Journal of Carcinogenesis
Covering Genesis to Therapeutics

Journal of Carcinogenesis, ISSN: 1477-3163
2025; Vol 24: Issue 3, https://carcinogenesis.com/

Decentralized Intelligence for Cardiac Health: A Federated Learning Approach to Privacy-Conscious Disease Prediction

Pyla Uma, P.V.Lakshmi

Department of Computer Science and Engineering, School of Technology, Gandhi Institute of Technology and Management (GITAM University), Visakhapatnam, India

Department of Computer Science and Engineering, School of Technology, Gandhi Institute of Technology and Management (GITAM University), Visakhapatnam, India

upyla@gitam.in , vpanga@gitam.edu

Corresponding author: Pyla Uma, pylauma2017@gmail.com

Clinical trial number: not applicable.

ABSTRACT

Cardiovascular disease (CVD) remains the leading cause of mortality worldwide, with early detection being vital for reducing life-threatening complications and improving patient prognosis. Traditional machine learning (ML) models for CVD prediction require centralized access to large-scale clinical data, but such access is often restricted due to privacy regulations and institutional barriers. To overcome these challenges, this study proposes a privacy-preserving federated learning (FL) framework that enables multiple healthcare institutions to collaboratively train predictive models without sharing sensitive patient data. The framework employs a multilayer neural network trained using the Federated Averaging (FedAvg) strategy, combined with data preprocessing techniques including standardization and class balancing via Synthetic Minority Oversampling Technique (SMOTE). Experimental evaluations using a real-world CVD risk dataset demonstrate that the proposed federated model achieves a classification accuracy of 94.0%, an F1-score of 93.5%, and outperforms centralized and local baselines. Communication overhead is kept minimal, averaging ~75 KB per round, and the model shows strong convergence across heterogeneous client data. An ablation study further confirms the critical role of SMOTE and local training configurations in model performance.

Keywords: *Cardiovascular Disease Detection; Federated Learning; Privacy Preservation; Deep Neural Networks; Class Imbalance Handling; Distributed Healthcare AI.*

How to Cite: Pyla Uma, P.V.Lakshmi, (2025) Decentralized Intelligence for Cardiac Health: A Federated Learning Approach to Privacy-Conscious Disease Prediction, *Journal of Carcinogenesis*, *Vol.24*, *No.3*, 273-291.

1. INTRODUCTION

Cardiovascular diseases (CVDs) remain the leading cause of death globally, accounting for approximately 17.9 million deaths annually, which represents about 32% of all global deaths according to the World Health Organization [1]. Early detection of CVDs is crucial as it significantly improves the prognosis, reduces treatment costs, and lowers mortality rates through timely interventions. Traditional diagnostic procedures such as electrocardiograms (ECG), echocardiography, and computed tomography (CT) scans have been widely used for CVD detection. However, the integration of artificial intelligence (AI), particularly machine learning (ML) techniques, has revolutionized disease diagnosis by enabling rapid, accurate, and scalable prediction models that can assist clinicians in identifying at-risk individuals early in the disease trajectory.

ML algorithms have demonstrated high performance in classifying CVDs using structured and unstructured clinical data, including vital signs, laboratory results, and medical imaging [2]. Deep learning (DL) models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), have further enhanced this field, especially in image and signal-based diagnostics. For example, studies leveraging ECG data for arrhythmia detection using deep neural networks have achieved accuracies exceeding 90% [3], indicating the viability of AI-assisted cardiac care. The deployment of these models in clinical settings promises to reduce diagnostic errors and improve overall healthcare efficiency.

Despite these advancements, several challenges hinder the full-scale deployment of ML in healthcare diagnostics. Concerns

about data security and privacy rank highest among them. Strict legal frameworks like the General Data Protection Regulation (GDPR) in Europe and the Health Insurance Portability and Accountability Act (HIPAA) in the US safeguard sensitive medical data. These regulations restrict the free sharing of patient data across institutions, thereby impeding the development of robust, generalized AI models that require large, diverse datasets for training [4].

Centralized data approaches, wherein data from multiple institutions are pooled into a single location for model training, are particularly susceptible to breaches and raise significant privacy concerns. Moreover, such approaches often fail to represent the heterogeneity of healthcare systems, including differences in patient demographics, diagnostic equipment, and medical practices. These discrepancies can lead to biased models with poor generalizability across different healthcare environments. Additionally, the high computational and communication costs associated with transmitting large volumes of medical data further constrain centralized strategies, making them inefficient and potentially impractical for real-world deployment.

