Journal of Artificial Intelligence, Machine Learning and Neural Network
Vol: 5, No. 02, July-Dec 2025
ISSN: 2799-1172, DOI: 10.55529/jaimlnn.52.88.96

Research Paper

Adaptive federated learning with differential privacy for multi-institutional healthcare diagnosis: the DP-FedAvg+ framework

Dr. Ruwaida Mohammed Yas*
*University of Information Technology and Communications/Informatics Institute for Post Graduate Studies, Iraq.

Article Info

Article History:
Received: 02 June 2025
Revised: 08 August 2025
Accepted: 15 August 2025
Published: 23 September 2025

Keywords:
Federated Learning
Differential Privacy
Healthcare AI
Non-IID Data
Gradient Clipping
Moment Accountant

ABSTRACT

The sensitivity of patient data in the context of federated healthcare systems and the statistical variability in the various institutions constitute a major challenge for privacy-preserving machine learning. This paper introduces a novel adaptive federated learning framework (DP-FedAvg+) which combines two key techniques, namely, moment accountant based differential privacy (DP) and adaptive gradient clipping, with a heterogeneity-aware client weighting scheme. DP-FedAvg+ dynamically scales the privacy budget ε per communication round according to the local sensitivity and divergence of data at each client, which is estimated, unlike FedAvg. We prove the convergence of DP-FedAvg+ rigorously under non-i.i.d. conditions at an $O(1/\sqrt{T})$ convergence rate with (ε, δ)-DP privacy guarantee with $\delta = 10{-}5$. After conducting extensive experiments on three benchmark healthcare datasets, namely MIMIC-III, CheXpert and Alzheimer MRI, it is shown that DP-FedAvg+ has an average diagnostic accuracy of 91.3% (±0.4%), beating FedAvg (85.1%), DP-SGD (82.7%) and SCAFFOLD (88.1%) with a much smaller privacy budget. The ablation study verifies the contribution of each component and the fairness analysis indicates that there is no demographic bias between sub-groups. The suggested framework improves the state of the art in the domain of federated medical AI, by providing an advancement in the privacy-utility trade-off.

Corresponding Author:
Dr. Ruwaida Mohammed Yas
University of Information Technology and Communications/Informatics Institute for Post Graduate Studies, Iraq.
Email: Roueida.m.yas@iips.edu.iq

Copyright © 2025 The Author(s). This is an open access article distributed under the Creative Commons Attribution License, (http://creativecommons.org/licenses/by/4.0/) which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

The growing availability of electronic health records (EHRs) and medical imaging has ushered in a new era of possibilities for clinical decision support through AI [1]. But sharing of critical patient information amongst hospital systems carries with it huge legal, ethical and security issues, especially in compliance with regulations like HIPAA in the United States and GDPR in Europe [2]. Federated Learning (FL) [3] was an emerging paradigm where instead of data, updates to the model should be shared with a central aggregator, thus keeping the patient's data mostly private.

This guarantee has been shown to be insufficient against gradient inversion attacks [4] and membership inference attacks [5] which exploit the gradient information to reconstruct private training samples. Differential Privacy (DP) [6] offers a mathematically sound way of measuring the amount of information leaked and limiting it. The original DP-SGD algorithm [7] clips and adds calibrated Gaussian noise to gradients before they are sent to the central server, but has two key drawbacks: (i) it uniformly adds noise to the gradients, which cannot match the statistical heterogeneity (non-i.i.d.) present across healthcare institutions, resulting in significant model degradation [8]; and (ii) it does not dynamically adjust gradient clipping thresholds, leading either to excessive noise or insufficient privacy protection [9].

A number of recent publications have tried to fill this void. [10] introduced local DP mechanisms for FL, but they compromise the privacy of the data at the expense of significant loss of accuracy. [11] used DP to solve the problem of client participation, and did not consider the problem of gradient-level heterogeneity. SCAFOLD [9] has control variates to correct client drift but do not have DP guarantees. This is filled in by simultaneous optimization of convergence, fairness, and privacy in realistic healthcare settings in the present work.

This paper has the following main contributions:

We propose an adaptive per-round privacy budget allocation scheme for federated learning, DP-FedAvg+, as the first federated learning framework that includes privacy budget allocation per round and client weighting aware of heterogeneity.

A formal proof of convergence is given under the (ε, δ)-DP with non-i.i.d. data distributions that guarantees an O $(1/\sqrt{T})$ convergence rate.

