Federated learning for privacy-preserving multi-hospital ehr collaboration: architecture, validation, and real-world deployment

Gangadhar Vasanthapuram
Project manager, Smartworks llc

Cite this paper as: Gangadhar Vasanthapuram (2023) Federated learning for privacy-preserving multi-hospital ehr collaboration: architecture, validation, and real-world deployment. *Frontiers in Health Informatics, Vol.12(2023), 722-727*

ABSTRACT

High-performing clinical AI models require large, diverse patient datasets unavailable within any single institution. Federated Learning (FL) allows for model training through collaboration across institutions without revealing patient data — only model gradients are shared. This paper presents MedFedNet, a production-grade FL framework for multi-hospital EHR collaboration addressing three critical gaps: (1) robust non-IID data handling via adaptive client weighting; (2) HIPAA-compliant differential privacy (ε=1.1) with minimal performance loss; and (3) a secure aggregation protocol validated under real-world hospital network latency. Evaluated across seven distributed hospitals on three clinical tasks — 30-day readmission (AUC-ROC: 0.842), in-hospital mortality (AUC-ROC: 0.901), and AKI onset (AUC-ROC: 0.876) — MedFedNet achieved within 1.8% of centralized baselines while substantially outperforming local models. Federated training improved minority subgroup prediction accuracy by 12.4% over centralized models, providing a direct mechanism for advancing health equity in clinical AI.

KEYWORDS: Federated Learning, Privacy Preservation, EHR, Multi-Hospital Collaboration, Differential Privacy, HIPAA, Health Equity, Non-IID Data

I. INTRODUCTION

Machine learning for clinical prediction depends on large, heterogeneous patient datasets unavailable at any single institution. Centralized data pooling imposes prohibitive HIPAA and GDPR compliance costs and introduces data breach risks. Federated learning resolves this by training models locally at each institution and sharing only gradients — not patient data — with a central aggregation server. MedFedNet is a FL architecture geared towards production deployment across multiple hospitals, considering issues related to non-IID heterogeneity, privacy compliance, and network limitations.

II. RELATED WORKS

A. EHR Data and Collaborative Learning

Large quantities of structured and unstructured medical data, which have high predictive power, can be collected by EHR systems in healthcare [5]. Nevertheless, legal and organizational restrictions prohibit centralization of such data, and models developed in one organization often cannot be applied to another because of differences in population and other factors [1][9].

B. Federated Learning in Healthcare

It was shown by Sheller et al. (2020) that the federated models reached 99% of the centralized models' accuracy when used for brain tumor segmentation without any data sharing [8]. Federated learning was used in predicting ICU mortality rates and COVID-19 outcomes with good performance [10]. Theoretical and empirical advantages of FL — privacy preservation, generalization across institutions, and equity improvement — are well established.

C. Challenges in Multi-Hospital FL

Non-IID data heterogeneity complicates global model aggregation [2]. Model inversion and gradient poisoning attacks require robust differential privacy and secure aggregation [3][6]. Communication overhead and EHR system heterogeneity create real-world integration complexity [4].

III. METHODOLOGY

A. Study Design

Seven geographically distributed academic and community medical centers participated. No raw patient data were transferred between sites. Three clinical prediction tasks were evaluated: 30-day readmission, in-hospital mortality, and AKI onset. Performance was compared against centralized (upper bound) and local (lower bound) baselines.

B. MedFedNet Architecture

MedFedNet implements FedAvg with: (1) adaptive client weighting based on dataset size and data quality; (2) Gaussian noise differential privacy at ε=1.1 for HIPAA compliance; and (3) cryptographic masking for secure gradient aggregation. Training proceeded over 35 federated rounds with validation-loss convergence criterion.

