SSR Journal of Engineering and Technology (SSRJET)

Homepage: https://ssrpublisher.com/ssrjet/

Volume 3, Issue 4, 2026 ISSN: 3049-0383

Design and Implementation of Adaptive Personalized Differential Privacy in Federated Learning

Prof. Akomolafe D.T. and Giwa Jesuloluwa Emmanuel

Olusegun Agagu University of Science and Technology, Okitipupa, Ondo State, Nigeria

Received: *01.01.2026* | ***Accepted:*** *06.01.2026* | ***Published:*** *24.04.2026*

***Corresponding author:** Giwa Jesuloluwa Emmanuel

DOI: *10.5281/zenodo.19734376*

Abstract **Original Research Article**

Federated Learning enables collaborative model training across distributed clients while keeping raw data local, making it attractive for privacy sensitive domains such as healthcare, finance, and edge intelligence. This paper proposes the Design and Implementation of Adaptive Personalized Differential Privacy framework for federated learning that dynamically allocates client specific privacy budgets based on quantified contribution scores. The proposed approach is implemented using PyTorch and Opacus and evaluated on benchmark datasets under non independent and identically distributed data settings. Experimental results demonstrate that the adaptive personalised strategy achieves higher global model accuracy, improved fairness across clients, and a more efficient privacy utility balance compared to uniform differential privacy baselines. Overall, the study confirms that adaptive personalised differential privacy significantly enhances the practicality and robustness of federated learning systems.

Keywords: Federated Learning, Differential Privacy, Personalized Privacy, Adaptive Privacy Budget.

1.0 Introduction

Federated Learning is now seen as a promising paradigm for privacy preserving machine learning because it allows collaborative model training across multiple clients without sharing raw data. Sensitive data are stored on local devices therefore, the risks associated with centralised data collection and storage is minimized. This paradigm is particularly relevant in domains where regulatory, ethical, or commercial constraints prohibit direct data sharing.

Although, it is a decentralised design, federated learning is not completely immune to privacy leakage. Previous studies have shown that gradients and model updates can be exploited through inference attacks to reconstruct and manipulate sensitive information about client data. To eliminate these risks, Differential Privacy has been integrated into federated learning systems to provide mathematically provable privacy guarantees by populating model updates with carefully calibrated noise.

Most existing differentially private federated learning frameworks adopt a uniform privacy budget for all participating clients. This assumption ignores real world heterogeneity in data distribution, client contribution, and privacy sensitivity. Clients with limited or highly sensitive data may require stronger privacy protection, while clients with larger or more informative datasets could tolerate weaker privacy constraints to improve overall model

performance. Uniform privacy mechanisms therefore lead to inefficient privacy utility tradeoffs and can exacerbate fairness issues among heterogeneous clients.

This paper addresses these limitations by proposing an adaptive personalised differential privacy framework that dynamically assigns privacy budgets to clients based on their quantified contribution to the global model. By aligning privacy allocation with client influence, the proposed approach improves model utility, fairness, and privacy efficiency.

2.0 Related Work

Federated Learning was formally introduced by McMahan *et al.* (2017) as a communication efficient framework for training machine learning models across decentralised data sources. Since then, extensive research has investigated its scalability, convergence behaviour, and applicability to privacy sensitive domains such as mobile analytics, healthcare, and finance. Despite its decentralised nature, subsequent studies demonstrated that federated learning is vulnerable to privacy leakage through shared gradients and model updates, motivating the integration of stronger privacy preserving mechanisms.

Differential Privacy provides a mathematically rigorous framework for limiting information leakage from statistical computations. Abadi *et al.* (2016) introduced differentially private stochastic gradient descent, establishing a foundation for privacy preserving deep learning. Building on this work, client level differential privacy was proposed for federated learning to protect the participation of entire clients rather than individual data records. Geyer *et al.* (2017) and McMahan *et al.* (2018) demonstrated that adding calibrated noise to client updates can provide formal privacy guarantees in federated settings, albeit at the cost of reduced model accuracy.

Rényi Differential Privacy were proposed to tighten privacy bounds and improve utility. Mironov (2017) showed that Rényi based accounting enables more accurate tracking of cumulative privacy loss compared to classical composition theorems.

Recent works have begun to explore contribution aware mechanisms in federated learning, including data valuation, gradient based influence estimation, and Shapley value approximations. While these methods aim to quantify client importance, they are rarely integrated with formal differential privacy mechanisms. Existing differentially private federated learning frameworks largely ignore client contribution heterogeneity, resulting in inefficient privacy budget utilisation.