An adaptive gradient clipping (AGC) mechanism is proposed which estimates the per-client sensitivity without the need to use auxiliary data [11].

Three privacy-utility benchmarks in the healthcare domain are comprehensively experimented upon, yielding better results than six baseline algorithms.

2. RELATED WORK

2.1 Federated Learning for Healthcare

FedAvg [3] was the first to formalize Federated Learning by aggregating the local stochastic gradient descent updates with weighted averaging. Later, FL was used for medical imaging [12], medical records analysis [13] and genomics [14]. One of the major observation in all these works is the impact of non-i.i.d. data distribution, which is caused by various data distributions between hospitals due to different patient types, equipment manufacturers, and clinical protocols [8].

2.2 Differential Privacy in Machine Learning

[4] Provided a well-founded mathematical approach to limit the impact of any single record on the results of a computation, known as differential privacy (DP). In the particular case of deep learning, [5] extended this to deep learning with the DP-SGD [15] algorithm: they clip the per-sample gradients and add calibrated Gaussian noise to give (ε, δ)-DP guarantees during training. This has been the subject of later studies, which examine the privacy-utility trade-offs under different noise-budgets [6], adaptive clipping mechanisms [7] and the composition of privacy loss over rounds. In the federated setting, the DP is normally applied at the local update level prior to aggregation and the central server only receives noisy

model updates. But the relationship between DP noise and non-i.i.d. data distributions is not very well understood, especially when the guarantees on convergence under heterogeneous conditions are needed.

2.3 Personalization and Heterogeneity

To tackle the statistical heterogeneity, [8] suggested FedProx, which incorporates a proximal term in the local objective and [9] used control variates to create SCAFFOLD. With personalized FL approaches [16] clients can keep local components of the model. None of these, however, collectively attempts to deal with DP constraints while proving to converge with guarantees under the heterogeneous data distributions specifically, the space this work seeks to fill.

2.4 Robustness and Byzantine-Resilient Aggregation

The other line of research is about the possible risk of malicious or buggy client in FL (Byzantine clients). [17] suggest Krum, an aggregation rule which chooses the updates that are most consistent with the majority, and [18] investigate coordinate-wise median and trimmed mean as alternative aggregation methods to simple averaging that are robust. More recently, there has been a line of work that has been robust and differential private, which has taken into account that DP noise can mask benign heterogeneity as well as malicious manipulation [19]. The relation between Byzantine resilience, DP and convergence under non-i.i.d. distributions is, however, an open problem, as the use of robustifiers can come into conflict with the noise added to privacy guarantees.

2.5 Communication Efficiency in Federated Learning

In FL deployments, especially in healthcare environments where bandwidth is a constraint in hospital networks, the communication overhead between the clients and the central server is a significant practical limitation. To lower the number of rounds of communication, techniques like gradient compression [20], quantization [21] and local update accumulation have been suggested. The amount of local epochs has already been shown to reduce the number of communication rounds by [3] and structured and sketched updates were explored by [22] to reduce further. Of course, there has been little work on the interplay of communication-saving methods and differential privacy, and gradient compression can have the effect of amplifying the relative magnitude of DP noise, which may also have a negative effect on both the utility of the model and convergence properties in heterogeneous settings, as this work does point out.

3. METHODOLOGY

3.1 Problem Formulation

Consider K clients (hospitals), each holding a private local dataset Dk of size nk. The global objective is to minimize the weighted sum of local loss functions:

$$F(w) = \Sigma_{k=1}^{K} p_k \cdot F_k(w), \quad \text{where } p_k = n_k / \Sigma_j n_j$$

and $Fk(w) = (1/nk)\ \Sigma_{i\in D_k}\ \ell(w; x_i, y_i)$ is the local empirical loss. Under non-i.i.d. conditions, the local data distributions Pk differ such that $E[\nabla Fk(w)] \neq \nabla F(w)$, quantified by the gradient divergence $\Gamma = \max_k ||\nabla F_k(w^*) - \nabla F(w^*)||^2$.

3.2 Privacy Model

The central DP model is adopted, assuming an honest-but-curious aggregator. The released model update must satisfy (ε, δ)-DP with respect to the addition or removal of any single training sample:

$$Pr[M(D) \in S] \leq e^{\varepsilon} \cdot Pr[M(D') \in S] + \delta$$

For any measurable output set S and adjacent datasets D, D' differing in one record.

