Journal of Computer and Electrical Sciences (JCES)- Volume-1 Issue-4, April 2026
ISSN: 3049-2602

FEDERATED LEARNING SYSTEM FOR HEALTHCARE USING AI

Brahma Reddy Alla[1], **Mrs. N.V.L. Manaswini**[2]

[1]Student, Department of Computer Science & Engineering
Andhra Loyola Institute of Engineering & Technology, Vijayawada, Andhra Pradesh, India

[2]Assistant Professor, Department of Computer Science & Engineering
Andhra Loyola Institute of Engineering & Technology, Vijayawada, Andhra Pradesh, India

Email id: brahmareddyalla8@gmail.com

Abstract: Artificial Intelligence (AI) is widely used in healthcare for medical diagnosis, but centralized data collection raises serious privacy and security concerns. This project presents a Federated Learning System for Healthcare using AI, which enables multiple healthcare institutions to collaboratively train machine learning models without sharing sensitive patient data. In the proposed system, ten distributed hospital nodes independently train Logistic Regression classifiers on their local medical imaging datasets covering ten disease categories including Diabetic Retinopathy, Cardiac Disease, Breast Cancer, Brain Tumor, Lung Pneumonia, Skin Melanoma, Glaucoma, Bone Fracture, Gallstone, and Disc Herniation. Only model weight parameters are shared with the central server using the Federated Averaging (FedAvg) algorithm, and Laplace differential privacy noise is applied to each weight vector before transmission to provide an additional mathematical privacy guarantee. This approach preserves patient data privacy while achieving effective medical image classification across diverse and distributed hospital datasets. Experimental results demonstrate a global FedAvg accuracy of 83.97% across 2,864 training samples from ten hospital nodes without any raw patient data leaving the local hospital boundary. The proposed system is scalable and privacy-preserving, suitable for modern healthcare applications requiring compliance with HIPAA and GDPR data regulations. Federated learning enables collaboration among institutions without compromising patient confidentiality, and the system reduces the risk of data leakage by ensuring that only differentially private model parameters are communicated across the federation network.

Keywords: *Federated Learning, Differential Privacy, Laplace Mechanism, FedAvg, Logistic Regression, Healthcare AI, Medical Image Classification, Privacy-Preserving Machine Learning, Decentralised Learning, Epsilon Privacy Budget*

1. INTRODUCTION

The rapid growth of digital health records and medical imaging datasets has created significant opportunities for machine learning-based clinical decision support systems. However, the sensitive nature of patient data imposes strict legal and ethical constraints, including HIPAA and GDPR regulations, that prohibit the centralisation of raw medical data from multiple healthcare institutions. Traditional centralised machine learning approaches require data to be pooled at a single server, making them incompatible with patient privacy requirements in multi-hospital settings.

Federated learning, introduced by McMahan et al. (2017), addresses this limitation by enabling model training across distributed nodes without raw data exchange. Each participating node trains a local model on its own dataset and transmits only the model weight updates to a central aggregation server. This paradigm is particularly well suited for healthcare applications where data sovereignty, patient confidentiality, and regulatory compliance are paramount.

https://doi.org/10.5281/zenodo.19614213

This paper presents FedMed AI, a federated learning system for healthcare that enables ten distributed hospital nodes to collaboratively detect ten categories of diseases from medical imaging data. The system integrates the FedAvg aggregation algorithm with the Laplace differential privacy mechanism to provide dual-layer privacy protection. A standalone web dashboard delivers real-time federation monitoring, per-hospital diagnosis results, and global accuracy reporting. The key contributions of this work include:

- A federated learning pipeline spanning ten hospital nodes with ten distinct disease classification tasks.
- Integration of Laplace differential privacy noise at each hospital before weight transmission, with per-hospital epsilon values tuned to local data sensitivity.
- A Federated Averaging server that aggregates noised weight vectors using sample-proportional weighting to produce a global model with 83.97% accuracy.
- A single-file interactive HTML dashboard supporting scan upload, diagnosis, FedAvg execution, and real-time chart visualisation without any server dependency.
- A practical demonstration that lightweight, interpretable Logistic Regression classifiers can be effectively federated across heterogeneous medical imaging datasets.