Federated Learning (FL) emerges as a good way to address the above limitations. FL is a decentralized machine learning paradigm that enables multiple institutions to collaboratively train a shared global model without transferring raw patient data outside their local systems [5]. Instead, model updates such as gradients or weights are exchanged and aggregated in a privacy-preserving manner, thus keeping patient data secure and compliant with regulatory standards. In healthcare, FL has demonstrated remarkable potential in balancing model performance with data privacy. For instance, [6] showcased that FL could achieve comparable accuracy to centralized models in brain tumor segmentation while maintaining complete data decentralization. Techniques such as differential privacy, secure multi-party computation, and homomorphic encryption can be integrated within FL frameworks to further enhance security and resilience against adversarial attacks [7]. Moreover, the heterogeneity of healthcare data, which is a challenge for centralized systems, can be better managed in FL through personalization strategies and robust aggregation mechanisms.

Motivated by the pressing need for early and privacy-preserving detection of CVD, this study contributes a technically robust and scalable solution through federated learning. First, we introduce a novel FL-based diagnostic framework tailored for CVD prediction, facilitating cooperative model training across several dispersed healthcare facilities. By ensuring that raw patient data never leaves local nodes, the framework adheres to stringent privacy regulations while maintaining data sovereignty and institutional autonomy.

Second, we implement and evaluate a multilayer neural network trained under the Federated Averaging (FedAvg) strategy, incorporating essential preprocessing techniques such as feature standardization and class rebalancing via SMOTE. Our model is rigorously compared with centralized and local baselines using a real-world CVD dataset. The federated model consistently outperforms these baselines, achieving a classification accuracy of 94.0% and an F1-score of 0.935, while demonstrating stable convergence across heterogeneous data environments.

Third, we provide an in-depth analysis of the system's communication efficiency and privacy trade-offs. Detailed metrics, including communication overhead (~75 KB per round) and encryption-based secure weight updates, validate the scalability and practicality of the proposed system. An ablation study further quantifies the impact of key variables such as class rebalancing and local training epochs on model performance. Together, these contributions position our FL framework as a promising, privacy-compliant solution for real-world deployment in CVD risk assessment and clinical decision support. Despite the growing success of AI in cardiovascular disease detection, the development and deployment of high-performance ML models are severely constrained by data privacy regulations and the limitations of centralized data storage approaches. Existing methods often fail to balance the dual needs of data utility and security, particularly in sensitive domains such as healthcare. There is a pressing need for privacy-preserving learning paradigms that can accommodate distributed data environments while delivering high predictive accuracy.

The aim of this study is to develop and validate a federated learning framework for early and accurate detection of cardiovascular diseases. The framework is designed to preserve patient data privacy, enhance communication efficiency, and ensure high model performance in a distributed healthcare setting. Through extensive experiments and analysis, we aim to demonstrate that FL can serve as a viable alternative to centralized AI approaches, offering a scalable, secure, and effective solution for collaborative CVD diagnostics.

Contributions made in this study:

- Presented a privacy-preserving federated learning (FL) framework for early cardiovascular disease (CVD) detection, enabling collaborative model training across decentralized clinical datasets without exposing sensitive patient information.
- Integrated a multilayer neural network with the Federated Averaging (FedAvg) strategy, achieving a robust balance between model accuracy and communication efficiency in non-IID and resource-constrained healthcare environments.

- Applied class balancing techniques during client-side preprocessing to address class imbalance issues, significantly improving model generalization and increasing accuracy by up to 6%.
- Conducted a comprehensive performance evaluation, where the proposed FL model achieved 94.0% accuracy and 0.935 F1-score, outperforming centralized and local baselines.
- Validated communication efficiency and scalability, showing stable overhead (~75 KB per round) and convergence within 50 communication rounds.

2. RELATED WORK

Federated learning in healthcare

In healthcare AI, federated learning has become a significant technique that allows for cooperative model training across several organisations without centralising private patient data. This paradigm is particularly valuable in healthcare, where privacy, data heterogeneity, and regulatory compliance pose significant barriers to data sharing. This subsection explores recent advancements in FL applications within clinical settings, focusing on its potential for secure, scalable, and privacy-aware medical AI systems, the summary of key related studies are presented in table 1.