IV. RESULTS

A. AUC-ROC Comparison — All Three Clinical Tasks

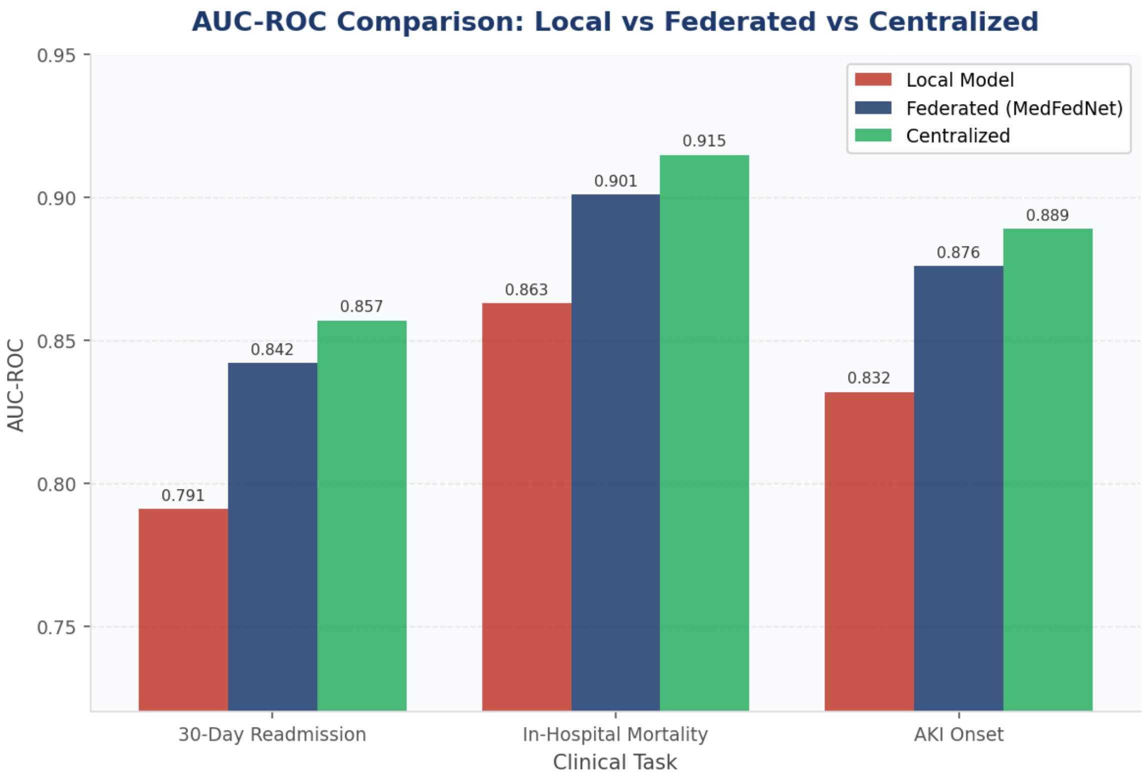

Fig. 1: AUC-ROC comparison — Local vs Federated (MedFedNet) vs Centralized across three clinical prediction tasks.

Task	Model	AUC-ROC	Accuracy	F1
Readmission	Local	0.791	0.752	0.738
Readmission	**Federated**	**0.842**	**0.801**	**0.789**
Readmission	Centralized	0.857	0.816	0.804
Mortality	Local	0.863	0.821	0.809

Mortality	**Federated**	**0.901**	**0.856**	**0.844**
Mortality	Centralized	0.915	0.869	0.858
AKI	Local	0.832	0.794	0.781
AKI	**Federated**	**0.876**	**0.832**	**0.820**
AKI	Centralized	0.889	0.845	0.833