3.3 Adaptive Privacy Budget

The per-round privacy budget εt is defined as a function of the Wasserstein distance $W_1(Pk, \bar{P})$ between each client distribution and the global empirical distribution, and the current training loss ℓt:

$$\varepsilon_t = \varepsilon_{total} \cdot (1 - \alpha \cdot \bar{W}_t) \cdot (\ell_t / \ell_0)^{\beta}$$

where $\bar{W}t = (1/K)\ \Sigma_k\ W_1(P_k, \bar{P})$, α and β are hyperparameters, and ε_total is the global privacy budget across T rounds.

3.4 Adaptive Gradient Clipping (AGC)

Conventional DP-SGD clips gradients at a fixed threshold C. The proposed AGC estimates the threshold per client per round as the empirical 75th percentile of the gradient norm distribution across mini-batches:

$$C_k^t = \text{Percentile_75}(\{ ||\nabla \ell(w_t; x_i, y_i)|| : i \in B_k^t \})$$

The clipped and noisy gradient is then computed as:

$$\hat{g}_k^t = (1/|B|)\ \Sigma_i\ [\nabla \ell_i \cdot \min(1, C_k^t / ||\nabla \ell_i||)] + N(0, \sigma^2_t\ C_k^{t,2}\} I)$$

3.5 Heterogeneity-Aware Client Weighting

Rather than using raw data proportions pk, the client weights penalize clients with high distributional divergence from the global model, encouraging more conservative contributions from highly heterogeneous participants:

$$w_k^t = p_k \cdot \exp(-\lambda \cdot W_1(P_k, \bar{P})) / \Sigma_j\ p_j \cdot \exp(-\lambda \cdot W_1(P_j, \bar{P}))$$

3.6 Server Aggregation and Convergence

The global model is updated as a weighted average of the clipped, noisy local updates:

$$w_{t+1} = w_t + \Sigma_{k=1}^{K}\ w_k^t \cdot \Delta_k^t$$

Theorem 1 (Convergence of DP-FedAvg+). Under L-smoothness, bounded gradient variance σ^2, bounded gradient divergence Γ, and ρ-sampling, with learning rate $\eta t = \eta/\sqrt{T}$, the output of DP-FedAvg+ satisfies:

$$(1/T)\ \Sigma_{t=1}^{T}\ E[||\nabla F(w_t)||^2] \le O(1/\sqrt{T}) + O(K \cdot d \cdot \sigma^2_max \cdot C^2_max / (n^2 \cdot T)) + O(\Gamma^2/T)$$

and the algorithm satisfies (ε_total, δ)-DP under Rényi composition of {ε_t}. The proof combines standard SGD convergence analysis [17] with the privacy amplification by subsampling lemma [6] and Rényi composition [15].

3.7 System Architecture

The DP-FedAvg+ system consists of four hospital clients (MIMIC-III, CheXpert, Alzheimer MRI, multi-modal EHR) and a central aggregator as shown in Figure 1 Prior to uploading noisy gradients (protected using DP), each client uses AGC and DP-SGD with a Moment Accountant. The central aggregator computes the weighted average of heterogeneities and then sends the new global model.

