

FedSecure: A Privacy-Preserving Federated Learning Framework for Heterogeneous Edge Networks

James R. Mitchell*, Sarah K. Park[†], and David L. Fernandez*

*Department of Electrical and Computer Engineering

Georgia Institute of Technology, Atlanta, GA 30332

Email: {jrmitchell, dfernandez}@gatech.edu

[†]Microsoft Research

Redmond, WA 98052

Email: sarahpark@microsoft.com

Abstract—Federated learning enables collaborative model training across distributed edge devices without centralizing raw data. However, existing frameworks face critical challenges in handling statistical heterogeneity across clients and ensuring robust privacy guarantees. In this paper, we present FedSecure, a novel federated learning framework that integrates differential privacy with adaptive aggregation strategies designed for heterogeneous edge networks. Our approach introduces a client-aware noise calibration mechanism that adjusts privacy budgets based on local data characteristics, achieving a favorable privacy-utility trade-off. We propose a hierarchical aggregation protocol that groups clients by data distribution similarity, reducing the impact of non-IID data on global model convergence. Extensive experiments on four benchmark datasets across diverse heterogeneity scenarios demonstrate that FedSecure improves accuracy by 8.3% over FedAvg while maintaining (ϵ, δ) -differential privacy with $\epsilon = 1.0$. Our framework also reduces communication overhead by 40% through gradient compression and selective participation.

Index Terms—Federated learning, differential privacy, edge computing, heterogeneous networks, distributed machine learning

I. INTRODUCTION

The proliferation of edge devices has created unprecedented opportunities for distributed machine learning. Mobile phones, IoT sensors, and autonomous vehicles generate vast amounts of data that can be leveraged to train powerful predictive models. However, transmitting this data to a central server raises significant privacy concerns and may violate data protection regulations such as GDPR and CCPA [2].

Federated learning (FL) has emerged as a promising paradigm to address these challenges by enabling model training directly on edge devices [1]. In the canonical federated learning setup, a central server coordinates the training process by distributing model parameters to participating clients, which perform local training on their private data and return model updates to the server for aggregation.

Despite significant progress, two fundamental challenges remain:

- 1) **Statistical heterogeneity**: Data distributions across clients are typically non-IID, leading to client drift and degraded global model performance [3].

- 2) **Privacy guarantees**: While federated learning avoids direct data sharing, model updates can still leak sensitive information through inference attacks [4].

In this paper, we propose FedSecure, a comprehensive framework that jointly addresses these challenges. Our key contributions include:

- A client-aware differential privacy mechanism that adaptively calibrates noise injection based on local data characteristics and sensitivity.
- A hierarchical aggregation protocol that clusters clients by distribution similarity to mitigate non-IID effects.
- A gradient compression scheme that reduces communication costs while preserving model quality under privacy constraints.
- Comprehensive evaluation on CIFAR-10, CIFAR-100, FEMNIST, and Shakespeare datasets demonstrating state-of-the-art performance.

II. RELATED WORK

A. Federated Learning

McMahan et al. [1] introduced FedAvg, the foundational algorithm for federated learning. Subsequent work has proposed improvements for handling non-IID data, including FedProx [5], SCAFFOLD [3], and FedNova [6]. These methods introduce regularization terms or variance reduction techniques to improve convergence under heterogeneous data distributions.

B. Differential Privacy in FL

Differential privacy (DP) provides formal privacy guarantees by adding calibrated noise to computations [7]. Several works have integrated DP into federated learning, including user-level DP [8] and record-level DP approaches. The primary challenge lies in balancing the privacy budget with model utility, particularly in heterogeneous settings.

III. PROPOSED FRAMEWORK

A. System Model

We consider a federated learning system with N clients $\{C_1, C_2, \dots, C_N\}$ coordinated by a central server S . Each client C_i holds a local dataset $\mathcal{D}_i$ of size n_i , drawn from a potentially unique distribution P_i . The global objective is:

Algorithm 1 FedSecure Training Protocol

```

1: Input:  $N$  clients, rounds  $T$ , local epochs  $E$ , privacy budget  $\epsilon$ 
2: Initialize global model  $w_0$ 
3: for  $t = 0$  to  $T - 1$  do
4:   Server selects subset  $S_t \subseteq \{1, \dots, N\}$ 
5:   Server broadcasts  $w_t$  to selected clients
6:   for each client  $i \in S_t$  in parallel do
7:      $g_i \leftarrow \text{LocalTrain}(w_t, \mathcal{D}_i, E)$ 
8:      $\hat{g}_i \leftarrow \text{Clip}(g_i, C) + \mathcal{N}(0, \sigma_i^2 I)$ 
9:     Send  $\hat{g}_i$  to server
10:  end for
11:  Server clusters clients into groups  $\{G_k\}$ 
12:   $w_{t+1} \leftarrow \text{HierarchicalAggregate}(\{\hat{g}_i\}, \{G_k\})$ 
13: end for

```

$$\min_w F(w) = \sum_{i=1}^N \frac{n_i}{n} F_i(w) \quad (1)$$

where $F_i(w) = \mathbb{E}_{\xi \sim P_i}[f(w; \xi)]$ is the local objective and $n = \sum_{i=1}^N n_i$.