Gaber et al. [8] developed FedCVD, a federated learning-based framework for cardiovascular disease prediction using logistic regression and SVM models. By leveraging federated learning, they enabled collaborative training across decentralized patient datasets while preserving data privacy. To address class imbalance in the data, they applied three resampling techniques: Random Over Sampling, Random Under Sampling, and SMOTE. Their experiments on the 10-year risk of coronary heart disease Kaggle dataset showed that the federated models performed competitively or even better than centralized models. Notably, the federated logistic regression with SMOTE achieved an AUC of 0.7048, while the federated SVM with Random Under Sampling reached an AUC of 0.7340, outperforming the centralized SVM model (AUC of 0.6962). Antunes et al., [9] conducted a systematic literature review to explore the current state of federated learning applied to electronic health records (EHR) for healthcare applications. Recognizing the growing use of machine learning with sensitive medical data, they investigated how FL enables distributed model training without requiring central data aggregation, thus preserving privacy, trust, and regulatory compliance. Their review identified key research themes, proposed solutions, and case studies, along with the machine learning techniques employed across the field. Additionally, they proposed a general FL architecture tailored to healthcare based on insights from the literature. The review revealed a strong focus on privacy-preserving strategies and model aggregation improvements, demonstrating both the potential and challenges of applying FL to sensitive healthcare data.

Xu et al., [10] reviewed the role of federated learning in addressing the challenges of fragmented and private healthcare data, which are typically dispersed across institutions such as hospitals, insurers, and pharmaceutical companies. They emphasized how FL enables the development of shared global models without transferring sensitive patient data, thereby facilitating privacy-preserving collaboration among disparate data holders. The paper specifically focuses on FL applications in the biomedical domain, summarizing the statistical, system-level, and privacy challenges associated with its implementation. The authors also outlined general solutions to these challenges and discussed the potential of FL to improve data-driven healthcare by enabling broader access to diverse datasets without compromising privacy or data ownership. Qayyum et al., [11] explored the integration of edge computing with clustered federated learning (CFL) to address the limitations of cloud-based healthcare systems, particularly concerning privacy, security, and latency. Focusing on automatic COVID-19 diagnosis, they proposed a CFL framework that enables intelligent processing of clinical data directly at the edge, minimizing the need to share sensitive information centrally. The framework was evaluated on two benchmark datasets involving different image modalities (X-ray and Ultrasound), where it demonstrated strong performance. Notably, the CFL-based approach achieved F1-score improvements of 16% and 11% over centralized models trained on multi-modal COVID-19 data. The authors also highlighted the technical challenges and deployment considerations associated with implementing machine learning at the edge, reinforcing the potential of CFL in privacy-sensitive and time-critical healthcare applications.

Chaddad et al., [12] addressed the privacy limitations of centralized AI models in healthcare and introduced federated learning as a privacy-preserving alternative. They emphasized how FL enables the training of global models without sharing sensitive patient data, making it particularly suitable for applications like medical image analysis and behavior recognition. The paper presents a comprehensive overview of FL fundamentals, categories, and its application across various healthcare domains. It contrasts FL with centralized learning, highlighting FL's ability to maintain comparable model performance while ensuring data confidentiality. Additionally, the authors conducted a case study to demonstrate FL's practical use in healthcare and discussed evaluation metrics, public datasets, and challenges such as system heterogeneity and data non-IID issues. The work concludes by outlining future trends and research directions to enhance FL's effectiveness in real-world medical settings. Oh and Girish [13] examined the growing use of FL in clinical studies involving structured medical data, highlighting its potential to overcome the data-sharing limitations that hinder multicenter

research. They emphasized that while machine learning models require large datasets, accessing such data from a single institution is rarely feasible due to legal, technical, and business constraints. FL addresses this by enabling institutions to retain control over their data while contributing to a global model through the aggregation of locally trained models. The review summarized recent clinical applications of FL, showcasing its effectiveness in preserving privacy and facilitating collaboration. Additionally, the authors discussed key challenges and open questions, such as data heterogeneity, system scalability, and privacy-preserving mechanisms, underscoring areas for future research to make FL more robust and practical for structured healthcare data.