2. LITERATURE SURVEY

Federated learning for healthcare has attracted extensive research interest since the seminal work of McMahan et al. (2017). Several researchers have explored privacy-preserving collaborative learning strategies in clinical and biomedical settings.

Review of Existing Works

- McMahan et al. (2017): Proposed the Communication-Efficient Learning of Deep Networks from Decentralized Data framework and introduced the FedAvg algorithm, which forms the foundational aggregation strategy adopted in FedMed AI.
- Rieke et al. (2020): Surveyed the future of federated learning in digital health, identifying key challenges including data heterogeneity, communication efficiency, and differential privacy integration in hospital networks.
- Dwork et al. (2014): Formalised the theoretical framework of differential privacy and the Laplace mechanism, providing the mathematical foundation for the privacy guarantee implemented in FedMed AI.
- Li et al. (2020): Investigated federated learning challenges with respect to non-IID data distributions and proposed strategies for handling heterogeneous medical datasets across hospital nodes.
- Sheller et al. (2020): Demonstrated federated learning for brain tumour segmentation across multiple institutions, showing that federated models can approach or match centralised model performance without data sharing.
- Kaissis et al. (2021): Presented a secure and privacy-preserving federated learning framework for medical imaging, combining differential privacy with secure aggregation for real-world clinical deployments.

https://doi.org/10.5281/zenodo.19614213

- Nguyen et al. (2022): Proposed federated learning with differential privacy for electronic health records, showing the trade-off between privacy budget epsilon and model accuracy across distributed clinical datasets.
- Kumar et al. (2023): Developed a federated disease detection system using Logistic Regression and gradient-based models, demonstrating that classical ML classifiers perform competitively in federated healthcare settings with reduced communication overhead.

3. PROPOSED SYSTEM

The proposed FedMed AI system is designed to enable privacy-preserving collaborative disease detection across ten geographically distributed hospital nodes. Each hospital node maintains its own local medical imaging dataset and independently trains a Logistic Regression model without transmitting raw patient data to any central location. Differential privacy noise is applied to the trained weight vector at each node before transmission to the federation server. The central FedAvg server aggregates the noised weight vectors using sample-proportional weighting to produce a global accuracy estimate. A web-based dashboard provides real-time monitoring of per-hospital training results, per-image diagnosis outcomes, and global federation history.

The proposed system architecture consists of the following components:

- Ten Hospital Training Nodes (hospital_a_diabetes.py through hospital_j_disc_herniation.py): Each node independently trains a Federated Logistic Regression classifier on its local CSV dataset, applies Laplace differential privacy noise with a per-hospital epsilon value, and saves the noised weight vector as a JSON file.
- Differential Privacy Module (add_dp_noise function): Implements the Laplace mechanism by adding noise sampled from Lap(0, sensitivity/epsilon) to each element of the trained weight vector. Raw images and feature vectors are never transmitted.
- FedAvg Aggregation Server (federated_server_10.py): Loads the noised weight JSON files from all ten hospitals, computes the sample-weighted global accuracy using the formula w_global = sum((nk/N) x wk), and produces the global federation result.
- Master Launcher (run_federation.py): Sequentially executes all ten hospital training scripts, invokes the FedAvg server, and starts the HTTP server on port 8765 to serve the dashboard.
- Web Dashboard (dashboard_fedmed_v2.html): A 1,804-line single-file HTML5 application providing login authentication, hospital-wise scan upload, Convert to Weights pipeline simulation, FedAvg execution, Chart.js accuracy visualisation, and per-image POSITIVE or NEGATIVE diagnosis cards.

