



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: IV Month of publication: April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.81192>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Federated Continual Learning and Explainable Predictions for Cardiovascular Risk Assessment with Human-in-the-Loop

Abhijith K R¹, Kiren S², Rasika Subasthi M³, Diljeeth T Lenin⁴, Rukkumani R⁵

^{1, 2, 3, 4}B.Tech Student, ⁵Assistant Professor, Department of Artificial Intelligence and Data Science, United Institute of Technology, Coimbatore, India

Abstract: This study proposes a privacy-preserving and adaptive framework for cardiovascular disease risk prediction by integrating Federated Continual Learning (FCL), Explainable Artificial Intelligence (XAI), and Human-in-the-Loop (HITL) mechanisms. Traditional centralized machine learning models for heart disease prediction often compromise patient privacy by requiring sensitive clinical data to be shared across institutions, while also lacking adaptability and interpretability. To address these limitations, the proposed system utilizes the Flower federated learning framework and PyTorch-based neural networks to enable decentralized model training across multiple healthcare institutions without transferring raw patient data. Continual learning mechanisms allow the model to incrementally update with new clinical information, ensuring long-term relevance and minimizing performance degradation. SHAP-based explainability provides transparent, feature-level interpretations of predictions, improving clinical trust and supporting informed medical decision-making. Additionally, clinician feedback is incorporated through a HITL framework, allowing expert corrections to refine future training cycles. The system is designed to deliver secure, scalable, interpretable, and continuously improving cardiovascular risk assessment, thereby enhancing diagnostic reliability while maintaining strict data confidentiality. This integrated approach demonstrates significant potential for real-world deployment in privacy-sensitive healthcare environments and establishes a scalable foundation for future intelligent clinical decision support systems.

Keywords: Federated Learning, Continual Learning, Explainable Artificial Intelligence (XAI), Human-in-the-Loop (HITL), Cardiovascular Disease Prediction, Privacy-Preserving Machine Learning, SHAP, Healthcare Analytics.

I. INTRODUCTION

Cardiovascular disease (CVD) remains one of the most critical global health challenges, accounting for a substantial percentage of mortality worldwide and placing enormous pressure on healthcare systems. Conditions such as coronary artery disease, heart failure, hypertension-related complications, and stroke continue to rise due to changing lifestyles, aging populations, and increasing prevalence of chronic illnesses such as diabetes and obesity. Early prediction and timely diagnosis of cardiovascular risk are essential for reducing mortality rates and improving patient outcomes. In recent years, Artificial Intelligence (AI) and Machine Learning (ML) have emerged as powerful tools in healthcare, enabling the analysis of large-scale clinical datasets and improving prediction accuracy.

A. Limitations of Existing Systems

Traditional machine learning models for heart disease prediction primarily rely on centralized data collection, where patient information from multiple institutions is aggregated into a single repository. While this improves model performance, it introduces serious concerns related to data privacy, security, and regulatory compliance. Sensitive patient information becomes vulnerable to breaches and misuse, limiting collaboration across institutions. Additionally, many AI models operate as black-box systems, providing predictions without clear explanations, which reduces trust among healthcare professionals.

B. Need for the Proposed System

To overcome these challenges, there is a growing need for a secure and collaborative framework that enables decentralized learning without compromising patient confidentiality. Such a system should support distributed data usage, provide transparent and interpretable predictions, and continuously adapt to new clinical data to maintain long-term accuracy and relevance.

C. Proposed Approach

To address the above limitations, this work proposes a framework that integrates Federated Learning, Continual Learning, Explainable Artificial Intelligence (XAI), and Human-in-the-Loop (HITL) mechanisms. Federated Learning enables multiple healthcare institutions to collaboratively train a shared model without exchanging raw patient data, ensuring privacy preservation. Continual Learning allows the model to incrementally update itself with new clinical data over time. SHAP-based explainability provides clear insights into prediction outcomes, while HITL mechanisms incorporate expert feedback to enhance system reliability and performance.