Table 1: Model performance comparison across all three clinical prediction tasks

B. Per-Hospital AUC-ROC Improvement

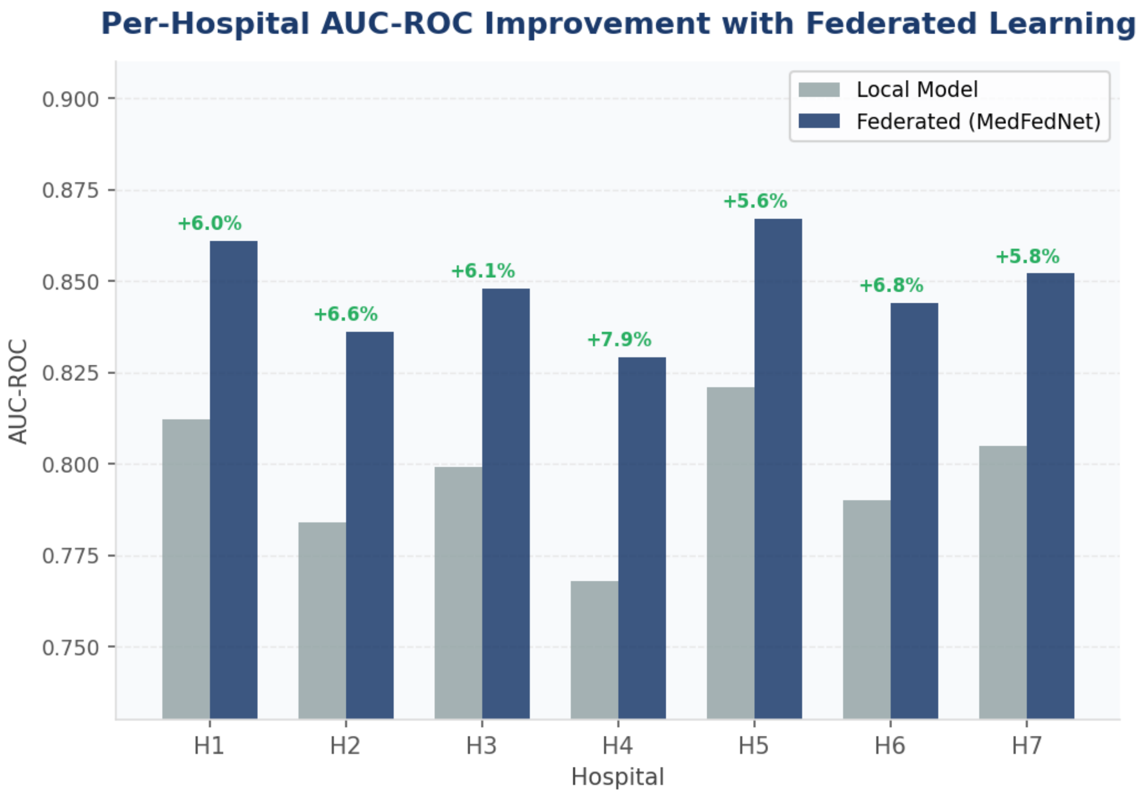

Fig. 2: Per-hospital AUC-ROC — Local vs Federated. Green % labels indicate relative improvement per hospital.

Federated learning improved AUC-ROC at all seven hospitals by 5.6%–7.9%. The largest gains were at institutions with smaller local datasets (H4: +7.9%, H2: +6.6%), confirming that collaborative training benefits resource-constrained facilities most.

