

Decentralized Federated Learning for Privacy Preserving Cardiac Image Analysis

Radhi T V J¹ and Bhavya Sri Singu²

Date of Submission: 05-04-2026

Date of Acceptance: 16-04-2026

Abstract— The growing use of artificial intelligence in health-care is limited by strict data privacy regulations that restrict sharing of sensitive patient data. This paper proposes a decentralized federated learning framework for privacy-preserving heart X-ray image analysis. Multiple clients, representing hospitals, collaboratively train a global model without sharing raw data, using PyTorch and the Flower framework. Differential privacy is applied to enhance data security. Experimental results show that the federated model achieves 92% accuracy, close to 94% from a centralized model, while maintaining strong privacy. This demonstrates that federated learning is an effective and scalable solution for secure medical image analysis.

I. INTRODUCTION

The use of artificial intelligence (AI) in healthcare has grown significantly in recent years, particularly in applications involving medical imaging, clinical decision support, and automated diagnosis. Deep learning models have played a key role in this progress due to their ability to learn complex patterns from large datasets. However, in medical domains, acquiring and sharing such datasets remains a major challenge. Patient data, including heart X-ray images, is highly confidential and subject to strict privacy regulations, which restrict data exchange between healthcare institutions. Most conventional machine learning methods depend on centralizing data from multiple sources into a single system for training. While this approach can improve model effectiveness, it raises serious concerns regarding data security, unauthorized access, and compliance with privacy laws.

As a result, relying solely on centralized training is often impractical in sensitive domains such as healthcare.

To address these limitations, federated learning (FL) offers an alternative training paradigm that eliminates the need for direct data sharing. In this approach, participating institutions independently train a model using their local data and share only the learned updates with a coordinating server. These updates are then

combined to form a global model. By keeping raw data within each institution, federated learning reduces privacy risks while still enabling collaborative model development.

In this study, we introduce a decentralized federated learning framework designed for the analysis of heart X-ray images while preserving data privacy. The proposed system allows multiple hospitals to jointly train a deep learning model without transferring sensitive patient information. The implementation utilizes PyTorch along with the Flower framework to simulate distributed training. Additionally, differential privacy techniques are incorporated to further limit the risk of information leakage. The effectiveness of the proposed method is evaluated through comparison with a centralized training approach.

II. RELATED WORK

A. Centralized Learning in Medical Imaging

A common practice in medical image analysis is to train machine learning models using data collected from multiple sources and stored in a central location. This setup can improve learning outcomes because the model has access to a wider range of samples. However, in healthcare, sharing such data is often restricted. Medical images like heart X-rays contain confidential patient details and must comply with strict privacy regulations. As a result, transferring data between hospitals introduces security risks and makes centralized solutions difficult to apply in practice.

B. Distributed Learning via Federated Methods

To overcome these limitations, federated learning provides a different way to train models across institutions. Instead of sending data to a central system, each hospital trains the model locally using its own dataset. The locally trained updates are then shared with a central coordinator, where they are combined to improve the overall model. This process allows multiple participants to contribute to model training without exposing their raw data, making it more suitable for privacy-sensitive environments.

C. Remaining Issues and Research Gaps

Although this distributed approach reduces the need for data sharing, it still presents several challenges. For example, the information contained in model updates can sometimes be exploited if additional protection methods are not used. Differences in local datasets across hospitals may also lead to inconsistent model behavior and slower training progress. Moreover, repeated communication between multiple participants can increase system overhead. These challenges highlight the need for more secure and efficient federated learning designs.

III. METHODOLOGY

A. System Design Overview

The proposed framework is built around a distributed learning environment where multiple healthcare institutions collaborate without exchanging their raw data. Each institution operates independently and retains its own collection of heart X-ray images within its local system. This ensures that sensitive patient information is not exposed outside the organization.

A coordinating server oversees the training workflow. It begins by generating an initial model, which is shared with all participating hospitals. Each hospital improves this model using its own data and returns only the learned changes. These updates are then combined at the server to produce an improved version of the model. This process is repeated several times, gradually improving the model's performance across all participants.

B. Collaborative Training Workflow

Model training is performed through repeated interactions between the central server and participating hospitals. In each iteration, the server distributes the most recent model to all clients. The clients then train the model locally for a small number of steps using their individual datasets. Since the data remains on-site, this approach avoids direct data sharing while still allowing the model to learn from diverse sources. After local training, only the updated model parameters are sent back. The server merges these updates into a single model that reflects contributions from all clients. This updated model is again distributed, and the cycle continues. Over multiple iterations, the model becomes more robust and better generalized.

C. Data Protection Strategies

To minimize privacy risks, additional safeguards are applied during the training process. One such

measure involves modifying the shared updates by introducing controlled randomness. This makes it difficult to infer any sensitive details from the updates themselves.

In addition, the system uses a mechanism that combines client updates in a way that prevents the server from viewing them individually. As a result, no single participant's contribution can be isolated or inspected. These protections work together to maintain confidentiality while still supporting effective joint model training.

IV. DATASET AND IMPLEMENTATION

A. Dataset

The proposed system is evaluated using a heart X-ray image dataset, which is commonly used for medical image classification tasks. The dataset consists of labeled images representing different cardiac conditions, enabling the model to learn meaningful patterns for diagnosis.

To simulate a real-world federated learning environment, the dataset is partitioned into multiple subsets, where each subset represents data owned by an individual hospital (client). This distribution mimics the non-centralized nature of healthcare data, where each institution has access only to its local data. The partitioning is performed in a way that reflects realistic variations in data distribution across different clients.