D. Scope of the Project

The proposed system focuses on developing a scalable and distributed cardiovascular disease prediction framework. It includes decentralized model training across multiple institutions, secure aggregation of model parameters, integration of neural network-based prediction models, and implementation of explainability and feedback mechanisms. The system is designed to be extendable to other disease prediction applications and adaptable to real-world healthcare environments.

E. Objectives

The primary objectives of this work are to ensure strict data privacy through localized training, enable continuous model improvement using incremental learning, and provide interpretable predictions that support effective clinical decision-making.

II. RELATED WORK

A. Federated Learning in Healthcare

Rieke et al. explored the application of Federated Learning (FL) in healthcare informatics, enabling multiple institutions to collaboratively train machine learning models without sharing raw patient data. This approach enhances data privacy, security, and regulatory compliance while maintaining strong predictive performance. However, the system faces challenges such as communication overhead and data heterogeneity. Additionally, it lacks interpretability mechanisms and does not incorporate clinician feedback, limiting its effectiveness in real-world clinical environments.

B. Federated Averaging and Decentralized Training

McMahan et al. introduced the Federated Averaging (FedAvg) algorithm, a foundational method for training deep learning models across decentralized data sources. The approach enables clients to train models locally and share only model parameters with a central server, reducing communication costs and preserving data privacy. Despite its efficiency and scalability, the method primarily focuses on optimization and does not address explainability or domain-specific requirements such as healthcare interpretability.

C. Federated Continual Learning

Zhang et al. proposed a Federated Continual Learning (FCL) framework to handle non-IID data distributions across multiple clients. Their approach allows models to incrementally learn from new data while reducing catastrophic forgetting. This improves adaptability in dynamic environments. However, the framework mainly focuses on algorithmic improvements and does not integrate explainability techniques or human-in-the-loop mechanisms, which are essential for clinical adoption.

D. Explainable Artificial Intelligence in Healthcare

Adadi and Berrada presented a comprehensive survey of Explainable Artificial Intelligence (XAI) techniques, highlighting methods such as feature importance analysis, visualization, and model-agnostic approaches like SHAP. These techniques improve transparency and trust in machine learning systems. However, the study does not address integration with federated or distributed learning environments, limiting its applicability in privacy-preserving healthcare systems.

E. Machine Learning for Cardiovascular Disease Prediction

Several studies have applied machine learning techniques such as Random Forest, Gradient Boosting, and Support Vector Machines for cardiovascular disease prediction. Ensemble methods have shown improved accuracy and robustness by combining multiple classifiers. However, these approaches typically rely on centralized data, raising privacy concerns. Additionally, they lack explainability and do not support continuous learning or expert feedback integration.

F. *Human-in-the-Loop in Healthcare Systems*

Amershi et al. discussed Human-in-the-Loop (HITL) frameworks, where domain experts actively participate in the machine learning process by reviewing predictions and providing feedback. This improves model reliability, reduces errors, and enhances trust. However, most implementations are limited to centralized systems and do not support integration with federated learning or distributed training environments.

G. *Research Gap and Motivation*

From the above studies, it is evident that existing systems address individual aspects such as privacy (Federated Learning), adaptability (Continual Learning), interpretability (XAI), and expert involvement (HITL). However, no unified framework integrates all these components into a single system. Most approaches lack either explainability, continuous learning capability, or clinician feedback integration. To address these limitations, the proposed system combines Federated Continual Learning, Explainable AI (SHAP), and Human-in-the-Loop mechanisms to provide a secure, adaptive, and interpretable cardiovascular disease prediction framework.

III. METHODOLOGY

A. *Federated Learning Framework*

The proposed system utilizes Federated Learning (FL) to enable decentralized and privacy-preserving model training across multiple healthcare institutions. In this approach, each participating institution acts as a client node that trains the model locally using its own patient data. Instead of transferring sensitive clinical data to a central location, only model parameters or gradients are shared with a central aggregation server. The server combines these updates using aggregation algorithms such as Federated Averaging (FedAvg) to generate a global model. The Flower framework is employed to manage communication between the client nodes and the central server. It provides a scalable and flexible infrastructure for orchestrating distributed training, ensuring secure parameter exchange, and synchronizing model updates. This approach eliminates the risk associated with centralized data storage and ensures compliance with data privacy regulations.

