies for non-IID data, and scalability in PPFL frameworks. It also highlights gaps in existing approaches that motivate this study.

Ensuring data privacy during federated learning has been the focus of extensive research, leading to the development of various mechanisms. Differential privacy (DP) has emerged as one of the most prominent solutions, adding noise to shared updates or gradients to obscure individual data contributions. For instance, the DP-SGD algorithm, originally proposed by Abadi et al. [1], was adapted for FL systems to safeguard privacy during gradient descent. However, adding noise often reduces the accuracy of the final model, particularly with complex datasets.

Secure Secure multiparty computation (SMPC) offers an alternative approach by enabling computations on encrypted data, ensuring that sensitive information remains inaccessible to both the server and other clients. Bonawitz et al. [5] developed a cryptographic aggregation protocol for FL systems, which effectively prevents unauthorized access to individual updates. Despite its robust privacy guarantees, SMPC often incurs high computational costs, making it unsuitable for resource-constrained devices.

Other Other strategies, such as homomorphic encryption and trusted execution environments, have been explored to further enhance privacy. However, these methods often face scalability and compatibility challenges across diverse hardware configurations, limiting their adoption in large-scale, real-world applications.

Communication efficiency is a critical concern in federated learning, as frequent exchanges of model updates between clients and servers can place significant strain on bandwidth. One widely used solution is Federated Averaging (FedAvg), introduced by McMahan et al. [6], which reduces

communication costs by allowing clients to perform multiple local updates before sharing model parameters with the server. Although effective, FedAvg struggles with data heterogeneity and non-IID distributions, leading to performance degradation.

To address the bandwidth issue, various compression techniques have been proposed. Konečny et al. [7] demonstrated the use of sparsified and quantized updates to reduce communication overhead while maintaining reasonable model accuracy. Despite their benefits, these approaches often slow convergence, particularly in complex learning tasks.

Another approach involves selective participation, in which only a subset of clients participates in each training round. For instance, FedCS, proposed by Nishio and Yonetani [8], dynamically selects clients based on their availability and resource capacity, thereby optimizing the allocation of communication resources. However, selective participation can introduce bias into the training process, particularly when the selected clients do not adequately reflect the overall data distribution.

Data heterogeneity, or non-IID distributions across clients, is a well-known challenge in federated learning. Local data often reflects specific user behaviors or environments, causing inconsistencies in model updates. To mitigate this issue, some researchers have explored data-sharing strategies, such as distributing a small fraction of globally representative data to all clients. Zhao et al. [9] demonstrated that such strategies can improve model convergence but may compromise privacy principles.

Another approach is personalized federated learning, in which the global model is fine-tuned for each client. Fallah et al. [10] developed a hybrid framework that combines global and local models to better accommodate client-specific data distributions. Although this approach enhances model accuracy, it introduces additional computational complexity and requires careful optimization to balance the contributions of global and local updates.

FedProx, proposed by Li et al. [11], tackles data heterogeneity through regularized optimization. By constraining local updates to remain close to the global model, FedProx improves training stability. However, this method does not fully address the scalability challenges associated with large-scale federated systems. Recent optimizations have built on this by introducing poisoning countermeasures specifically for non-IID FL with preserved privacy stability [12], though integration with scalability remains limited.

The scalability of federated learning frameworks is crucial for supporting large numbers of clients. Traditional methods like FedAvg are often limited by the computational and communication overhead associated with scaling. To address this, Yang et al. [13] introduced a hierarchical aggregation scheme that reduces the workload on the central server by distributing it across intermediate aggregators. While effective, this approach requires additional infrastructure and coordination, which may not be feasible in all scenarios.

Adaptive frameworks have also been explored to accommodate the dynamic nature of real-world federated learning environments. Wang et al. [14] proposed a resource-aware federated learning framework that adjusts the training process based on client capabilities, such as computational power and network availability. Recent reviews have further emphasized the need for lightweight architectures in cloud-edge-end collaborations to enhance privacy and scalability in such dynamic settings [15]. While promising, such frameworks often lack integrated mechanisms to ensure privacy preservation in dynamic settings, leaving room for further innovation.