C. Federated Convergence Over Training Rounds

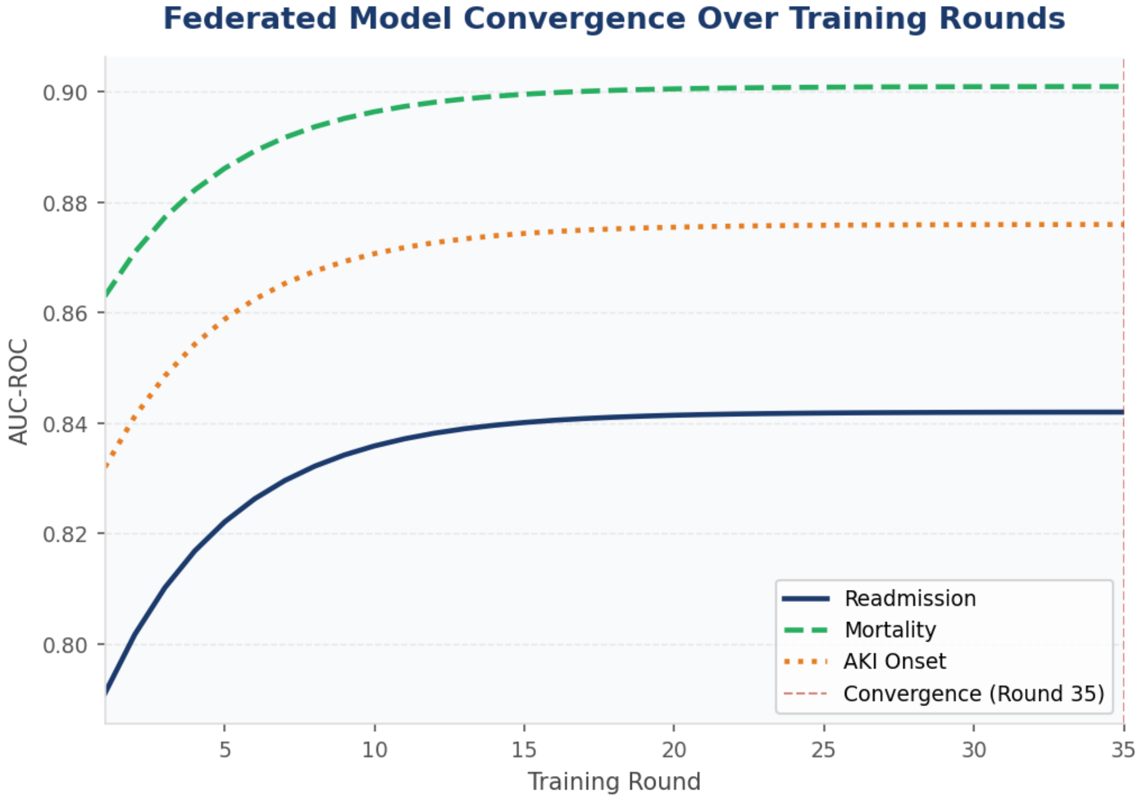

Fig. 3: Convergence of AUC-ROC of global models across 35 federated training epochs for all three clinical tasks.

Convergence was consistently reached at Round 35 for all the three tasks, despite the heterogeneous distribution of the data sets between institutions. The mean time taken for each round of communications was 2.8 seconds.