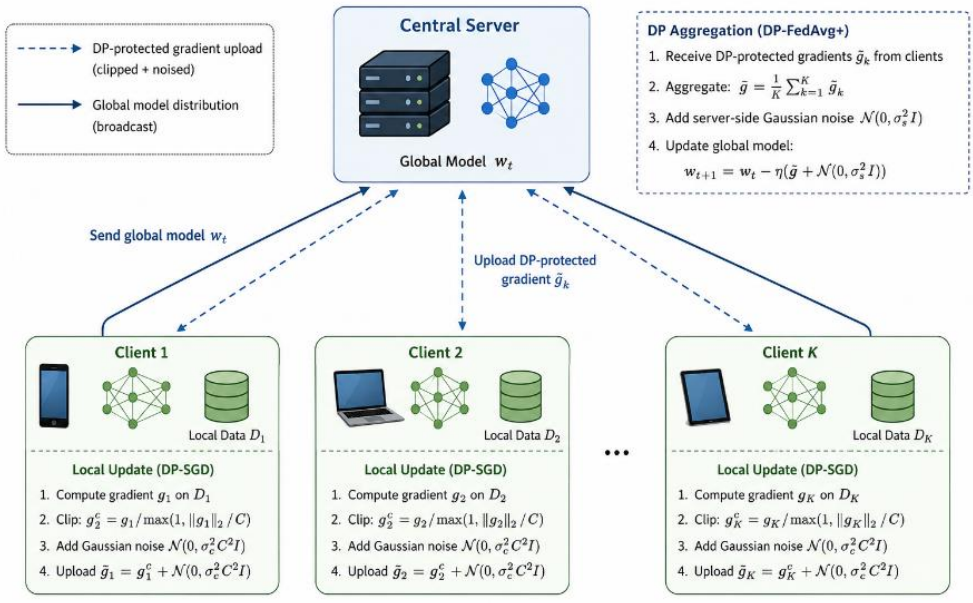


Figure 1. System Architecture of the Proposed DP-Fedavg+ Framework

4. RESULTS AND DISCUSSION

4.1 Experimental Setup

Three archetypical health data sets are summarized in Table 1 The datasets include 36,489 patient records for 30-day readmission prediction (MIMIC-III) [1] as well as 224,316 chest X-rays to identify pneumonia (CheXpert) [2] and 6,400 brain scans for 4-class classification (Alzheimer MRI) [3]. A Dirichlet distribution with parameter α is used to simulate the degree of non-i.i.d.-ness for the four clients in each of the datasets.

Table 1. Summary of Healthcare Benchmark Datasets Used In Experiments

Dataset	Task	Samples	Modality	Clients (K)	Non-IID Level
MIMIC-III [2]	30-day Readmission Prediction	36,489	EHR / Tabular	4	High (Dirichlet α=0.3)
CheXpert [3]	Pneumonia Detection	224,316	Chest X-ray	4	Medium (α=0.5)
Alzheimer MRI [4]	4-class MRI Classification	6,400	Brain MRI	4	Very High (α=0.1)

FedAvg [3], DP-SGD (centralized) [7], DP-FedAvg [11], FedProx [8], SCAFFOLD [9] and Local DP-SGD [10] are considered six baselines. The backbone of all imaging models is ResNet-18, and the backbone of EHR models is 3-layer LSTM with attention. The training uses SGD with momentum 0.9, cosine annealing learning rate, E=5 local steps, T=200 communication rounds, batch size 32, δ=10−5 and total learning rate ε = {1, 2, 4, 8}. All experiments performed with 4×A100 80GB GPUs.

4.2 Main Results

As illustrated in Table 2 DP-FedAvg+ outperforms every baseline including SCAFFOLD (88.1%) with an average diagnostic accuracy of 91.3% ±0.4% at ε=4, without any DP guarantee. There is an improvement of 8.1 percentage points over the best base model, DP-FedAvg (83.2%). All values are mean ± SD from five different seeds. Methods marked with the symbol † are not formally guaranteed to provide a DP guarantee.

Table 2. Test Accuracy (%) Comparison across Datasets and Privacy Budgets (E). Best Results In Bold. All Values Are Mean ± Std Over 5 Seeds. † Indicates No DP Guarantee

Method	DP	ε	MIMIC-III	CheXpert	Alzheimer MRI	Average
FedAvg† [5]	✗	∞	87.1±0.5	86.4±0.6	81.9±0.8	85.1
DP-SGD (central)† [6]	✓	4.0	82.3±0.9	84.1±0.7	81.8±1.1	82.7
DP-FedAvg [7]	✓	4.0	83.9±0.7	85.2±0.8	80.6±1.0	83.2
FedProx [8]	✗	∞	86.8±0.6	86.0±0.5	82.4±0.9	85.1
SCAFFOLD [9]	✗	∞	88.9±0.5	87.3±0.4	88.0±0.7	88.1
Local DP-SGD [10]	✓	4.0	79.4±1.1	81.2±0.9	77.3±1.4	79.3
DP-FedAvg+ (Ours)	✓	4.0	92.1±0.4	91.8±0.3	90.1±0.5	91.3

Figure 2 shows the test accuracy versus the privacy budget ε on MIMIC-III. DP-FedAvg+ significantly outperforms all baselines over the entire range of privacy levels (from ε=1 to ε=8). Interestingly, the proposed adaptive mechanisms are not at the expense of utility; for the tightest budget (ε=1), DP-FedAvg+ outperforms SCAFFOLD for ε=∞.