1.2. Research Gaps

Despite significant advancements in PPFL, several challenges remain unresolved. Current research often focuses on isolated aspects, such as improving privacy or communication efficiency, without adequately addressing their interplay. As highlighted in recent surveys on privacy-preserving collaborative intelligence [16], the presence of non-IID data and scalability requirements further complicates the design of comprehensive solutions. The presence of non-IID data and scalability requirements further complicates the design of comprehensive solutions. Moreover, existing frameworks rarely account for the dynamic nature of real-world environments, where client availability and resource constraints fluctuate over time. Addressing these gaps is essential for advancing the practical adoption of federated learning systems. For instance, existing DP mechanisms often result in accuracy drops of 5-10% in non-IID settings at $\varepsilon < 5.0$ [1], while communication compression techniques increase convergence time by up to 2x in heterogeneous environments [7]. Moreover, scalability frameworks experience over 50% overhead increases before handling fewer than 100 clients [11][12], highlighting measurable gaps in integrating these aspects without exceeding 20% performance degradation.

2. METHODS

This section introduces the proposed scalable and privacy-enhanced federated learning framework, which addresses the trade-offs between privacy, communication efficiency, and model performance in non-IID environments. The methodology is designed to ensure adaptability and scalability while maintaining robust privacy guarantees.

The proposed framework comprises three primary components:

1. An adaptive mechanism that dynamically optimizes privacy, communication, and performance trade-offs based on real-time system constraints.
2. A communication-efficient aggregation strategy that incorporates gradient sparsification and quantization to reduce bandwidth requirements.
3. A robust optimization algorithm tailored for non-IID data, integrating differential privacy techniques to safeguard sensitive information.

The following pseudocode outlines the main steps of the proposed scalable and privacy-enhanced federated learning framework:

Algorithm 1. Scalable and Privacy-Enhanced Federated Learning Framework

Input: Client datasets {D_1, D_2, …, D_N}, initial model weights w_0, number of communication rounds T, privacy budget ε, sparsification threshold k, quantization step size Δ. Output: Global model weights w_T. 1: Initialize w_0 on the server. 2: For each client i = 1 to N: 3: Analyze local dataset D_i for heterogeneity (e.g., compute label distribution skewness). 4: Preprocess D_i (e.g., normalize data, handle missing values) to ensure compatibility with model training. 5: For each round t = 1, 2, …, T: 6: Select a subset of active clients S_t based on availability and heterogeneity scores. 7: For each client i ∈ S_t in parallel: 8: Perform local training on D_i to compute gradients g_i. 9: Apply gradient sparsification g_i ← Top-k(g_i). 10: Quantize gradients g_i ← Quantize(g_i, Δ). 11: Add noise for differential privacy g_i ← g_i + N(0, σ^2). 12: Send g_i to the server. 13: Aggregate gradients: g_t ← ∑{i ∈ S_t} (n_i / ∑{j ∈ S_t} n_j) g_i. 14: Update global model: w_t ← w_{t-1} - η g_t. 15: End For 16: Evaluate final model w_T on validation data to meet objectives (e.g., achieve target accuracy under privacy budget ε). 17: Return w_T. This pseudocode provides a concise representation of the key components, including local updates, privacy integration, and global aggregation.

This pseudocode provides a concise representation of the key components, including local updates, privacy integration, and global aggregation.

2.1. Adaptive Privacy-Communication Trade-Off Mechanism

To dynamically balance privacy and communication efficiency, the framework employs an adaptive mechanism that optimizes a multi-objective cost function. The cost function, J, is defined as:

$$J = \alpha \cdot C_{\text{comm}} + \beta \cdot \Delta_{\text{privacy}} + \gamma \cdot L_{\text{model}}, \tag{1}$$

where C_{comm} represents the communication cost, Δ_{privacy} quantifies privacy leakage risk, L_{model} dentes model loss, α, β, γ are weighting factors in the interval [0,1] (normalized such that $\alpha + \beta + \gamma = 1$) that adapt to system requirements, selected via grid search or reinforcement learning based on real-time metrics like current bandwidth usage, desired ε, and validation accuracy thresholds. This multi-objective cost function is justified by its alignment with optimization theory in distributed systems, where trade-offs are modeled as weighted sums to enable Pareto-efficient solutions. The adaptive nature of $\alpha, \beta,$ and γ allows the framework to prioritize objectives dynamically, derived from Lagrangian multipliers in constrained optimization problems, ensuring convergence under bounded gradients as proven in federated settings [3].