This work differentiates itself from prior research by unifying personalised differential privacy with contribution aware federated learning. Unlike existing approaches that apply uniform privacy budgets or ad hoc personalisation, the proposed framework introduces a principled, adaptive privacy allocation strategy grounded in gradient norm based contribution scoring and formal privacy accounting. This integration addresses both privacy and fairness concerns, advancing the state of the art in privacy preserving federated learning.

3.0 Methodology

This study adopts a design science research approach to develop, implement, and evaluate an adaptive personalized differential privacy framework within a federated learning (FL) environment. The methodology is structured into five key phases: system design, client contribution assessment, privacy budget allocation, model training, and evaluation.

3.1 System Architecture

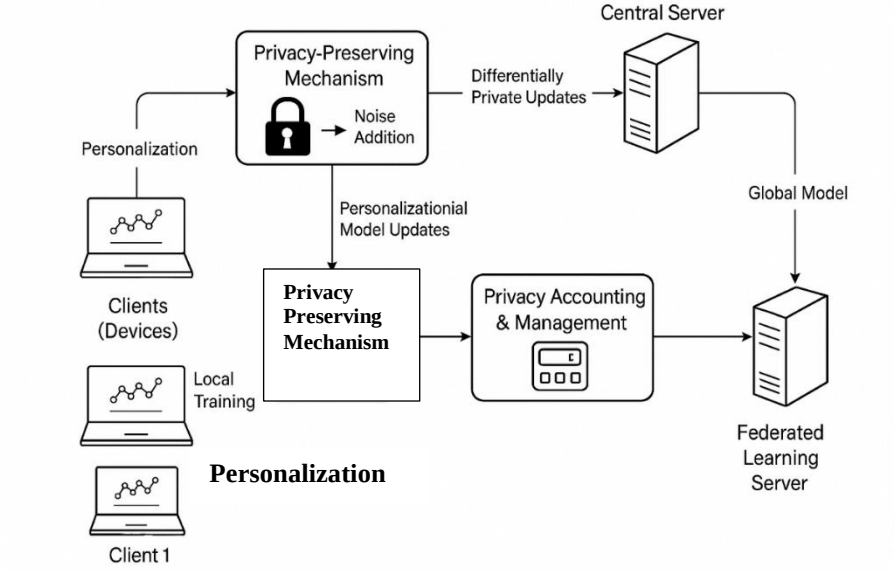

Fig 1: Architectural Design of the Model

3.2 Components of the architecture

3.2.1 Clients (Devices): These are decentralized data holders such as smartphones, laptops, or edge devices. Each client performs local training (training on its own dataset without sharing raw data) and personalization (tailoring the model to the unique distribution of the local data (non-IID)). The system supports client-specific model customization while still contributing to a global model.

3.2.2 Privacy-Preserving Mechanism: This module appears twice in the architecture, reflecting dual roles. During Local Training, it ensures that noisy updates are generated before being transmitted, implementing differential privacy through noise addition and, before global aggregation: It applies an additional layer of protection to prevent sensitive leakage, particularly from personalized model updates.

This helps limit the possibility of inference attacks on individual client data.

3.2.3 Noise Addition: This is a central component of Differential Privacy (DP). It ensures (ε, δ)-DP guarantees by perturbing model updates before sharing, protecting against adversaries attempting to reconstruct original client data from gradients.

Mathematically:

$$\tilde{w}_i = w_i + \mathcal{N}(0, \sigma^2 I)$$

Where ω̃ is the privatized weight update sent to the server.

3.2.4 Differentially Private Updates: These are the noised model parameters or gradients shared by clients with the Central Server. They contain enough information for global model learning while masking specific data contributions.

3.2.5 Central Server: It collects DP updates from clients, sends global model back after

aggregation via FedAvg. but, cannot reconstruct individual data due to DP. This ensures central orchestration without central data collection, a core tenet of FL.

3.2.6 Federated Learning Server: It performs global aggregation and model updates. It may include modules for client sampling, global optimization and robust aggregation (e.g., against poisoned clients). It acts as the final model distributor after integrating all privacy-preserving client contributions.

3.2.7 Privacy Accounting & Management: It is responsible for **t**racking the cummulative privacy loss (ε) per client across rounds. It also ensures that no client exceeds their allocated privacy budget and also helps regulate adaptive noise levels based on personalized risk thresholds.

3.2.8 Global Model Loopback: At this point, the updated global model is sent back to all clients for further training (next round) and personalization (local fine-tuning). This loop allows continual learning while adapting to each user's context.

4.0 Implementation and Result

Implementation was performed using PyTorch for model definition, Opacus for DP enforcement, NumPy, Pandas, Matplotlib for analytics and visualization and LEAF for data distribution simulation.