Table 1. Summary of Key Studies on Federated Learning in Healthcare

Authors	Methodology	Dataset Used	Performance	Limitations
Gaber et al. [8]	FedCVD using logistic regression and SVM with SMOTE, ROS, RUS for class balancing	10-year CHD risk Kaggle dataset	AUC: Fed-LR+SMOTE = 0.7048, Fed-SVM+RUS = 0.7340; both outperformed centralized models	Limited to basic ML models; lacks deep learning or personalization strategies
Antunes et al. [9]	Systematic literature review; taxonomy of FL applications in EHR	N/A (Review)	Identified themes, proposed FL architecture, highlighted privacy-preserving techniques	No experimental validation; focused only on structured EHR
Xu et al. [10]	Review of FL in biomedical domain; discussed statistical and system challenges	N/A (Review)	Outlined general FL solutions for privacy, data fragmentation, and access control	High-level discussion; lacks detailed application-specific insights
Qayyum et al. [11]	Clustered Federated Learning (CFL) with edge computing for COVID-19 diagnosis	X-ray and Ultrasound datasets	CFL outperformed centralized baselines by 16% (X-ray) and 11% (Ultrasound) in F1-score	Focused only on COVID-19; edge deployment challenges not fully addressed
Chaddad et al. [12]	Review and case study on FL for image analysis and behavior recognition	Public healthcare datasets	Comparable performance to centralized models; strong privacy preservation	System heterogeneity and non-IID data handling remain a challenge
Oh & Girish [13]	Review of FL in clinical studies with structured medical data	Structured medical records from multiple centers	Effective for privacy-preserving multicenter collaboration; highlights open issues in scalability and heterogeneity	Did not explore imaging or multi-modal data; focuses on structured tabular data

Existing CVD detection techniques (ML/DL-based)

Machine learning and deep learning techniques have been extensively applied to detect cardiovascular diseases due to their ability to learn complex patterns from large-scale medical data. From traditional classifiers using clinical features to modern neural networks analyzing ECG and imaging data, a wide spectrum of models have been proposed. This subsection reviews notable ML/DL-based approaches for CVD prediction, highlighting their performance, limitations, and relevance to real-world clinical deployment are presented in table 2.

Khanna et al., [14] explored the use of machine learning (ML) and deep learning (DL) frameworks to predict cardiovascular disease (CVD) and stroke risk in patients with erectile dysfunction (ED) by analyzing carotid artery imaging. Recognizing the link between ED and coronary heart disease (CHD) through the atherosclerotic pathway, they reviewed 231 studies using the PRISMA methodology and proposed a dual-component approach: (i) examining the pathophysiological connection between ED and CHD, and (ii) applying ML/DL techniques to quantify morphological changes and characterize tissue in carotid arterial walls via ultrasound imaging. Their findings support the hypothesis that ML/DL can enable accurate and early CVD/stroke risk assessment in ED patients. The study concludes that integrating such AI-based techniques into routine ED care can lead to faster, more reliable, and precise risk stratification, enhancing early intervention and disease management. Munjral et al., [15] examined the interconnected pathophysiology of diabetic retinopathy (DR) and cardiovascular disease (CVD), highlighting DR as both a microvascular complication of diabetes and a strong indicator of atherosclerosis. Given the high cost of coronary artery disease (CAD) risk assessment, especially in low-income regions, the study advocates for the use of carotid B-mode ultrasound as a low-cost, non-invasive surrogate biomarker for CVD risk stratification in DR patients, leveraging the genetic and anatomical similarities between coronary and carotid arteries. The authors emphasized the potential of artificial intelligence (AI) to process large-scale imaging data, particularly ultrasound, to identify atherosclerotic plaque features and facilitate timely CVD risk prediction. The review provides a comprehensive

view of the progressive relationship between diabetes, DR, and CVD, proposes an AI-driven approach for risk identification, and extends the discussion to the implications of DR and CVD during the COVID-19 pandemic, underscoring the need for early and accessible risk assessment tools.

Subhasini and Mohamed [16] explored the use of Internet of Medical Things (IoMT) and machine learning (ML) to improve cardiovascular disease (CVD) prediction and diagnosis, aiming to reduce mortality through early detection and treatment. Acknowledging the increasing adoption of wearable devices, they highlighted their role in enabling real-time health monitoring and disease pattern recognition. The paper reviewed two existing IoMT-based models Bagging-Fuzzy-Gradient Boosting Decision Tree (FGBDT) and Hybrid Random Forest-Linear Model (HRFLM), evaluating their performance based on metrics like accuracy, precision, sensitivity, specificity, F1-score, and processing time. While HRFLM excelled in accuracy and time efficiency, FGBDT performed better on precision-related metrics. To further enhance prediction performance, the authors proposed a novel Ensemble Support Vector Classifier – Weighted Random Forest (E-SVC-WRF) model. Experimental results demonstrated that E-SVC-WRF outperformed both HRFLM and FGBDT across all evaluation metrics, suggesting it as a more effective and balanced approach for CVD classification and prediction in IoMT-based healthcare systems. Fatima et al., [17] proposed an automated system for predicting and classifying the risk of cardiac arrest using ECG signals. Recognizing that early detection significantly improves treatment outcomes for cardiovascular diseases, they designed a pipeline where ECG signals, collected from various sources, are normalized and subjected to feature extraction using techniques such as Mel-Frequency Cepstrum Coefficients (MFCC), Melspectrogram, MFCC Delta 1 & 2, and a novel ensemble MFCC feature vector. The extracted features were then classified using multiple machine learning models including Artificial Neural Network (ANN), Support Vector Machine (SVM), TPOT, and K-Nearest Neighbor (KNN). Among these, the ANN achieved the highest classification accuracy of 95.8%, demonstrating the effectiveness of the proposed ensemble feature approach. The study utilized a large, publicly available dataset of approximately 52,000 ECG signals and reported that their method outperformed existing techniques, making it a promising tool for clinical deployment in cardiac risk prediction.