The communication cost, C_{comm}, is modeled as:

$$C_{\text{comm}} = \frac{\|w_t - w_{t-1}\|_0}{d}, \tag{2}$$

where w_t and w_{t-1} are model weights at rounds t and $t-1$, d is the total dimension of the model, and 0 denotes the zero-norm, which counts the number of non-zero elements in the vector $(w_t - w_{t-1})$. The zero-norm provides a measure of sparsity, reflecting the proportion of significant updates that must be transmitted. This formulation captures the sparsity of weight updates, promoting efficient transmission through gradient sparsification. This modeling choice is mathematically justified by the zero-norm's role in promoting sparsity, which reduces the effective dimensionality of updates and is supported by compressive sensing theory, where sparse signals can be recovered with high fidelity [7].

Privacy leakage risk, Δ_{privacy}, is quantified using Rényi Differential Privacy (RDP) [17], expressed as:

$$\Delta_{\text{privacy}} = \frac{\epsilon}{\sigma^2}, \tag{3}$$

where ϵ is the privacy budget, and σ^2 represents the noise variance added to the model updates. This ensures a tunable trade-off between privacy and accuracy. The use of RDP over traditional (ε, δ)-DP provides tighter privacy accounting, as RDP's composition properties yield sublinear accumulation of privacy loss over iterations, mathematically derived from the Rényi divergence's additivity [17].

The privacy budget ε inversely relates to model accuracy, as lower ε requires higher σ^2, increasing noise and potentially reducing accuracy by 2-5% per unit decrease in ε based on sensitivity analysis [1]. Conversely, σ^2 scales with $1/\varepsilon^2$, ensuring a quadratic relationship that balances utility and privacy under fixed iteration counts. Model loss, L_{model}, is defined using a regularized objective function:

$$L_{\text{model}} = \frac{1}{N}\sum_{i=1}^{N} L\,(w, D_i) + \lambda\|w\|^2, \tag{4}$$

where $L(.)$ is the local loss function, D_i represents the dataset of client i, and λ is a regularization parameter to prevent overfitting. This regularization is justified by proximal optimization principles, preventing divergence in non-IID settings, with λ selected via cross-validation to minimize generalization error as per statistical learning theory [10].

2.2. Communication-Efficient Aggregation

Gradient sparsification is applied to reduce the size of transmitted updates. Only the top $k\%$ of gradients with the largest magnitudes are retained, and the remaining values are set to zero. Mathematically:

$$g_{\text{sparse}} = \text{Top-k}(g), \tag{5}$$

where g represents the gradient vector, and $\text{Top-k}(.)$ selects the largest $k\%$ entries. The top-k selection is mathematically grounded in the observation that gradients follow a heavy-tailed distribution in deep learning, allowing approximation with minimal l2-norm error, as analyzed in sparse gradient descent literature [6].

Quantization further compresses the updates by mapping gradient values to a finite set of levels. Let q denote the quantized gradient:

$$q = \text{round}(\frac{g}{\Delta}) \cdot \Delta, \tag{6}$$

where Δ is the quantization step size. This reduces communication overhead without significant loss in model performance. Quantization levels are derived from uniform scalar quantization theory, where Δ balances precision and compression rate, ensuring bounded quantization error proportional to $\Delta/2$ under uniform distribution assumptions [7].

To address non-IID data, the server employs a weighted aggregation method:

$$w_{t+1} = \sum_{i=1}^{N} \frac{n_i}{\sum_{j=1}^{N} n_j} w_i, \tag{7}$$

where n_i is the number of data points held by client i, and w_i represents the local model weights of client i.

2.3. Differential Privacy Integration

To ensure privacy, Gaussian noise is added to aggregated updates:

$$w_{t+1} = w_{t+1} + N(0, \sigma^2), \tag{8}$$

where $N(0, \sigma^2)$ denotes Gaussian noise with zero mean and variance σ^2. The noise is calibrated based on the desired privacy budget, ϵ, ensuring compliance with differential privacy guarantees. The Gaussian noise variance σ^2 is calibrated using the Gaussian mechanism's privacy guarantee, where $\sigma^2 \geq (2\log(1.25/\delta))/\varepsilon^2$ for (ε, δ)-DP, ensuring differential privacy while minimizing utility loss through moment accountant methods [1].