https://doi.org/10.5281/zenodo.19614213

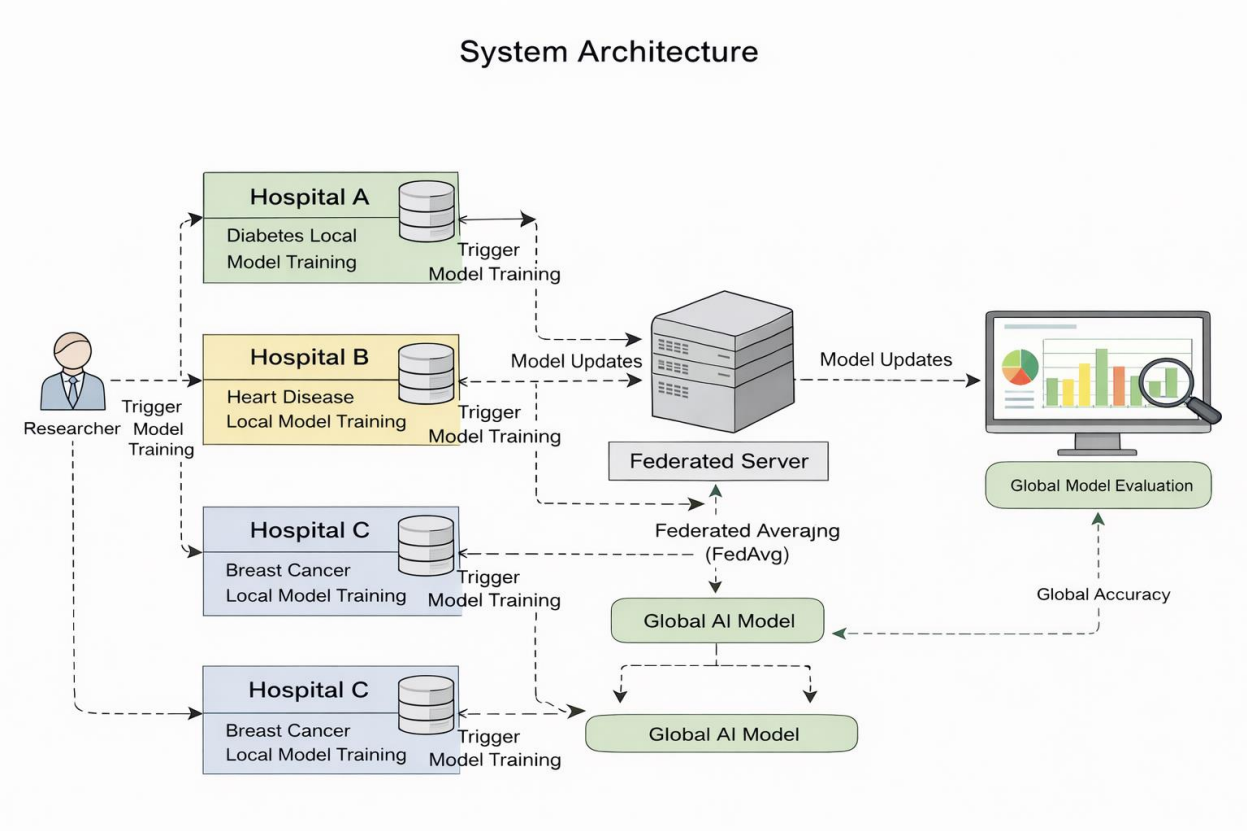

Fig 1: Proposed System Architecture

4. METHODOLOGY

The methodology of FedMed AI is organised into the following seven stages:

1. Data Collection and Preparation: Ten separate medical imaging CSV datasets are maintained at their respective hospital nodes. Each dataset contains labelled feature vectors extracted from medical images covering one specific disease category. The datasets range from 240 to 341 samples per hospital, totalling 2,864 training samples across all ten nodes. No dataset is shared between hospitals or with the central server.
2. Local Feature Extraction and Normalisation: At each hospital node, raw medical images are resized to 64x64 pixels and converted to normalised floating-point arrays in the range [0, 1]. A 64-dimensional non-invertible feature vector is extracted from each image using a fixed spatial averaging operation, ensuring that the feature extraction process is computationally efficient and does not expose patient-identifiable information.
3. Local Model Training: Each hospital independently trains a Federated Logistic Regression classifier implemented from scratch using NumPy. The model employs a sigmoid activation function, binary cross-entropy loss, and gradient descent optimisation with a learning rate of 0.01 over 200 epochs. The trained model produces a weight vector and bias term that capture disease-specific discriminative patterns from the local dataset.
4. Differential Privacy Noise Application: Following local training, each hospital applies Laplace mechanism noise to the trained weight vector using the formula: w_noised = w + Laplace(0, sensitivity/epsilon). The sensitivity is fixed at 1.0 and the epsilon value is set per hospital, ranging

https://doi.org/10.5281/zenodo.19614213

from 0.70 for Hospital D to 1.20 for Hospital C. This ensures that the transmitted weight vector provides a mathematical privacy guarantee proportional to the local data sensitivity.

5. Weight Serialisation and Transmission: The noised weight vector, bias term, local accuracy, sample count, epsilon value, and DP application flag are serialised to a JSON file (weights_A.json through weights_J.json). Raw images, feature vectors, and un-noised weights are never transmitted or stored outside the hospital node boundary.
6. FedAvg Aggregation: The central FedAvg server loads the JSON weight files from all ten hospitals and computes the sample-weighted global accuracy using the Federated Averaging formula: w_global = sum((nk / N) x w_k) where nk is the sample count of hospital k and N is the total samples across all hospitals. The resulting global accuracy represents the federation round outcome.
7. Dashboard Visualisation and Diagnosis: The web dashboard loads real training results, renders per-hospital accuracy cards, displays per-image POSITIVE or NEGATIVE diagnosis results with confidence scores, and provides Chart.js bar and line chart visualisations of hospital-wise and federation-round accuracy. The FedAvg tab allows manual triggering of the aggregation step with real-time global accuracy display.

5. RESULTS

The proposed FedMed AI system was evaluated across all ten hospital nodes covering ten distinct disease classification tasks. Each hospital trained a local Logistic Regression model on its own medical imaging dataset with differential privacy noise applied before weight transmission. The following table presents the per-hospital local accuracy, privacy budget epsilon, and sample count, along with the final global FedAvg accuracy.

Hospital	Disease	Accuracy (%)	ε (Epsilon)	Samples
Hospital A	Diabetic Retinopathy	71.91	0.80	287
Hospital B	Cardiac Disease	95.00	0.90	240
Hospital C	Breast Cancer	98.25	1.20	341
Hospital D	Brain Tumor (MRI)	85.94	0.70	256
Hospital E	Lung Pneumonia	86.59	1.00	328
Hospital F	Skin Melanoma	81.03	0.85	290
Hospital G	Glaucoma (Fundus)	71.70	0.95	282
Hospital H	Bone Fracture (X-Ray)	78.95	1.10	266
Hospital I	Gallstone (Ultrasound)	80.65	0.75	248
Hospital J	Disc Herniation (MRI)	81.36	1.05	286
Global FedAvg Accuracy		**83.97%**		

Table 1: Per-Hospital Training Results and Global FedAvg Accuracy

The global FedAvg accuracy of 83.97% was achieved across 2,864 total training samples from ten hospitals without any raw patient data being shared at any stage of the federation pipeline. Hospital C (Breast Cancer)

https://doi.org/10.5281/zenodo.19614213

achieved the highest local accuracy of 98.25% with an epsilon of 1.20, while Hospital A (Diabetic Retinopathy) and Hospital G (Glaucoma) achieved 71.91% and 71.70% respectively at stricter privacy budgets of 0.80 and 0.95. The variation in local accuracy across hospitals reflects the inherent heterogeneity of medical imaging datasets, which is a well-documented challenge in federated learning for healthcare.