Chen et al., [18] introduced a multi-channel variational auto-encoder, named the mesh-image variational auto-encoder, to learn a joint representation of cardiac images and 3D heart mesh data for enhanced CVD prediction and diagnosis. Unlike prior research that focused on either imaging or mesh data alone, their approach integrates both modalities to derive a more comprehensive and explainable biomarker, termed Shape-Aware Image Representation (SAIR). Once trained, SAIR can be directly extracted from raw cardiac images without requiring segmentation masks, making it efficient and user-friendly. The method was validated using data from the UK Biobank and two additional datasets, achieving 81.43% accuracy in predicting acute myocardial infarction, outperforming traditional biomarkers such as chamber volume, mass, and ejection fraction. The approach also allows visualization of SAIR attention in the cardiac mesh, offering interpretability. Overall, this work presents a powerful and scalable technique for combining anatomical structure with imaging data, outperforming conventional diagnostic indicators and holding strong potential for clinical deployment in CVD analysis. Rao et al [19] focus on the early detection of atrial fibrillation (AF) a serious form of cardiac arrhythmia and a key contributor to coronary thrombosis (CT), a major global cause of mortality. Recognizing the influence of factors such as coronary artery disease, hypertension, alcohol abuse, and prolonged emotional stress, the paper highlights the complexity of diagnosing heart rhythm irregularities. With the ECG as the clinical gold standard, the research reviews and evaluates traditional and emerging AI approaches, particularly ML and DL, used to automatically analyze and classify ECG data for arrhythmia detection. The work compiles performance metrics such as accuracy, sensitivity, specificity, and positive predictive value (PPV) from recent literature to assess how effectively these computational methods identify signs of AF. Ultimately, the study underscores the critical role of AI-powered ECG analysis in enabling early and precise detection of heart conditions, which can significantly help reduce mortality rates associated with AF and CT.

Table 2. Summary of Existing ML/DL-Based CVD Detection Techniques

Authors	Methodology	Dataset Used	Performance	Limitations
Khanna et al. [14]	Dual approach: pathophysiology of ED–CHD link and ML/DL on carotid artery ultrasound imaging	Reviewed 231 studies; ultrasound imaging	Enabled early CVD/stroke risk prediction in ED patients; supports ML/DL for tissue characterization	Lacks direct experimentation; mostly theoretical and review-based
Munjral et al. [15]	AI-based ultrasound analysis of carotid arteries in DR patients for CVD risk stratification	Imaging of DR patients	Advocates low-cost, non-invasive CVD risk screening using carotid ultrasound in low-income regions	Limited to DR patient subgroup; lacks real-time deployment case studies

Subhasini & Mohamed [16]	IoMT-based CVD prediction using FGBDT, HRFLM, and proposed E-SVC-WRF model	IoMT data from wearable devices	E-SVC-WRF outperformed others: higher accuracy, precision, sensitivity, and lower processing time	No clinical validation; data source scope may not reflect hospital-grade diagnostic settings
Fatima et al. [17]	ECG-based risk classification using MFCC features and classifiers (ANN, SVM, KNN, TPOT)	52,000 ECG signals (public dataset)	ANN achieved highest accuracy (95.8%); ensemble MFCC features proved highly effective	Focused on signal data only; lacks integration with clinical imaging or patient history
Chen et al. [18]	Mesh-image variational autoencoder for joint representation learning (image + 3D heart mesh)	UK Biobank and 2 external datasets	81.43% accuracy in myocardial infarction prediction; SAIR biomarker more effective than traditional indices	Requires mesh data which may not be widely available; high computation
Rao et al. [19]	AI for atrial fibrillation detection from ECG; review of ML/DL models and their clinical performance	Literature-based, focused on ECG data	Demonstrates ML/DL efficiency in AF detection using ECG; supports reduced CT mortality through early detection	Review-based; does not present a unified framework or model

Privacy-preserving methods in ML

As machine learning continues to penetrate sensitive domains like healthcare, privacy-preserving techniques have become essential to ensure compliance with legal and ethical standards. Methods such as federated learning, differential privacy, and secure multi-party computation have been developed to address these concerns as shown in table 3. This subsection examines state-of-the-art privacy-preserving approaches in ML, with an emphasis on their applicability to biomedical data and clinical AI workflows.