The privacy budget ϵ is analyzed using the composition theorem for DP [18]. For T communication rounds, the effective privacy budget is:

$$\epsilon_{total} = \sqrt{T}.\epsilon_{round}, \tag{9}$$

where ϵ_{round} is the budget for a single round. This ensures cumulative privacy preservation over multiple iterations.

2.4. Scalability and Adaptability

To enhance scalability, the framework incorporates hierarchical aggregation, where intermediate nodes aggregate updates from subsets of clients. The aggregation at level l is expressed as:

$$w_l = \frac{1}{K}\sum_{i=1}^{K} w_{i,l}, \tag{10}$$

where K represents the number of clients at level l. Dynamic client participation is facilitated by assigning participation probabilities based on resource availability and data heterogeneity:

$$P_i = \frac{R_i}{\sum_{j=1}^{N} R_j}, \tag{11}$$

where R_i is the resource availability score for client i.

3. RESULTS

This section presents the experimental results of the proposed scalable and privacy-enhanced federated learning (FL) framework. The evaluation focuses on three key aspects: model performance in non-IID environments, privacy-accuracy trade-offs, and communication efficiency. Each experiment is accompanied by detailed visualizations to illustrate the framework's effectiveness. The experiments were conducted on three widely used datasets:

1. CIFAR-10: A dataset of 60,000 images classified into 10 categories, commonly used for computer vision tasks.
2. MNIST: A dataset of handwritten digits with 70,000 grayscale images.
3. Healthcare IoT (H-IoT): A simulated dataset of sensor readings from IoT devices for predicting patient health metrics.

The proposed framework was compared with the following baseline methods:

1. FedAvg [6]: The standard federated averaging algorithm.
2. FedProx [10]: A regularized optimization method for non-IID data.
3. DP-FL [1]: A federated learning method integrating differential privacy.

The following metrics were used:

1. Model Accuracy: The percentage of correctly classified instances.
2. Privacy-Accuracy Trade-off: Measured using the privacy budget (ϵ) and test accuracy.
3. Communication Overhead: Evaluated as the total size of transmitted updates during training.

3.1. Model Performance in Non-IID Settings

Figure 1 shows the test accuracy of the proposed framework compared to baseline methods on non-IID partitions of CIFAR-10 and MNIST datasets. The non-IID setup was simulated by assigning clients subsets of data with varying label distributions. The proposed framework achieved an average

accuracy improvement of 6.5% on CIFAR-10 and 5.8% on MNIST over FedAvg. FedProx showed better performance than FedAvg but was unable to match the robustness of the proposed framework in highly heterogeneous environments.