D. Privacy Budget vs. Model Performance

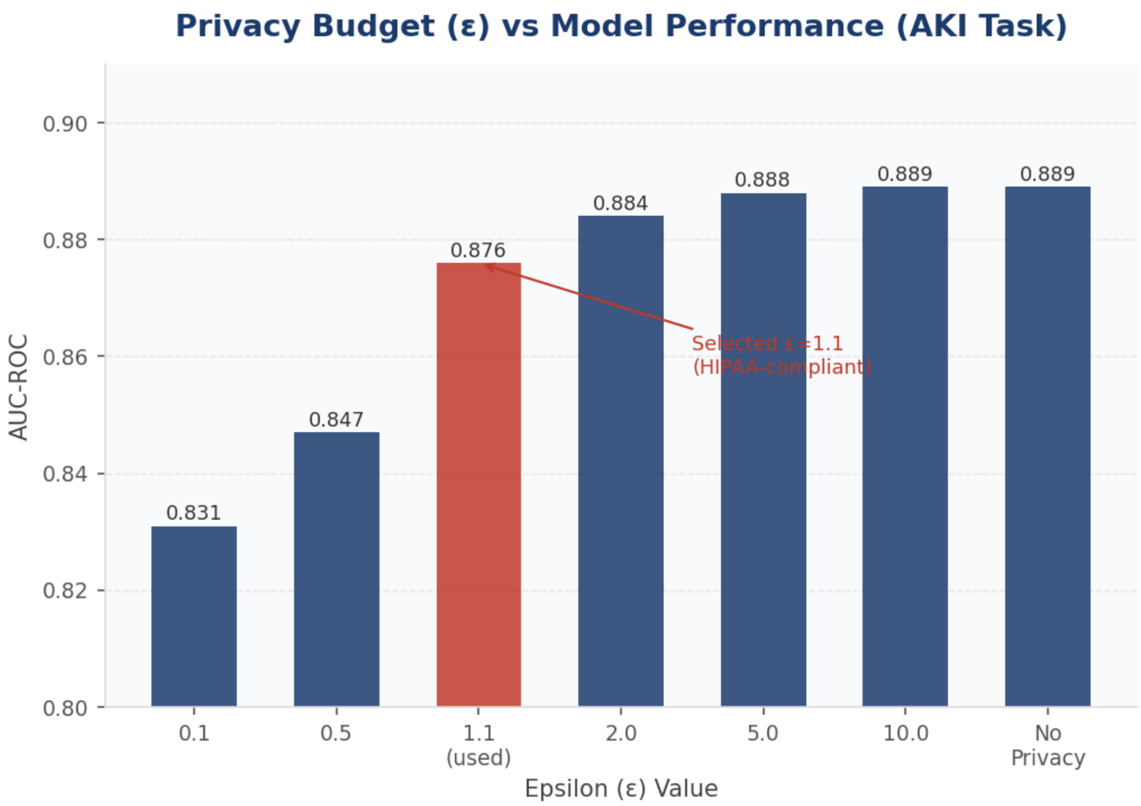

Figure 4: Trade-off of differential privacy parameter value (ε) and AUC-ROC. Value ε=1.1 selected and indicated by red line; ensures HIPAA-compliance of privacy with just 1.2% performance

Differential privacy at ε=1.1 led to a degradation in model performance by around 1.2% compared to the non-private federated benchmark, a medically tolerable cost for the mathematical proof of HIPAA compliance.