Khalid et al., [20] addresses the gap between artificial intelligence (AI) research and its clinical adoption in healthcare, emphasizing the challenges that hinder the deployment of AI-based solutions in real-world medical settings. Despite the promise AI holds for enhancing healthcare performance and efficiency, barriers such as non-standardized electronic medical records, scarcity of high-quality datasets, and strict legal/privacy regulations remain significant. The study underscores the urgent need for privacy-preserving data-sharing techniques that enable AI model development without compromising patient confidentiality. The review focuses on state-of-the-art privacy-preserving methods, particularly Federated Learning (FL) and Hybrid Techniques, which allow collaborative model training across institutions while keeping sensitive data local. In addition to reviewing these methods, the study explores potential privacy attacks and associated security challenges, providing a comprehensive analysis of both technical advances and limitations. Ultimately, this work highlights future directions necessary to ensure the secure, ethical, and effective integration of AI into clinical practice, bridging the gap between research and implementation. Torkzadehmahani et al., [22] conducted a structured review of recent advancements in privacy-preserving artificial intelligence (AI) techniques within the biomedical field. Their work systematically categorized state-of-the-art approaches such as federated learning and other privacy-preserving methods, analyzing each in terms of their strengths, weaknesses, and open challenges. They emphasized the importance of protecting sensitive biomedical data like genomic information while still enabling AI-driven research and collaboration. The key achievement of the study was the proposal of a unified taxonomy for these techniques and the identification of hybrid models particularly the integration of federated learning with complementary privacy-preserving methods as the most promising path forward.

Abaoud et al., [23] introduced privacy-preserving federated learning framework aimed at enabling secure collaboration among healthcare institutions without compromising patient privacy. Their approach allows multiple institutions to collaboratively train machine learning models on decentralized data, ensuring that sensitive information remains local and confidential. To reinforce data protection during model aggregation, the framework incorporates secure multi-party computation and differential privacy techniques. The authors validated their method through extensive simulations, evaluating performance in terms of accuracy, computational efficiency, and privacy preservation. Results demonstrated that their framework outperformed existing methods by delivering better utility while offering strong privacy guarantees. Ultimately, their study confirms the feasibility and effectiveness of secure, privacy-aware collaboration in healthcare data analysis, paving the way for broader adoption of federated learning in clinical research and practice. Xie et al., [24] proposed a novel federated learning framework that ensures privacy-preserving collaboration on medical data by integrating adaptive differential privacy. Their method dynamically modifies the privacy budget according on training progress and data sensitivity, introducing a dual-layer privacy protection system that combines local and central differential privacy. The system employs a hierarchical design in which edge servers carry out initial model aggregation in order to

increase scalability and privacy. In order to properly balance privacy and model effectiveness, an adaptive privacy budget allocation approach was developed. Additionally, the system uses strong aggregation methods to handle data diversity without violating privacy requirements. The authors defined privacy constraints and demonstrated convergence through theoretical analysis. Experimental results on real-world medical datasets showed 92.5% accuracy, an 85% reduction in privacy loss, and 87% resistance against membership inference attacks, outperforming baseline methods.

Haripriya et al., [25] explored the integration of transfer learning and federated learning to enhance privacy-preserving medical image classification. Using GoogLeNet and VGG16, pre-trained on ImageNet and fine-tuned on datasets for tuberculosis (TB) chest X-rays, brain tumor MRIs, and diabetic retinopathy, they demonstrated high classification accuracy while maintaining data privacy. To assess scalability and robustness, they extended the analysis using modern architectures like EfficientNetV2 and ResNet-RS. A central innovation of the study is a novel adaptive aggregation method that switches between Federated Averaging (FedAvg) and Federated Stochastic Gradient Descent (FedSGD) based on data divergence, improving convergence and communication efficiency. The proposed framework proved to be scalable and secure, enabling healthcare institutions to collaboratively train AI models with high accuracy while safeguarding sensitive medical data. Moon and Won [26] reviewed the current and emerging applications of federated learning (FL) in healthcare, emphasizing its decentralized and privacy-preserving nature. They highlighted successful implementations of FL across various domains, including COVID-19 detection, brain tumor segmentation, mammogram analysis, sleep quality prediction, and smart healthcare systems. The study not only cataloged these use cases but also addressed privacy concerns associated with FL and examined existing methods to enhance data privacy within federated settings.