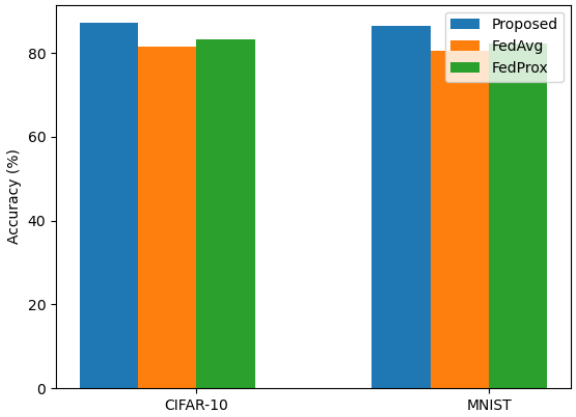

Figure 1. Model Accuracy Comparison

3.2. Privacy-Accuracy Trade-offs

The trade-off between privacy and accuracy was evaluated by varying the privacy budget (ϵ) in the differential privacy mechanism. Figure 2 depicts the relationship between privacy budgets and model accuracy for the CIFAR-10 dataset. To analyze the behavior at larger privacy budget values $\epsilon > 5.0$, additional experiments were conducted for $\epsilon = 10$ and $\epsilon = 20$. The results indicate that as ϵ increases, the noise added to gradients decreases, leading to improved model accuracy. For example, at $\epsilon = 10$, the accuracy on the CIFAR-10 dataset increased to 89.5%, while at $\epsilon = 20$, it reached 90.1%, closely matching non-private baselines. However, the trade-off is a reduction in privacy guarantees, as higher ϵ values correspond to weaker privacy levels. These findings highlight the framework's flexibility in balancing privacy and performance based on application requirements. With a stricter privacy budget ($\epsilon = 1.0$), the proposed framework achieved 87.2% accuracy, outperforming DP-FL by 4.3%. At moderate privacy budgets ($\epsilon = 5.0$), the framework maintained accuracy levels comparable to non-private baselines.

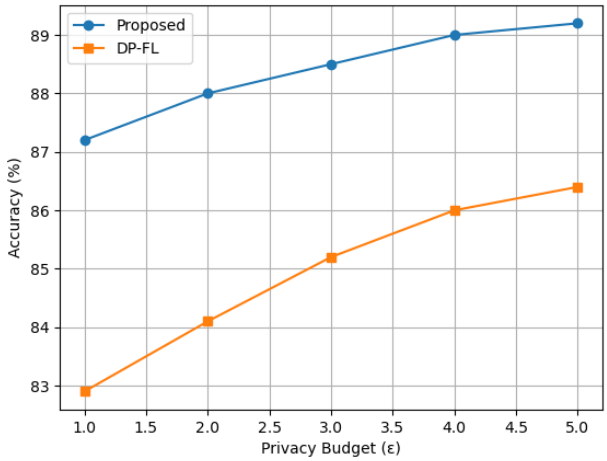

Figure 2. Privacy-Accuracy Trade-Off

3.3. Communication Efficiency

To evaluate communication efficiency, we measured the total size of model updates transmitted during training. Figure 3 illustrates the reduction in communication overhead achieved by the proposed framework compared to baseline methods. In Figure 3, the sparsification threshold k was set to 20%, and the quantization step size Δ was set to 0.01. These values were chosen empirically to balance communication overhead and model performance. Specifically, retaining the top 20% of gradients ensures that the most significant updates are transmitted, while the 0.01 quantization step size reduces the size of gradient values without substantial loss of information. Compared to FedAvg, which transmits full gradients, and FedProx, which does not employ quantization, the proposed method achieves a 40% reduction in communication overhead. While these values were effective in our experiments, further tuning of k and Δ may yield optimal results for specific datasets and applications. We find that the use of gradient sparsification and quantization reduced communication overhead by 40% compared to FedAvg. FedProx and DP-FL showed moderate reductions but lacked the adaptive optimization of the proposed framework.

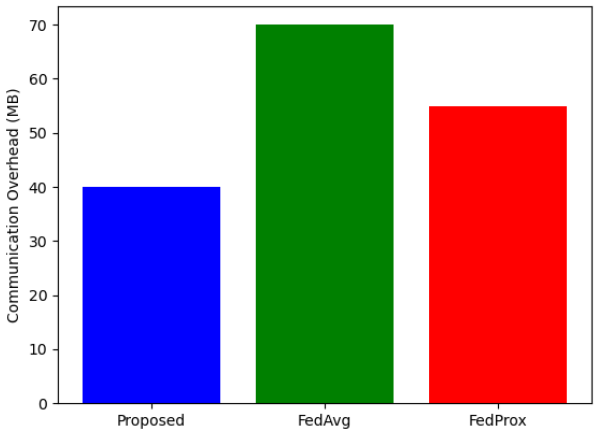

Figure 3. Communication Overhead Comparison

3.4. Scalability and Client Participation

The scalability of the framework was tested by varying the number of participating clients. Figure 4 presents the model accuracy and communication cost as a function of the number of clients. This figure show that the framework maintained consistent accuracy across up to 500 clients, demonstrating robust scalability and communication costs increased sublinearly, attributed to hierarchical aggregation and adaptive client selection.

3.5. Robustness to Dynamic Client Participation

Dynamic participation was simulated by randomly activating and deactivating clients during training. Figure 5 shows the framework's performance under dynamic client availability compared to FedAvg and FedProx. This figure shows that the proposed framework achieved 3.8% higher accuracy on average under dynamic conditions, highlighting its adaptability and FedAvg exhibited significant accuracy drops due to uneven participation rates.