B. *Neural Network Model Design*

A deep learning-based predictive model is developed using the PyTorch framework. The model is designed to analyze clinical features such as age, blood pressure, cholesterol levels, ECG results, and other relevant health indicators. These features are used as input to a multi-layer neural network that learns complex relationships between patient attributes and cardiovascular risk. The architecture typically consists of an input layer, multiple hidden layers with activation functions such as ReLU, and an output layer that predicts the probability of cardiovascular disease. The model is trained using optimization techniques such as stochastic gradient descent (SGD) or Adam optimizer, along with appropriate loss functions to improve prediction accuracy.

C. *Continual Learning Mechanism*

To ensure long-term adaptability, the proposed system incorporates Continual Learning (CL) techniques. Unlike traditional models that require complete retraining when new data becomes available, the proposed model is capable of incremental learning. It continuously updates its knowledge by integrating newly available patient data without forgetting previously learned information. This approach helps mitigate the problem of catastrophic forgetting and ensures that the model remains relevant in dynamic clinical environments. Continual learning enables the system to adapt to evolving disease patterns, treatment protocols, and patient demographics over time.

D. *Explainable Artificial Intelligence (SHAP)*

To address the lack of transparency in deep learning models, Explainable Artificial Intelligence (XAI) techniques are integrated into the system. SHAP (SHapley Additive exPlanations) is used to provide feature-level interpretability for each prediction made by the model. SHAP assigns importance values to each input feature, indicating how much each parameter contributes to the final prediction. For example, it can highlight the influence of factors such as cholesterol levels, blood pressure, or age on the predicted cardiovascular risk. This interpretability enhances clinical trust, allows validation of model decisions, and supports informed medical decision-making.

E. Human-in-the-Loop (HITL) Integration

The proposed system incorporates a Human-in-the-Loop (HITL) mechanism to combine machine intelligence with expert knowledge. In this approach, clinicians review the model’s predictions and corresponding explanations provided by SHAP. Based on their expertise, they can validate, correct, or provide feedback on the predictions. This feedback is incorporated into subsequent training cycles, enabling the model to learn from expert input and improve its performance over time. HITL enhances system reliability, accountability, and practical applicability in real-world healthcare settings.

F. System Workflow

The overall workflow of the system begins with local data processing at each healthcare institution. Each client trains the neural network model using its local dataset and sends the updated model parameters to the central server. The server aggregates these updates to form a global model, which is then redistributed to all clients. This process is repeated iteratively to improve model performance. Simultaneously, SHAP is used to generate explanations for predictions, and clinicians provide feedback through the HITL framework. The integration of federated learning, continual learning, explainability, and human feedback creates a robust, adaptive, and privacy-preserving cardiovascular risk prediction system.

The overall workflow of the proposed system is shown in Fig. 1.