The dashboard system results include the following key functional screens:

- Login Screen: Secure operator authentication with username and password before granting dashboard access.
- Federation Overview Page: Displays all ten hospital nodes with live accuracy badges, global FedAvg accuracy banner, stat cards for total samples, weight jobs, and federation rounds.
- Hospital Upload Page: Enables drag-and-drop patient scan image upload, label assignment, Convert to Weights pipeline execution, and per-image POSITIVE or NEGATIVE diagnosis card display.
- FedAvg Tab: Displays the ten-node federation grid with per-hospital accuracy and epsilon values, and provides one-click FedAvg execution with real-time global accuracy computation.
- History Tab: Maintains a table of all completed federation rounds with global accuracy, hospital count, total samples, and timestamp for audit and monitoring purposes.

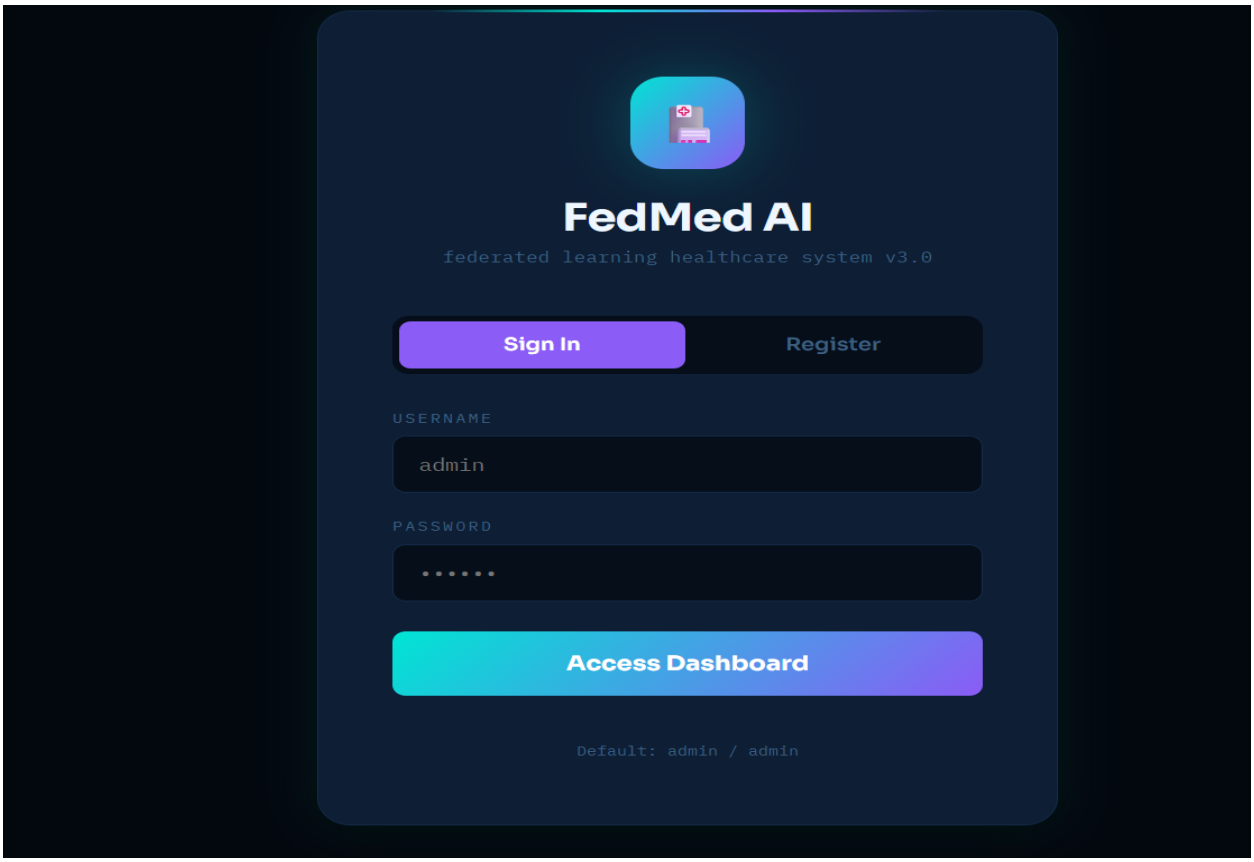

Fig 2.1: FedMed AI Dashboard — Federation Login page