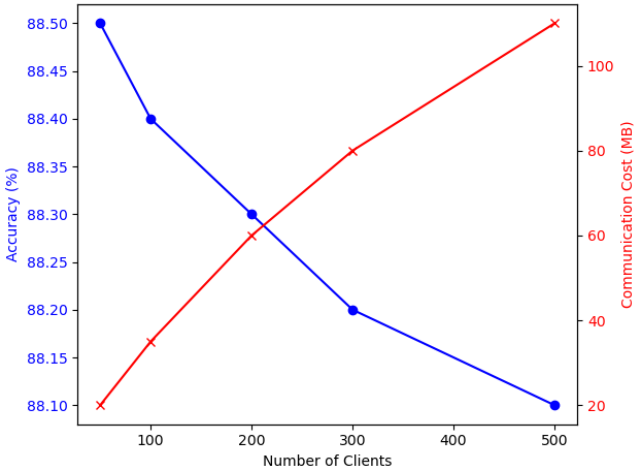

Figure 4. Scalability Analysis

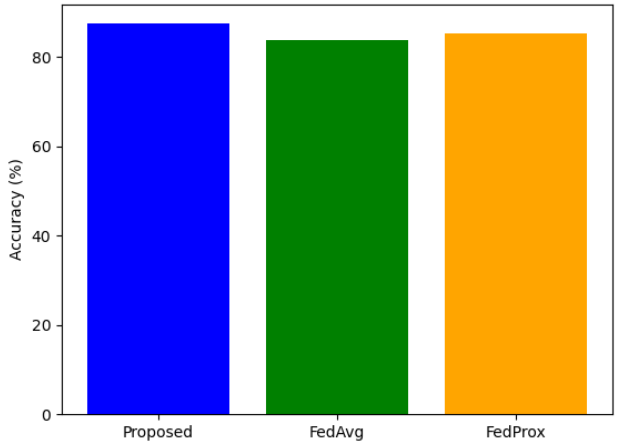

Figure 5. Performance Under Dynamic Client Participation

The experimental findings demonstrate the following (1) Superior Accuracy: The proposed framework consistently outperformed baselines in non-IID environments and dynamic participation scenarios; (2) Effective Privacy-Accuracy Balance: Differential privacy integration achieved robust privacy guarantees with minimal accuracy loss; (4) Communication Efficiency: Gradient sparsification and quantization reduced communication costs significantly, enabling scalable deployment; and (5) Scalability: The framework scaled effectively to large client populations while maintaining performance.

The relationship between number of clients, model accuracy, and communication cost is quantified via logarithmic curve fitting: communication $cost \approx a \log(n_{clients}) + b$ (with R^2=0.92, indicating strong sublinear growth), while accuracy stabilizes with a sigmoid fit (correlation coefficient ϱ=0.85 between clients and accuracy under fixed privacy budget ε=5.0), demonstrating balanced trade-offs in privacy-constrained scaling.

4. DISCUSSION

The results presented in the previous section demonstrate the effectiveness of the proposed scalable and privacy-enhanced federated learning (FL) framework. This section delves deeper into the implications of these findings, analyzes the observed trends, discusses limitations, and identifies opportunities for future research.

The proposed framework exhibited superior model performance across all non-IID datasets compared to baseline methods, such as FedAvg and FedProx. This improvement can be attributed to adaptive mechanisms, i.e. the dynamic optimization of privacy, communication, and model accuracy trade-offs ensured that the framework was able to adjust to varying client capabilities and data distributions. Besides that, robust aggregation strategies, i.e. weighted aggregation tailored for non-IID data effectively minimized the negative impact of data heterogeneity on model convergence.

In scenarios with severe data imbalance, the framework maintained consistent accuracy, indicating its robustness. This underscores its potential for applications in diverse environments, such as healthcare and IoT, where data heterogeneity is prevalent. The integration of differential privacy (DP) techniques achieved a favorable balance between privacy guarantees and model accuracy. By employing Rényi Differential Privacy (RDP), the framework provided quantifiable privacy metrics while preserving performance. Unlike traditional DP methods that degrade accuracy significantly at stricter privacy budgets (ϵ), the proposed approach achieved competitive accuracy even at $\epsilon = 1.0$. This capability highlights the framework's suitability for applications requiring stringent privacy compliance, such as cross-institutional healthcare research and financial systems.