B. Client-Aware Noise Calibration

Standard DP mechanisms apply uniform noise across all clients, which is suboptimal when clients have varying data characteristics. We propose an adaptive noise calibration strategy that adjusts the noise scale σ_i for each client:

$$\sigma_i = \frac{\Delta_i \cdot \sqrt{2 \ln(1.25/\delta)}}{\epsilon_i} \quad (2)$$

where Δ_i is the sensitivity of client i 's gradient update and ϵ_i is the allocated privacy budget. We distribute the total privacy budget ϵ across clients using an optimization-based allocation that minimizes the expected aggregation error.

C. Hierarchical Aggregation

To address non-IID data, we introduce a two-level aggregation scheme. First, we compute the pairwise distribution distance between clients using the Wasserstein distance estimated from model updates. Clients are then clustered into K groups $\{G_1, \dots, G_K\}$ using spectral clustering. Aggregation proceeds in two phases: intra-group aggregation followed by inter-group aggregation.

IV. EXPERIMENTAL EVALUATION**A. Setup**

We evaluate FedSecure on four datasets: CIFAR-10, CIFAR-100, FEMNIST (62-class character recognition), and Shakespeare (next-character prediction). We simulate non-IID distributions using a Dirichlet allocation with concentration parameter $\alpha \in \{0.1, 0.5, 1.0\}$ across 100 clients. We use ResNet-18 for CIFAR experiments and an LSTM model for Shakespeare.

TABLE I

TEST ACCURACY (%) COMPARISON UNDER NON-IID SETTING ($\alpha = 0.5$) WITH $\epsilon = 1.0$ DIFFERENTIAL PRIVACY.

Method	CIFAR-10	CIFAR-100	FEMNIST	Shakespeare
FedAvg + DP	61.2	33.8	72.4	45.1
FedProx + DP	63.5	35.2	74.1	46.8
SCAFFOLD + DP	65.1	37.4	75.3	47.9
DP-FedAvg	59.8	32.1	70.9	43.6
FedSecure	69.5	41.7	78.6	51.2

B. Main Results

Table I presents the comparison of FedSecure with baseline methods under non-IID settings ($\alpha = 0.5$) with privacy guarantee $\epsilon = 1.0$.

C. Communication Efficiency

FedSecure achieves a 40% reduction in communication overhead compared to FedAvg through the combination of gradient compression and selective client participation. To reach 65% accuracy on CIFAR-10 with $\epsilon = 1.0$, FedSecure requires 120 communication rounds compared to 280 for FedAvg + DP and 195 for SCAFFOLD + DP.

D. Privacy-Utility Trade-off

We analyze the privacy-utility trade-off by varying ϵ from 0.1 to 10.0. FedSecure consistently outperforms baselines across all privacy levels. At stringent privacy ($\epsilon = 0.1$), FedSecure achieves 52.3% accuracy on CIFAR-10 compared to 41.7% for the best baseline, demonstrating the effectiveness of client-aware noise calibration.

V. CONCLUSION

We presented FedSecure, a privacy-preserving federated learning framework that addresses the dual challenges of statistical heterogeneity and privacy protection in edge networks. Through client-aware noise calibration and hierarchical aggregation, FedSecure achieves state-of-the-art performance across diverse datasets while providing formal differential privacy guarantees. Future work will explore extending FedSecure to cross-silo settings and incorporating secure aggregation protocols for enhanced protection against server-side attacks.

ACKNOWLEDGMENT

This work was supported by NSF grants CNS-2143868 and CNS-2106891, and by a gift from Microsoft Research.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agueria y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, pp. 1273–1282, 2017.
- [2] P. Voigt and A. Von dem Bussche, "The EU general data protection regulation (GDPR)," *Springer International Publishing*, vol. 10, pp. 10–5555, 2017.
- [3] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "SCAFFOLD: Stochastic controlled averaging for federated learning," in *Proc. ICML*, pp. 5132–5143, 2020.
- [4] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning," in *Proc. IEEE S&P*, pp. 739–753, 2019.

- [5] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. MLSys*, 2020.
- [6] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," in *Proc. NeurIPS*, vol. 33, pp. 7611–7623, 2020.
- [7] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [8] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, "Learning differentially private recurrent language models," in *Proc. ICLR*, 2018.