The performance of the proposed Adaptive Personalized Differential Privacy Federated Learning framework is evaluated against standard Federated Averaging without differential privacy and differentially private federated learning with uniform privacy budgets. Experimental results demonstrate that the proposed approach achieves superior global model accuracy under equivalent or stronger privacy constraints. The following graphs were generated:

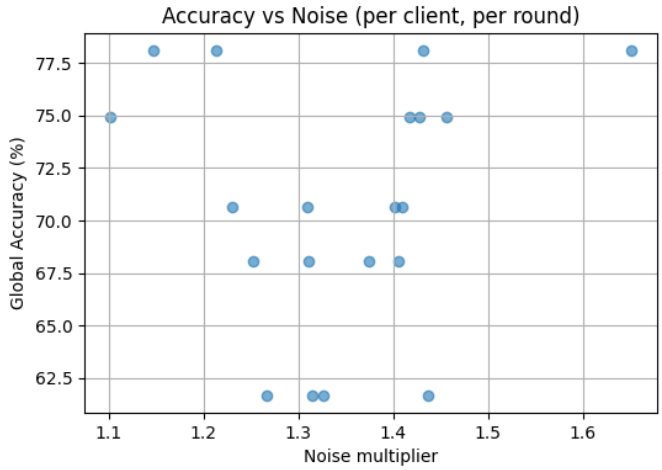

Figure 2

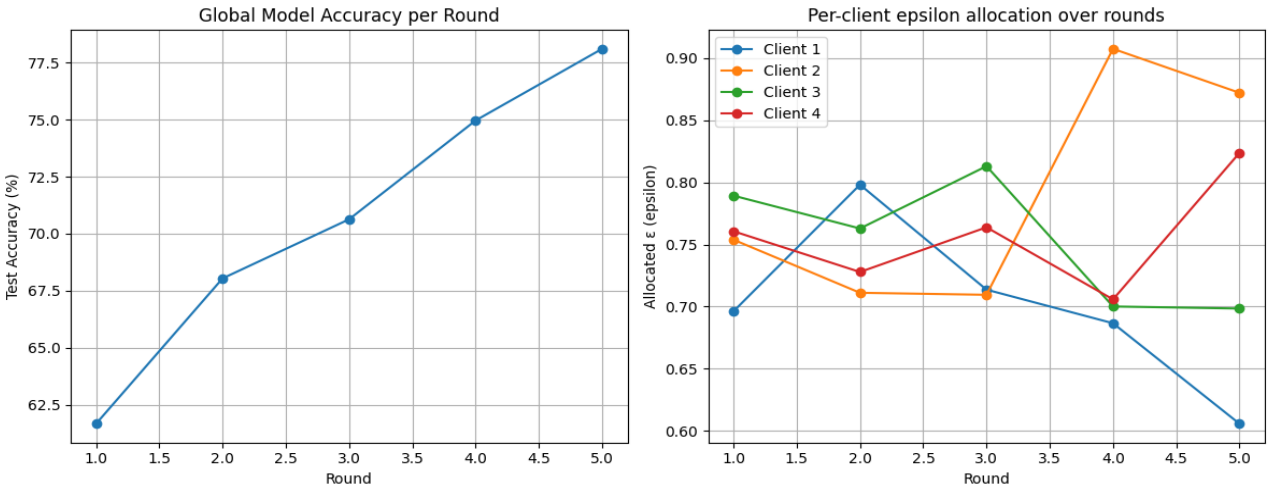

Figure 3

The resulting graphs are interpreted as follows:

4.1 Accuracy vs Noise: In figure 2, the scatter plot shows that accuracy fluctuates with different noise multipliers. Noise values lie mostly between 1.1 and 1.6, and accuracy ranges from 61% to ~78%. Some points show high accuracy even at relatively higher noise levels (~1.4), while others show accuracy drop at similar noise. This indicates that noise alone does not determine global performance; instead, how noise is allocated across clients and rounds matters.

4.2 Global Accuracy per Round: On the left hand side of figure 3, the accuracy improves steadily from 61% (Round 1) to 78% (Round 5). This trend shows that despite fluctuating noise (figure 2), the system's adaptive mechanism ensures consistent learning and convergence. The improvement demonstrates that noise injection is being managed effectively by the personalized allocation strategy.

4.3 Per-client Epsilon Allocation per Round: The right hand side of figure 3 shows how clients receive different ε allocations across rounds: Client 1 steadily decreasing ε (stricter privacy, less importance); Client 2 spikes in Round 4 (~0.91), showing a major contribution in that round; Client 3; fairly stable ε, showing balanced contributions; and Client 4 increases towards Round 5 (~0.82), indicating rising influence. This validates that privacy is not uniform, but dynamically tuned to client contribution and sensitivity.