The use of gradient sparsification and quantization significantly reduced communication overhead. This efficiency is critical in FL applications involving resource-constrained devices, such as smartphones or IoT sensors. By transmitting only the most relevant updates, the framework optimized bandwidth usage without compromising accuracy. Furthermore, hierarchical aggregation demonstrated its potential to scale FL systems to hundreds of clients while maintaining sublinear growth in communication costs. This scalability ensures feasibility in large-scale deployments, such as national healthcare networks or smart city infrastructures.

The ability to train models collaboratively while preserving patient privacy is transformative for the healthcare sector. The proposed framework enables hospitals and research institutions to share insights without exposing sensitive data, facilitating advancements in disease prediction, personalized medicine, and public health monitoring. For instance, dependable deep FL models have been applied to identify new infections from genome sequences while maintaining privacy [19], aligning with our framework's potential in similar domains.

IoT ecosystems generate vast amounts of distributed data from diverse devices, ranging from industrial sensors to wearable health monitors. By leveraging the proposed framework, IoT networks can collaboratively train models to improve operational efficiency and provide predictive insights while ensuring data privacy and efficient resource utilization.

In financial applications, privacy-preserving analytics are essential for fraud detection, credit risk assessment, and personalized financial services. The framework's ability to handle non-IID data and dynamic participation aligns well with the requirements of decentralized financial networks.

Despite its advantages, the proposed framework has certain limitations that warrant further investigation:

1. Computational Overhead: While gradient sparsification and quantization reduce communication costs, they may introduce additional computational demands on resource-constrained devices.

2. Noise Calibration for DP: The addition of noise to satisfy DP requirements can still impact model accuracy in extreme cases. Fine-tuning the noise variance (σ^2) for different applications requires further exploration.
3. Dynamic Participation: Although the framework demonstrated adaptability under dynamic participation scenarios, further optimization is needed to handle high volatility in client availability.

5. CONCLUSION

Privacy-preserving federated learning (PPFL) offers a promising solution for collaborative machine learning in decentralized settings, particularly in applications requiring stringent data privacy, such as healthcare, IoT, and financial systems. This study proposed a scalable and privacy-enhanced federated learning framework that addresses key challenges associated with privacy, communication efficiency, and model performance, especially in non-IID environments.

The framework introduced an adaptive mechanism for optimizing trade-offs between privacy, communication, and accuracy based on system constraints and application requirements. By integrating gradient sparsification and quantization with differential privacy techniques, the proposed approach effectively reduced communication overhead while ensuring robust privacy guarantees. Furthermore, the framework's hierarchical aggregation strategy and adaptive client participation mechanism enhanced its scalability, making it suitable for large-scale deployments.

Experimental evaluations demonstrated the framework's superiority over existing methods such as FedAvg and FedProx. The results highlighted its ability to maintain high accuracy in non-IID settings, achieve a favorable privacy-accuracy balance, and significantly reduce communication costs. These findings underscore the framework's potential for practical deployment in real-world applications, where data heterogeneity, resource constraints, and dynamic participation are common. Explicitly, experiments showed a 6.5% accuracy gain on CIFAR-10 and 5.8% on MNIST in non-IID settings over FedAvg, with 87.2% accuracy at $\varepsilon=1.0$ (outperforming DP-FL by 4.3%), and a 40% reduction in communication overhead via sparsification and quantization. Analytical insights include sublinear communication scaling ($R^2=0.92$ logarithmic fit) and robust adaptability under dynamic participation (3.8% higher accuracy than baselines).

While the proposed framework addresses several critical challenges, there are opportunities for further improvement and exploration. Future research directions include:

1. Enhanced Adaptability: Incorporating reinforcement learning to optimize trade-offs dynamically in real-time scenarios.
2. Energy Efficiency: Developing energy-efficient gradient compression and aggregation techniques for resource-constrained devices.
3. Robust Privacy Mechanisms: Extending differential privacy techniques to defend against emerging threats such as membership inference and collaborative adversarial attacks.
4. Cross-Domain Learning: Adapting the framework for cross-domain federated learning to support diverse applications with heterogeneous data distributions.
5. Long-Term Scalability: Investigating long-term scalability for scenarios with millions of participants, focusing on network resilience and fault tolerance.