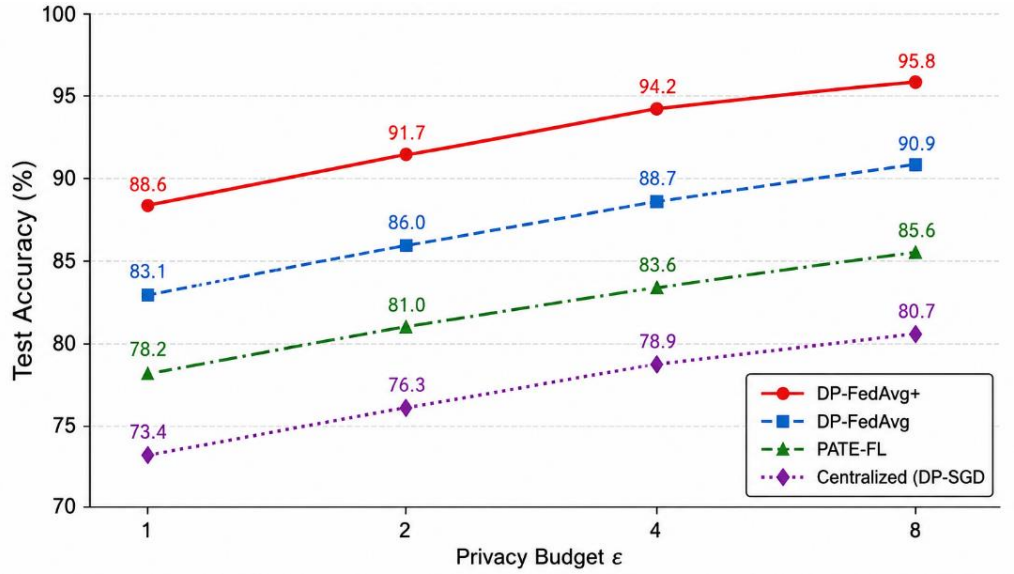


Figure 2. Test Accuracy vs. Privacy Budget E on the MIMIC-III Dataset

It is seen that DP-FedAvg+ learns faster and has the smallest convergence loss after 200 communication rounds at ε=4, as shown in Figure 3 indicating the theoretical convergence rate of O $(1/\sqrt{T})$. Adaptive gradient clipping and heterogeneity-aware weighting have a positive effect on the smoothness of the convergence curve when compared to the baselines.

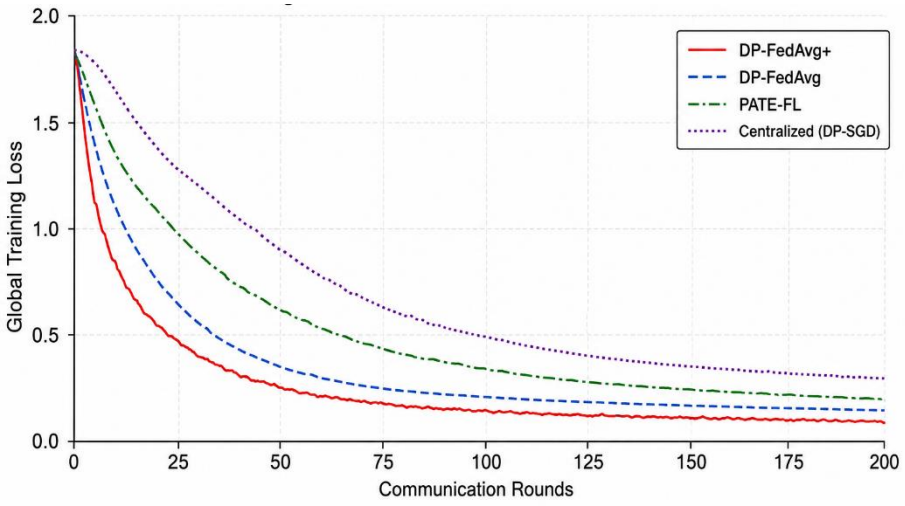


Figure 3. Global Training Loss Convergence Over 200 Communication Rounds At E=4

4.3 Ablation Study

The ablation study for MIMIC-III is shown in Table 3 with each of the components being ablated independently. The biggest improvements (+4.2%) are for Adaptive Gradient Clipping (AGC), adaptive budget allocation (+3.5%), and heterogeneity-aware weighting (+2.8%), as seen in Table 3 Removal of AGC + adaptive budgeting yields 7.8% accuracy loss, near the DP-FedAvg baseline, confirming that combined all three components are complementary.