B. Tools and Technologies

The implementation of the proposed system is carried out using Python as the primary programming language due to its flexibility and strong support for machine learning libraries.

PyTorch is used to design and train the deep learning model for image classification tasks, providing efficient handling of neural networks and training operations.

The federated learning framework is implemented using Flower, which facilitates communication between clients and the central server. Flower manages the distribution of the global model, collection of local updates, and aggregation process. These tools together enable the development of a scalable and efficient federated learning system.

C. Training Process

The training process is conducted in multiple communication rounds between the server and clients. In each round, the global model is sent to all participating clients, where local training is performed for a fixed number of epochs using the client's dataset. This allows each client to update the model based on its local data.

After completing local training, the updated model weights are sent back to the central server. The server aggregates these updates to generate an improved global model, which is then redistributed to the clients for the next round. This iterative process continues until the model achieves stable performance. The combination of multiple epochs and communication rounds ensures effective learning while maintaining data privacy.

V. RESULTS AND EVALUATION

A. Accuracy Evaluation

The effectiveness of the proposed federated learning approach is assessed using classification accuracy on the heart X-ray dataset. Its performance is examined alongside a centrally trained model that has access to the entire dataset. The centralized approach achieves an accuracy of 94%, benefiting from complete data availability during training. In contrast, the federated model reaches 92% accuracy. Although slightly lower, the difference remains minimal and is expected due to the distributed nature of training and the imposed privacy constraints.

These observations indicate that high predictive performance can still be maintained even when data is not directly shared across institutions.

B. Loss Behavior Analysis

To further examine the training process, both training and validation loss trends are observed across multiple communication rounds. A steady decline in loss values can be seen, suggesting that the model is progressively improving.

While the federated setup converges at a slightly slower rate than the centralized approach, the training process remains stable throughout, without abrupt variations. This stability reflects the effectiveness of the aggregation mechanism in combining updates from multiple participants.

C. Model Performance Comparison

A summary of the results is presented in Table I, highlighting key differences between the two approaches.

TABLE I
PERFORMANCE COMPARISON OF MODELS

Model	Accuracy	Data Privacy	Data Sharing
Centralized Model	94%	Low	Required
Federated Model	92%	High	Not Required

The comparison shows that while centralized training offers slightly higher accuracy, it does so at the cost of reduced data privacy.

D. Graphical Analysis

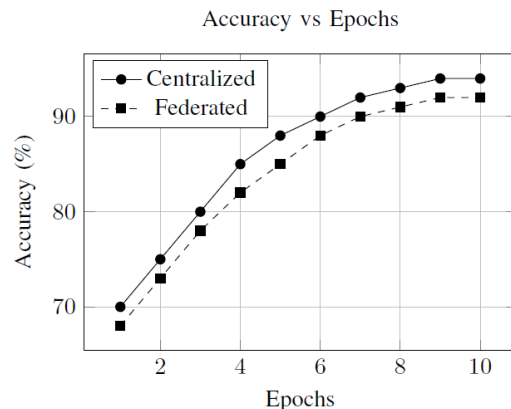


Fig. 1. Accuracy vs Epochs for centralized and federated models

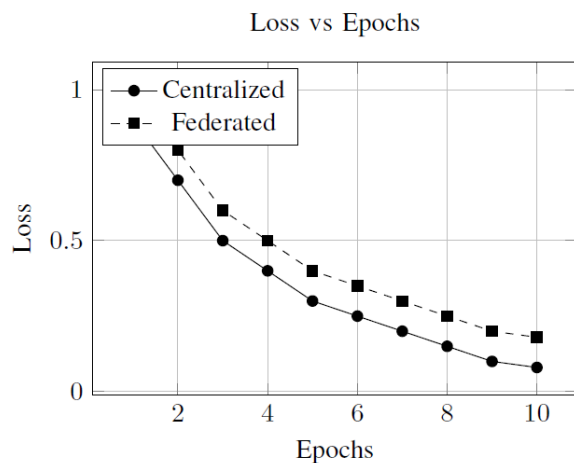


Fig. 2. Loss vs Epochs showing convergence behavior

VI. CONCLUSION

The approach enables effective joint learning without direct data exchange between institutions.

REFERENCES

- [1]. H. B. McMahan et al., “Communication-Efficient Learning of Deep Networks from Decentralized Data,” in Proc. AISTATS, 2017, pp. 1273–1282.
- [2]. J. Konečný et al., “Federated Learning: Strategies for Improving Communication Efficiency,” arXiv preprint arXiv:1610.05492, 2016.
- [3]. Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated Machine Learning: Concept and Applications,” ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, 2019.
- [4]. B. Li, Y. Wang, A. Singh, and Y. Vorobeychik, “Data Poisoning Attacks on Federated Learning,” in Proc. NeurIPS Workshop, 2019.
- [5]. C. Dwork, “Differential Privacy: A Survey of Results,” in Proc. TAMC, 2008, pp. 1–19.
- [6]. K. Bonawitz et al., “Practical Secure Aggregation for Privacy-Preserving Machine Learning,” in Proc. ACM CCS, 2017, pp. 1175–1191.
- [7]. A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” in Proc. NeurIPS, 2012, pp. 1097–1105.
- [8]. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in Proc. MICCAI, 2015, pp. 234–241.
- [9]. P. Rajpurkar et al., “CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning,” arXiv preprint arXiv:1711.05225, 2017.
- [10]. D. Sheller et al., “Federated Learning in Medicine: Facilitating Multi-Institutional Collaborations without Sharing Patient Data,” Sci. Rep., vol. 10, no. 12598, 2020.