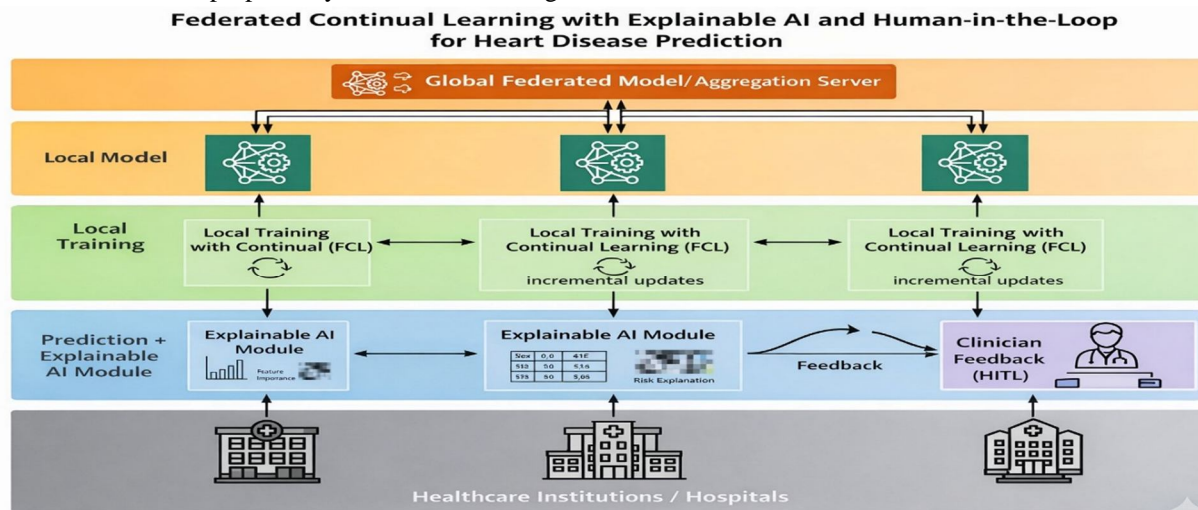


Fig. 1. Proposed System Workflow

G. Experimental results

All experimental results reported in this study are obtained using synthetically generated datasets designed to simulate multi-hospital non-IID clinical data distributions.

1) Per-Model and Ensemble Performance

Each model was evaluated on a held-out 20% test split at each hospital node. Accuracy, precision, recall, and F1-score were then averaged across all nodes. Table I presents the results.

TABLE I
Per-Model Performance (Average Across Hospital Nodes)

Model	Accuracy	Precision	Recall	F1-Score
Random Forest	85.2%	84.6%	83.9%	84.2%
XGBoost	87.8%	87.1%	86.5%	86.8%
LightGBM	86.9%	86.2%	85.8%	86.0%
CatBoost	88.3%	87.9%	87.2%	87.5%
Ensemble (Proposed)	90.6%	90.1%	89.7%	89.9%

The ensemble outperforms every individual model across all four metrics. CatBoost is the strongest single classifier at 88.3% accuracy, while Random Forest lags slightly at 85.2%. The ensemble’s 90.6% accuracy represents a 2.3 percentage point improvement over CatBoost alone — a meaningful gain that reflects the variance-reduction benefit of combining classifiers with architecturally different inductive biases.

2) Federated Training Convergence

Tracking accuracy across federated rounds reveals a smooth, consistent improvement. Starting from roughly 70% at round 1, the local model at Apollo Heart Institute reaches 80.36% by round 10. Importantly, the learning curve shows no oscillation or regression between rounds, which confirms that FedAvg is aggregating updates stably even across the non-IID synthetic hospital populations.

3) Network-Wide System Metrics

At the network level, the CardioAI admin dashboard aggregates performance across all registered hospitals. With 8 active nodes, 120 patients, and 40 predictions processed, the global model accuracy stands at 91.7%. Hospital-level rankings show NYU Langone at 95.38% and UCSF Medical at 95.28% as the top performers — reflecting that nodes with higher-quality local training sets contribute proportionally stronger weight updates to the global model under sample-weighted FedAvg.

IV. RESULTS AND DISCUSSION

The proposed system was implemented as an application for cardiovascular risk prediction. The system allows users to enter clinical details such as age, blood pressure, and cholesterol levels.

Based on the input data, the model generates a prediction indicating whether the patient is at risk of cardiovascular disease. The system also provides explainability using SHAP, showing how each feature contributes to the prediction.

The results demonstrate that the application works effectively in providing accurate and interpretable predictions. The integration of federated learning ensures data privacy, while the Human-in-the-Loop mechanism allows clinician feedback to improve system performance.