4.4 Link between Accuracy vs Noise in figure 2 and Global Accuracy Trend on the left hand side of figure 3: Although higher noise typically reduces accuracy, the global accuracy still improves per round. This means the adaptive allocation prevents noise from overwhelming the model, balancing utility and privacy.

4.5 Link between Epsilon Allocation (Right Hand Side of figure 3) and Accuracy vs Noise (Figure 2): Clients with higher ε (less noise) tend to stabilize or improve global accuracy, which explains why the model performs better even when some clients face stricter noise. The scattered performance in figure 3 reflects the heterogeneity of noise assignments across clients, which Graph 3 clarifies.

4.6 Fairness Implications: The per-client ε allocation shows personalization rather than one-size-fits-all. Figure 2's variance in accuracy vs noise confirms this, clients experiencing higher noise don't dominate performance, as others balance the training

5.0 Discussion and Conclusion

This work developed an adaptive personalized differential privacy mechanism integrated into a federated learning framework to enhance privacy protection, model performance, and fairness

among heterogeneous clients. The core innovation involved designing a client contribution scoring method based on proxy gradient norms, which informed a personalized privacy budget allocation algorithm. This approach enabled clients contributing more to the global model to receive proportionally higher privacy budgets, balancing privacy and utility adaptively.

References:

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L.

(2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (pp. 308–318). Association for Computing Machinery. https://doi.org/10.1145/2976749.2978318

Aono, Y., Wang, L., Hayashi, T., & Moriai, S. (2019). Privacy-preserving distributed deep

learning with homomorphic encryption and differential privacy. IEICE Transactions on Information and Systems, E103-D(2), 280–288. https://doi.org/10.3390/electronics8040411

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D.,

Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 1175–1191). Association for Computing Machinery. https://doi.org/10.1145/3133956.3133982

Caldas, S., Wu, P., Li, T., Konečný, J., McMahan, H. B., Smith, V., & Talwalkar, A. (2018).

LEAF: A benchmark for federated settings [Preprint]. arXiv. https://doi.org/10.48550/arXiv.1812.01097

Curran Associates, Inc.Truex, S., Liu, L., Gursoy, M. E., Yu, L., & Wei, W. (2019).

Demystifying differential privacy for machine learning. In 2019 IEEE International Conference on Big Data (Big Data) (pp. 2182–2191). IEEE. https://doi.org/10.1109/BigData47090.2019.9005476

Dwork, C. (2006). Differential privacy. In Automata, languages and programming: 33rd

International Colloquium, ICALP 2006, Venice, Italy, July 10–14, 2006, Proceedings, Part II (pp. 1–12). Springer. https://doi.org/10.1007/11787006_1

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in

private data analysis. In Theory of cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4–7, 2006, Proceedings (pp. 265–284). Springer. https://doi.org/10.1007/11681878_14

Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations

And Trends® in Theoretical Computer Science, 9(3–4), 211–407.https://doi.org/10.1561/0400000042

Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client

level perspective [Paper presentation]. NIPS Workshop on Private Multi-Party Machine Learning, Long Beach, CA, United States.

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K.,

Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., Gruteser, M., ... Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends® in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017).

Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 1273–1282). PMLR.

McMahan, H. B., Ramage, D., Talwar, K., & Zhang, L. (2018). Learning differentially private

recurrent language models. International Conference on Learning Representations (ICLR). https://openreview.net/forum?id=BJ0hF1Z0b

Mironov, I. (2017). Rényi differential privacy. In 2017 IEEE 30th Computer Security

Foundations Symposium (CSF) (pp. 263–275). IEEE. https://doi.org/10.1109/CSF.2017.11

Nasr, M., Shokri, R., & Houmansadr, A. (2019). Comprehensive privacy analysis of deep

learning: Passive and active white-box inference attacks against centralized and federated learning. In 2019 IEEE Symposium on Security and Privacy (SP) (pp. 739–753). IEEE. https://doi.org/10.1109/SP.2019.00065

Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In Proceedings of the

22nd ACM SIGSAC Conference on Computer and Communications Security (pp. 1310–1321). Association for Computing Machinery. https://doi.org/10.1145/2810103.2813687

Smith, V., Chiang, C.-K., Sanjabi, M., & Talwalkar, A. (2017). Federated multi-task learning.

In Advances in neural information processing systems 30 (NeurIPS 2017).

Zhao, Y., Zhang, J., Xu, J., & Zhu, Y. (2020). Local differential privacy-based federated

learning for Internet of Things. IEEE Internet of Things Journal, 7(10), 9530–9538. https://arxiv.org/pdf/2004.08856