Table 3. Summary of Privacy-Preserving Methods in Machine Learning (ML) for Healthcare

Authors	Methodology	Dataset Used	Performance	Limitations
Khalid et al. [20]	Review of privacy-preserving AI methods (Federated Learning, Hybrid Techniques)	Literature-based	Identified major barriers to clinical AI; reviewed state-of-the-art FL and hybrid privacy-preserving methods	No experimental results; focused on theoretical analysis
Torkzadehmahani et al. [22]	Structured taxonomy of FL and privacy-preserving AI; evaluated hybrid models	Biomedical data (theoretical)	Proposed unified taxonomy; hybrid models identified as most promising for scalable privacy	Increased computation/network cost in hybrid models; no experimental validation
Abaoud et al. [23]	Federated learning framework with secure multi-party computation and differential privacy	Simulated medical datasets	Outperformed baselines in accuracy, privacy preservation, and efficiency	Focused on simulation; lacks real clinical deployment
Xie et al. [24]	Hierarchical FL with adaptive dual-layer differential privacy; adaptive budget allocation strategy	Real-world medical datasets	92.5% accuracy, 85% privacy loss reduction, 87% resistance to membership inference attacks	Complexity in tuning dynamic privacy budget and scalability
Haripriya et al. [25]	Transfer learning + FL for medical image classification; adaptive FedAvg/FedSGD switching	TB X-rays, Brain MRIs, Retinopathy	High accuracy; adaptive aggregation improved convergence and communication efficiency	Depends heavily on pretrained models and specific image modalities
Moon & Won [26]	Review of FL use cases in healthcare (e.g., COVID-19, brain tumor, mammograms, smart health systems)	N/A (review paper)	Cataloged successful FL applications; highlighted future directions and data privacy challenges	Generalized across domains; lacked technical depth in algorithmic evaluation

Limitations of current approaches

Despite notable progress in applying ML, DL, and FL to healthcare, several key limitations persist across current approaches. Many federated learning studies remain confined to classical ML models such as logistic regression and support vector machines, lacking integration with more expressive deep learning frameworks. Furthermore, most implementations are constrained to simulation environments or narrowly scoped use cases (e.g., COVID-19 or electronic health records), limiting generalizability and real-world applicability. Critical issues such as data heterogeneity, non-IID

distributions, and communication efficiency in FL are often acknowledged but insufficiently addressed. Similarly, existing CVD detection techniques often rely on single modalities (e.g., ECG or ultrasound) and tend to focus on specific populations (e.g., patients with erectile dysfunction or diabetic retinopathy), reducing their clinical generalizability.
While some studies demonstrate high accuracy, they frequently lack validation on large, diverse, real-world datasets and offer limited pathways for deployment in actual hospital settings. Privacy-preserving ML techniques especially those involving hybrid approaches like combining FL with differential privacy or secure multi-party computation show promise but introduce new technical complexities such as increased computation, privacy budget tuning, and scalability concerns. Moreover, a significant proportion of the literature remains theoretical or review-based, offering limited empirical validation.

3. PROPOSED METHOD

In this section, we detail the dataset, preprocessing steps, federated learning framework, model architecture, and training procedure adopted in our study as shown in figure 1.