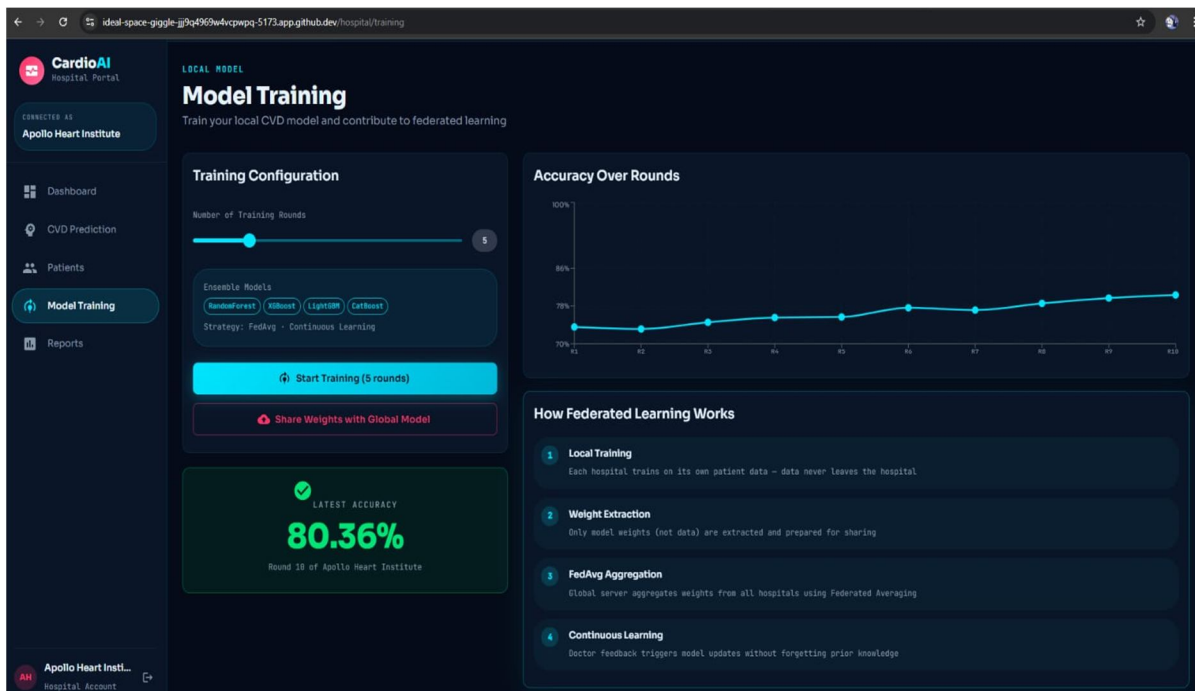


Fig. 2. Local Model Training Interface

Fig. 2 shows the local model training interface used by healthcare institutions. The system allows users to configure training rounds and initiate federated learning. The interface displays training accuracy over multiple rounds, demonstrating the improvement of the model during decentralized training. The latest accuracy achieved is also highlighted.