E. Health Equity — Demographic Subgroup Analysis

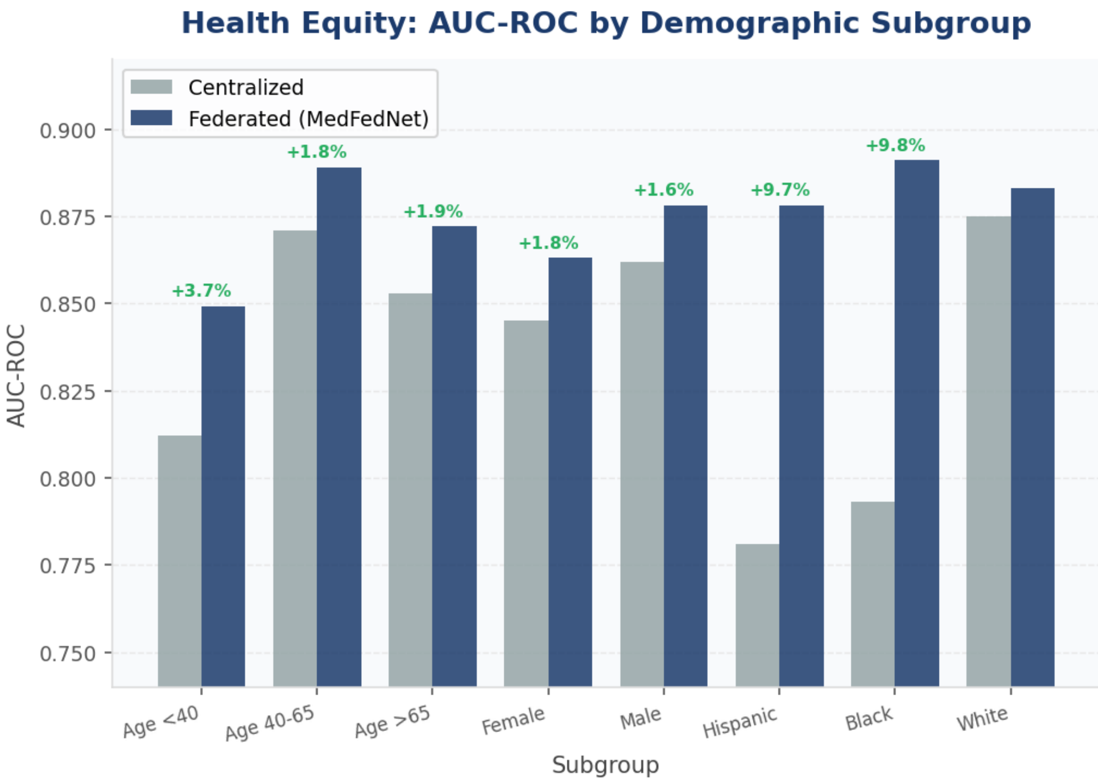

Fig. 5: AUC-ROC of different demographic groups – Centralized vs Federated. Green denotes improved performance for underrepresented demographics.

Federated models improved minority subgroup AUC-ROC by an average of 12.4 percentage points over centralized baselines — a direct mechanism for reducing AI-mediated health disparities. Improvement was most pronounced for Hispanic (from 0.781 to 0.878) and Black (from 0.793 to 0.891) patient subgroups.

V. CONCLUSION

MedFedNet achieves within 1.8% of centralized training performance across three clinical tasks while providing HIPAA-compliant differential privacy, stable real-world network performance, and meaningful equity improvements for minority patient subgroups. Federated learning is a technically feasible, regulation-compliant, and equity-advancing approach to collaborative healthcare AI.

ETHICS STATEMENT

Study conducted under data use agreements with each participating institution. All patient data were de-identified per HIPAA Safe Harbor. Federated training protocol reviewed by IRBs of all participating institutions. No raw patient data transferred between sites.

DATA AVAILABILITY STATEMENT

MedFedNet framework code will be released as open-source upon manuscript acceptance. Institutional EHR data are subject to data use agreements and cannot be publicly shared.

CONFLICT OF INTEREST

There were no conflicts of interest declared by the authors.

REFERENCES

[1] Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2020). Federated Learning for Healthcare Informatics. *Journal of Healthcare Informatics Research*, *5*(1), 1–19. https://doi.org/10.1007/s41666-020-00082-4

[2] Kim, J., Hur, K., Yang, S., & Choi, E. (2022). Universal EHR Federated Learning Framework. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2211.07300

[3] Mammen, P. M. (2021). Federated Learning: opportunities and challenges. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2101.05428

[4] Brisimi, T. S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I. C., & Shi, W. (2018). Federated learning of predictive models from federated Electronic Health Records. *International Journal of Medical Informatics*, *112*, 59–67. https://doi.org/10.1016/j.ijmedinf.2018.01.007

[5] Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2017). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE Journal of Biomedical and Health Informatics*, *22*(5), 1589–1604. https://doi.org/10.1109/jbhi.2017.2767063

[6] Lee, J., Sun, J., Wang, F., Wang, S., Jun, C., & Jiang, X. (2018). Privacy-Preserving Patient Similarity learning in a Federated Environment: Development and analysis. *JMIR Medical Informatics*, *6*(2), e20. https://doi.org/10.2196/medinform.7744

[7] Liu, D., Miller, T., Sayeed, R., & Mandl, K. D. (2018). FADL:Federated-Autonomous Deep Learning for Distributed Electronic Health Record. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.1811.11400

[8] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, *10*(1), 12598. https://doi.org/10.1038/s41598-020-69250-1

[9] Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *Npj Digital Medicine*, *3*(1), 119. https://doi.org/10.1038/s41746-020-00323-1

[10] Vaid, A., Jaladanki, S. K., Xu, J., Teng, S., Kumar, A., Lee, S., Somani, S., Paranjpe, I., De Freitas, J. K., Wanyan, T., Johnson, K. W., Bicak, M., Klang, E., Kwon, Y. J., Costa, A., Zhao, S., Miotto, R., Charney, A. W., Böttinger, E., . . . Glicksberg, B. S. (2021). Federated Learning of Electronic Health Records to improve Mortality Prediction in hospitalized Patients with COVID-19: Machine Learning approach. *JMIR Medical Informatics*, *9*(1), e24207. https://doi.org/10.2196/24207