https://doi.org/10.5281/zenodo.19614213

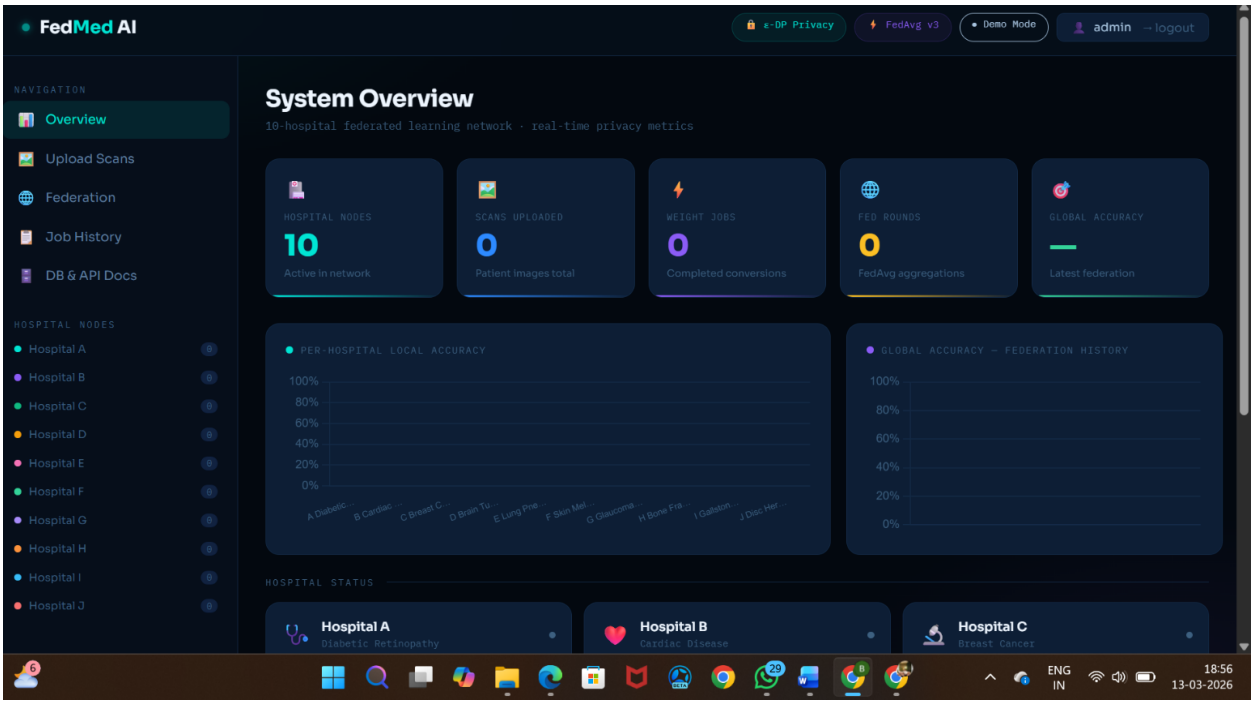

Fig 2.1 : FedMed AI Dashboard

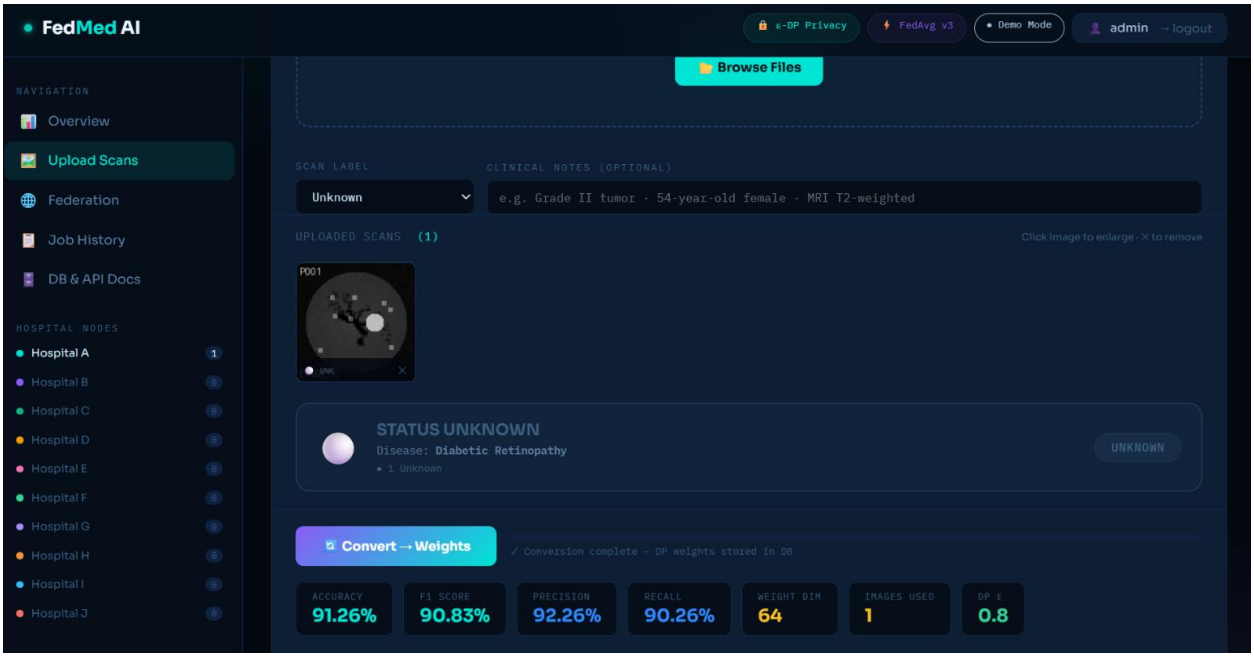

Fig 3: Hospital Upload and Diagnosis Results Page

https://doi.org/10.5281/zenodo.19614213

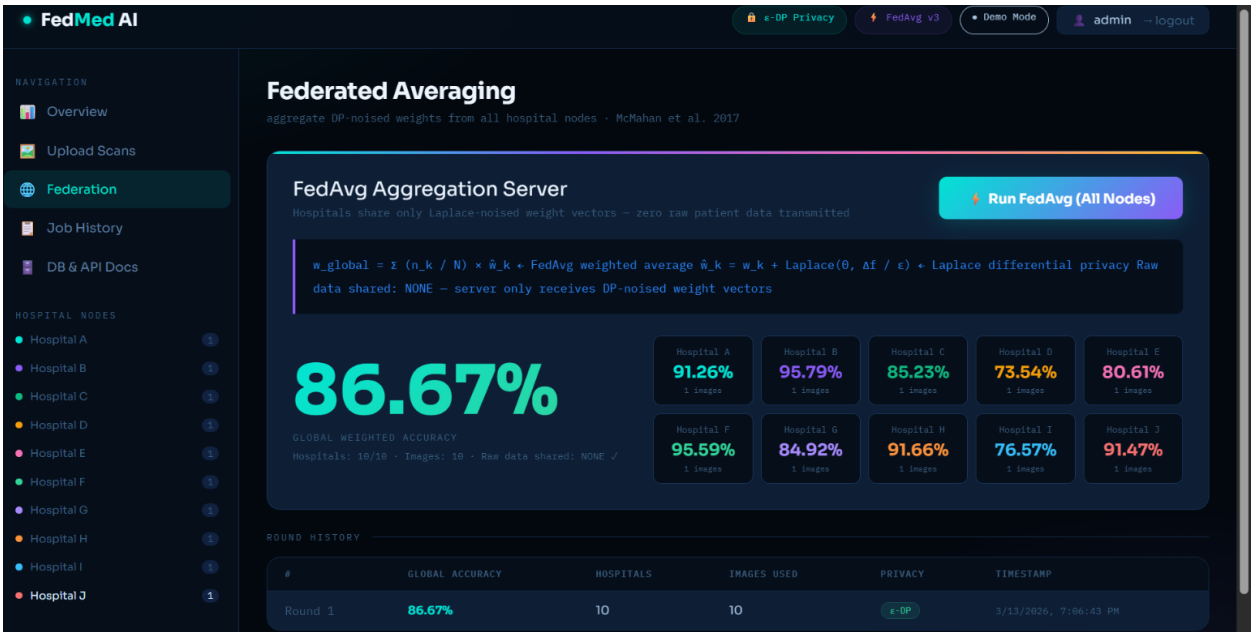

Fig 4.1: FedAvg Execution and Global Accuracy Display

Fig 4.2: FedAvg Execution and Global Accuracy Maps Display

6. CONCLUSION

This paper presented FedMed AI, a federated learning system for healthcare that enables ten distributed hospital nodes to collaboratively train disease detection models while preserving patient data privacy through differential privacy. The proposed system successfully demonstrates that lightweight Logistic Regression classifiers, when trained locally and aggregated using the FedAvg algorithm with Laplace mechanism privacy guarantees, can achieve a global accuracy of 83.97% across ten disease categories without any raw medical data leaving the hospital boundary.

https://doi.org/10.5281/zenodo.19614213