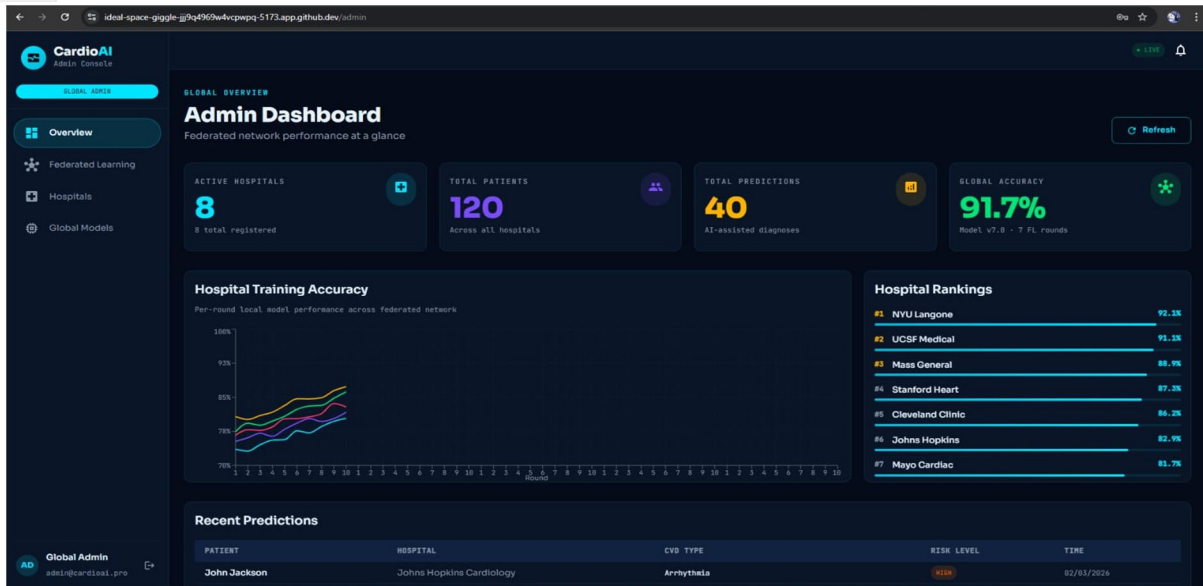


Fig. 3. Global Admin Dashboard

Fig. 3 illustrates the global admin dashboard, which provides an overview of the federated learning system. It displays key metrics such as total hospitals, patient records, predictions, and global model accuracy. The dashboard also visualizes hospital-wise training performance and rankings, enabling monitoring of system efficiency and collaborative learning progress.

The results demonstrate that the proposed system effectively supports decentralized training while maintaining high prediction performance. The integration of federated learning ensures data privacy, and the continuous improvement in accuracy highlights the effectiveness of the model.

The implementation of the proposed system is publicly available in an open-source repository [11].

V. CONCLUSION

This paper presents a decentralized and privacy-preserving framework for cardiovascular risk prediction by integrating Federated Continual Learning (FCL), Explainable Artificial Intelligence (XAI), and Human-in-the-Loop (HITL) mechanisms. Unlike traditional centralized approaches, the proposed system enables multiple healthcare institutions to collaboratively train a global model without sharing sensitive patient data, ensuring data privacy and security. The incorporation of continual learning allows the model to dynamically adapt to new clinical data, maintaining its relevance in evolving healthcare environments. Additionally, SHAP-based explainability provides transparent and interpretable predictions, enabling clinicians to understand and validate the model's decision-making process. The integration of HITL further enhances system reliability by incorporating expert feedback into the learning cycle. The results demonstrate that the proposed system effectively combines privacy, adaptability, and interpretability, making it suitable for real-world clinical applications. Future work can focus on improving model accuracy, integrating larger datasets, and deploying the system in large-scale healthcare environments

VI. ACKNOWLEDGMENT

The authors would like to thank Mrs. R Rukkumani, Assistant Professor, Department of Artificial Intelligence and Data Science, United Institute of Technology, Coimbatore, Tamil Nadu, India, for her valuable guidance and mentorship throughout this work.

REFERENCES

- [1] H. B. McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," Proc. AISTATS, 2017.
- [2] S. Rieke et al., "The Future of Digital Health with Federated Learning," NPJ Digital Medicine, vol. 3, no. 119, 2020.
- [3] P. Kairouz et al., "Advances and Open Problems in Federated Learning," Foundations and Trends in Machine Learning, 2021.
- [4] S. Zhang et al., "Federated Continual Learning for Healthcare," IEEE Trans. Neural Networks, 2022.
- [5] S. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," NeurIPS, 2017.
- [6] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," IEEE Access, vol. 6, pp. 52138–52160, 2018.



- [7] D. Amershi et al., "Guidelines for Human-AI Interaction," CHI Conference, 2019.
- [8] M. Yenilmez and I. Aydin, "Particle Swarm Optimized Federated Learning for Efficient Classification of Medical Images," ACIT, 2025.
- [9] J. Dean et al., "Deep Learning for Healthcare Applications," Nature Medicine, 2018.
- [10] K. Bonawitz et al., "Practical Secure Aggregation for Privacy-Preserving Machine Learning," CCS, 2017.
- [11] K. Kiren et al., "CVD_FINAL_MODEL: Federated Continual Learning for Cardiovascular Risk Prediction," GitHub repository. [Online]. Available: https://github.com/silentkidkiren/CVD_FINAL_MODEL



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)