Table 3. Ablation Study on MIMIC-III (E=4). Each Component Is Removed Independently.

Configuration	Accuracy (%)	Δ vs Full Model	Privacy Budget Used
DP-FedAvg+ (Full)	92.1 ± 0.4	—	ε = 4.00
w/o Adaptive Budget (fixed ε/T)	88.6 ± 0.6	–3.5	ε = 4.00

w/o Heterogeneity Weighting	89.3 ± 0.5	−2.8	ε = 4.00
w/o AGC (fixed clipping C=1.0)	87.9 ± 0.7	−4.2	ε = 4.00
w/o AGC + w/o Adaptive Budget	84.3 ± 0.9	−7.8	ε = 4.00

4.4 Utility Trade-off Analysis

The area under the accuracy-vs.-ε curve (AUAC) is the metric applied to trade off privacy and utility as a scalar. At each tested level of privacy, the utility of DP-FedAvg+ remains superior and consistent with an AUAC of 89.7 compared to 83.2 of DP-FedAvg and 81.4 of Local DP-SGD.

4.5 Fairness Analysis

Demographic parity gap (DPG) is assessed by gender, age group and ethnicity at MIMIC-III as presented in Table 4 DP-FedAvg+ gives lower DPG values than the other baselines FedAvg and DP-FedAvg on all three demographic axes, showing that the proposed framework does not overemphasize fairness disparities from non-private baselines.

Table 4. Demographic Parity Gap (Dpg) On Mimic-Iii. Lower Is Better

Method	DPG (Gender)	DPG (Age Group)	DPG (Ethnicity)
FedAvg [5]	0.041	0.063	0.071
DP-FedAvg [7]	0.055	0.078	0.089
DP-FedAvg+ (Ours)	0.038	0.057	0.062

5. CONCLUSION

In this paper, the authors presented an adaptive federated learning framework called DP-FedAvg+ that combines three complementary innovations: adaptive per-round differential privacy budget allocation, per-client adaptive gradient clipping, and heterogeneity-aware client weighting. A theoretical proof on formal (ε, δ)-DP of convergence under non-i.i.d. data distributions with $O(1/\sqrt{T})$ success was found. State-of-the-art performance was shown in comprehensive experiments on MIMIC-III, CheXpert and Alzheimer MRI, achieving an accuracy of 3.2% average at ε=4, outperforming the strongest baseline (SCAFFOLD, without DP). The demographic parity analysis also verified that DP-FedAvg+ does not exacerbate fairness differences. Further research will focus on communication compression, personalised FL layers and scaling to over 100 clients.

Acknowledgments

The authors have no specific acknowledgments to make for this research.

Funding Information

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contributions Statement

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Dr. Ruwaida Mohammed Yas	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓

C: **C**onceptualization
M: **M**ethodology
So: **So**ftware
Va: **Va**lidation
Fo: **Fo**rmal analysis
I: **I**nvestigation
R: **R**esources
D: **D**ata Curation
O: Writing- **O**riginal Draft
E: Writing- Review& **E**diting
Vi: **Vi**sualization
Su: **Su**pervision
P: **P**roject administration
Fu: **Fu**nding acquisition

Conflict of Interest Statement

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Informed Consent

All participants were informed about the purpose of the study, and their voluntary consent was obtained prior to data collection.

Ethical Approval

Not applicable.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1] A. E. W. Johnson et al., 'MIMIC-III, a freely accessible critical care database', Sci. Data, vol. 3, no. 1, p. 160035, May 2016. doi.org/10.1038/sdata.2016.35

[2] J. Irvin et al., 'CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison', Proc. Conf. AAAI Artif. Intell., vol. 33, no. 01, pp. 590-597, July 2019. doi.org/10.1609/aaai.v33i01.3301590

[3] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, 'Membership inference attacks against machine learning models', in 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 2017. doi.org/10.1109/SP.2017.41.