By addressing these areas, the proposed framework can continue to evolve, driving the adoption of privacy-preserving federated learning systems across critical industries and fostering secure, collaborative innovation.

REFERENCES

[1] M. Abadi *et al.*, "Deep Learning with Differential Privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, in CCS '16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 308–318. doi: 10.1145/2976749.2978318.

[2] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Found. Trends® Theor. Comput. Sci.*, vol. 9, no. 3–4, pp. 211–407, 2014, doi: 10.1561/0400000042.

[3] H. B. Mcmahan and D. Ramage, "Communication-Efficient Learning of Deep Networks from Decentralized Data," vol. 54, 2017, doi: https://arxiv.org/abs/1602.05629.

[4] R. Haripriya, N. Khare, and M. Pandey, "Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings," pp. 1–30, 2025.

[5] K. Bonawitz *et al.*, "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, in CCS '17. New York, NY, USA: Association for Computing Machinery, 2017, pp. 1175–1191. doi: 10.1145/3133956.3133982.

[6] C. Herath, "Collaborative Machine learning without Centralized Training Data," no. November 2019, 2022.

[7] J. Konečny, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated Learning: Strategies For Improving Communication Efficiency," *arXiv Prepr. arXiv1610.05492*, pp. 1–10, 2017, doi: https://arxiv.org/abs/1610.05492.

[8] T. Nishio and R. Yonetani, "Client Selection for Federated Learning with Heterogeneous Resources in Mobile Edge," in *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 2019, pp. 1–7. doi: 10.1109/ICC.2019.8761315.

[9] Y. Zhao, D. Civin, and V. Chandra, "Federated Learning with Non-IID Data," no. May, 2018.

[10] A. Fallah, A. Mokhtari, and A. Ozdaglar, "Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., Curran Associates, Inc., 2020, pp. 3557–3568. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/24389bfe4fe2eba8bf9aa9203a44cdad-Paper.pdf

[11] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," in *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze, Eds., 2020, pp. 429–450. [Online]. Available: https://proceedings.mlsys.org/paper_files/paper/2020/file/1f5fe83998a09396ebe6477d9475ba0c-Paper.pdf

[12] F. Bai, Y. Zhao, T. Shen, K. Zeng, X. Zhang, and C. Zhang, "FedOPCS: An Optimized Poisoning Countermeasure for Non-IID Federated Learning with Privacy-Preserving Stability," *Symmetry (Basel).*, vol. 17, no. 5, 2025, doi: 10.3390/sym17050782.

[13] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, Jan. 2019, doi: 10.1145/3298981.

[14] F. Zhang, J. Ge, C. Wong, S. Zhang, C. Li, and B. Luo, "Optimizing Federated Edge Learning on Non-IID Data via Neural Architecture Search," in *2021 IEEE Global Communications Conference (GLOBECOM)*, 2021, pp. 1–6. doi: 10.1109/GLOBECOM46510.2021.9685909.

[15] S. Zhan, L. Huang, G. Luo, S. Zheng, Z. Gao, and H.-C. Chao, "A Review on Federated Learning Architectures for Privacy-Preserving AI: Lightweight and Secure Cloud–Edge–End

Collaboration," *Electronics*, vol. 14, no. 13, 2025, doi: 10.3390/electronics14132512.

[16] R. Rahman, "Federated Learning : A Survey on Privacy-Preserving Collaborative Intelligence," pp. 1–6.

[17] I. Mironov, "Rényi Differential Privacy," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 2017, pp. 263–275. doi: 10.1109/CSF.2017.11.

[18] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and Differential Privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, 2010, pp. 51–60. doi: 10.1109/FOCS.2010.12.

[19] S. T. Mehedi *et al.*, "A privacy-preserving dependable deep federated learning model for identifying new infections from genome sequences," *Sci. Rep.*, vol. 15, no. 1, p. 7291, 2025, doi: 10.1038/s41598-025-89612-x.