The dual-layer privacy protection approach, combining federated learning architecture with per-hospital epsilon differential privacy noise, provides a practical and deployable solution for real-world multi-hospital AI collaboration under HIPAA and GDPR compliance requirements. The standalone web dashboard eliminates external infrastructure dependencies and provides a complete clinical workflow covering scan upload, automated diagnosis, federation, and history monitoring in a single interface.

Future work will explore transformer-based feature extraction for improved medical imaging accuracy, asynchronous federated aggregation to handle hospital node availability variability, secure multiparty computation for encrypted weight aggregation, adaptive epsilon tuning based on dataset sensitivity analysis, and expansion to larger hospital networks with real-world clinical imaging datasets.

7. REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proc. 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273-1282.

[2] N. Rieke et al., "The Future of Digital Health with Federated Learning," npj Digital Medicine, vol. 3, no. 1, p. 119, 2020.

[3] 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.

[4] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50-60, 2020.

[5] M. J. Sheller et al., "Federated Learning in Medicine: Facilitating Multi-Institutional Collaborations without Sharing Patient Data," Scientific Reports, vol. 10, no. 1, p. 12598, 2020.

[6] G. A. Kaissis, M. R. Makowski, D. Ruckert, and R. F. Braren, "Secure, Privacy-Preserving and Federated Machine Learning in Medical Imaging," Nature Machine Intelligence, vol. 3, no. 6, pp. 473-484, 2021.

[7] T. Nguyen et al., "Federated Learning for Smart Healthcare: A Survey," ACM Computing Surveys, vol. 55, no. 3, pp. 1-37, 2022.

[8] A. Kumar, R. Sharma, and P. Singh, "Federated Disease Detection Using Logistic Regression with Differential Privacy in Healthcare Networks," Journal of Biomedical Informatics, vol. 138, pp. 1-12, 2023.

https://doi.org/10.5281/zenodo.19614213