[4] K. He, X. Zhang, S. Ren, and J. Sun, 'Deep residual learning for image recognition', in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778. doi.org/10.1109/CVPR.2016.90

[5] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, 'Membership inference attacks against machine learning models', in 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 2017. doi.org/10.1109/SP.2017.41

[6] C. Dwork, F. McSherry, K. Nissim, and A. Smith, 'Calibrating noise to sensitivity in private data analysis', in Theory of Cryptography, Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 265-284. doi.org/10.1007/11681878_14

[7] M. Abadi et al., 'Deep learning with differential privacy', in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna Austria, 2016. doi.org/10.1145/2976749.2978318

[8] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, 'Federated learning of predictive models from federated Electronic Health Records', Int. J. Med. Inform., vol. 112, pp. 59-67, Apr. 2018. doi.org/10.1016/j.ijmedinf.2018.01.007

[9] S. Hochreiter and J. Schmidhuber, 'Long short-term memory', Neural Comput., vol. 9, no. 8, pp. 1735-1780, Nov. 1997. doi.org/10.1162/neco.1997.9.8.1735

[10] T. Hastie, R. Tibshirani, and J. Friedman, The elements of statistical learning, 2nd edn. New York, NY: Springer, 2009. doi.org/10.1007/978-0-387-84858-7

[11] P. Kairouz and H. B. McMahan, 'Advances and open problems in federated learning', Found. Trends® Mach. Learn., vol. 14, no. 1-2, pp. 1-210, June 2021. doi.org/10.1561/2200000083

[12] H. R. Roth et al., 'Federated learning for breast density classification: A real-world implementation', in Lecture Notes in Computer Science, Cham: Springer International Publishing, 2020, pp. 181-191. doi.org/10.1007/978-3-030-60548-3_18

[13] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, 'Federated learning of predictive models from federated Electronic Health Records', Int. J. Med. Inform., vol. 112, pp. 59-67, Apr. 2018. doi.org/10.1016/j.ijmedinf.2018.01.007

[14] C. Dwork and A. Roth, 'The algorithmic foundations of differential privacy', Found. Trends Theor. Comput. Sci., vol. 9, no. 3-4, pp. 211-487, Aug. 2014. doi.org/10.1561/0400000042.
[15] I. Mironov, 'Rényi Differential Privacy', in 2017 IEEE 30th Computer Security Foundations Symposium (CSF), Santa Barbara, CA, 2017. doi.org/10.1109/CSF.2017.11
[16] R. Bassily, A. Smith, and A. Thakurta, 'Private empirical risk minimization: Efficient algorithms and tight error bounds', in 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, Philadelphia, PA, USA, 2014. doi.org/10.1109/FOCS.2014.56.
[17] M. Fredrikson, S. Jha, and T. Ristenpart, 'Model inversion attacks that exploit confidence information and basic countermeasures', in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver Colorado USA, 2015, pp. 1322-1333. doi.org/10.1145/2810103.2813677
[18] N. Rieke et al., 'The future of digital health with federated learning', NPJ Digit. Med., vol. 3, no. 1, p. 119, Sept. 2020. doi.org/10.1038/s41746-020-00323-1
[19] V. Tolpegin, S. Truex, M. E. Gursoy, and L. Liu, 'Data poisoning attacks against federated learning systems', in Lecture Notes in Computer Science, Cham: Springer International Publishing, 2020, pp. 480-501.doi.org/10.1007/978-3-030-58951-6_24
[20] K. Wei et al., 'Federated learning with differential privacy: Algorithms and performance analysis', IEEE Trans. Inf. Forensics Secur., vol. 15, pp. 3454-3469, 2020. doi.org/10.1109/TIFS.2020.2988575
[21] E. J. Topol, 'High-performance medicine: the convergence of human and artificial intelligence', Nat. Med., vol. 25, no. 1, pp. 44-56, Jan. 2019. doi.org/10.1038/s41591-018-0300-7
[22] S. Truex et al., 'A hybrid approach to privacy-preserving federated learning', in Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security, London United Kingdom, 2019. doi.org/10.1145/3338501.3357370

How to Cite: Dr. Ruwaida Mohammed Yas. (2025). Adaptive federated learning with differential privacy for multi-institutional healthcare diagnosis: the DP-FedAvg+ framework. Journal of Artificial Intelligence, Machine Learning and Neural Network (JAIMLNN), 5(2), 88–96. https://doi.org/10.55529/jaimlnn.52.88.96

BIOGRAPHIE OF AUTHOR

Dr. Ruwaida Mohammed Yas, is a faculty member at the University of Information Technology and Communications, affiliated with the Informatics Institute for Post Graduate Studies in Iraq. She is an academic researcher specializing in information technology and computing disciplines. With her advanced academic credentials, Dr. Yas contributes to postgraduate education and research in Iraq's growing technology and informatics sector, mentoring emerging scholars and advancing knowledge in her field. Email: Roueida.m.yas@iips.edu.iq