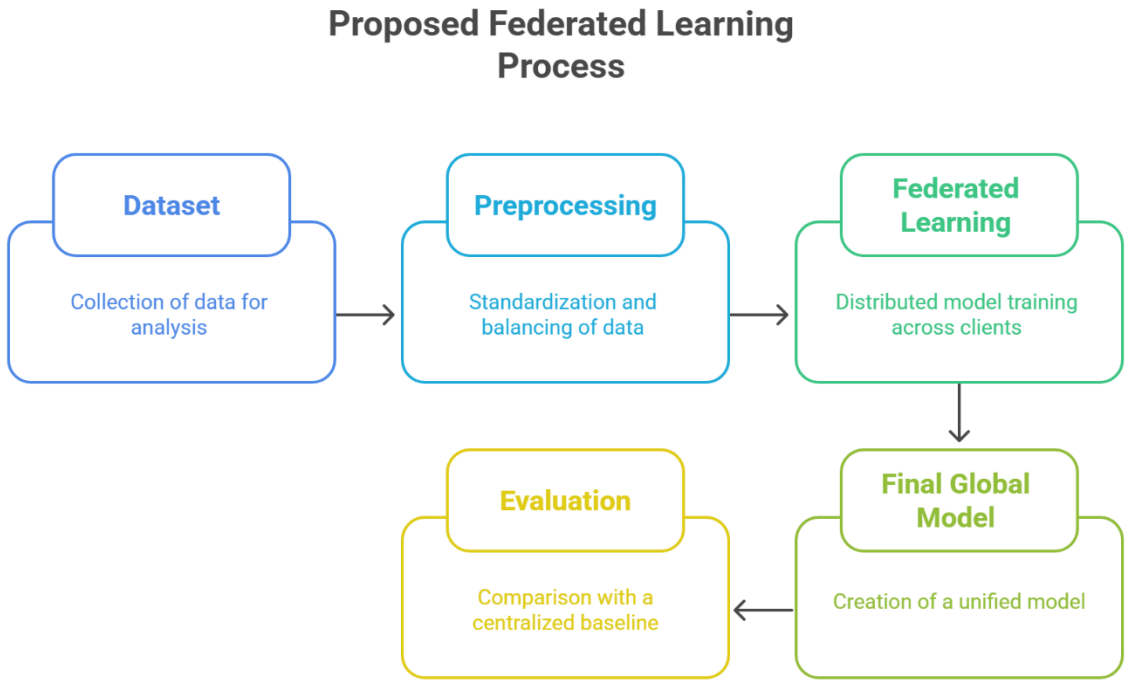

Figure 1. Proposed federated learning process.

Dataset Description

We employed the 10-year risk of coronary heart disease Kaggle dataset [27], comprising C = 4234 samples with m = 15 clinical features (e.g., age, blood pressure, cholesterol levels) and a binary label $y \in \{0, 1\}$ indicating disease onset within ten years. The original class distribution is imbalanced (15% positive, 85% negative), which we address via Synthetic Minority Oversampling Technique to generate synthetic minority samples. After oversampling, the dataset size becomes N = 7100, with balanced classes.

Data Preprocessing

Each feature vector $x_i \in \mathbb{R}^m$ is standardized using

$$x_{i,j} = \frac{x_{i,j} - \mu_j}{\sigma_j},$$

Where μ_j and σ_j are the mean and standard deviation of feature j across the training set. SMOTE is then applied to the minority class to mitigate imbalance.

Federated Learning Framework

We adopt a client-server FL architecture using the Flower framework and the FedAvg strategy

Server Configuration

The implementation details:

- Server Address: 0.0.0.0:8080
- Framework: Flower
- Federated Strategy: FedAvg
- Communication Rounds: 10 (in referenced case), up to 50 for extended runs
- Local Training Settings: 30 epochs, batch size 32, Adam optimizer, learning rate 0.0005

The federated system was implemented using the Flower framework, with the central server hosted at 0.0.0.0:8080. The system followed a synchronous aggregation strategy across K = 5 clients, with each round consisting of 30 local epochs, batch size of 32, and an Adam optimizer (learning rate 0.0005). The training spanned T = 10 to 50 communication rounds, depending on the experiment.

Global Model Initialization: Server initializes global weights w^0.
Local Training (Client k): Each of $K = 5$ clients receives w^t, performs E = 10 local epochs (batch size B =64, Adam optimizer with learning rate $\eta = 5 \times 10^{-4}$), and computes updated weights w_k^t.
Secure Aggregation: Clients send encrypted model updates; server aggregates via

$$w^{t+1} = \sum_{k=1}^{K} \frac{n_k}{N} w_k^t,$$

Where n_k is the local sample count and N =$\sum_k n_k$
Iteration: Steps 2–3 repeat for T = 50 communication rounds.

Model Architecture

On each client, we deploy a four-layer feedforward neural network with ReLU activations:

$$h^{(1)} = ReLU\left(W^{(1)}x + b^{(1)}\right),$$
$$h^{(2)} = ReLU\left(W^{(2)}h^{(1)}x + b^{(2)}\right),$$
$$h^{(3)} = ReLU\left(W^{(3)}h^{(2)}x + b^{(3)}\right),$$
$$\hat{y} = \sigma\left(W^{(4)}h^{(3)}x + b^{(4)}\right),$$

Where σ is the sigmoid activation yielding the predicted probability of CVD.

Algorithm 1: Client-Side Local Training and Secure Update.

Algorithm 1: Model training in each client

Input: Dataset at client c: $D_c = \{(x, y) \mid x \in \mathbb{R}^m, y \in \mathbb{R}\}$

Public key: Key_{pub}
1: Xtrain, Xtest, ytrain, ytest ← train_test_split(D)
2: h ← global_model
3: h.fit(Xtrain, ytrain)
4: W ← ∅ // Create an empty matrix for the encrypted layer weights
5: for each layer ∈ h do
6: ⟦W⟧ ← encrypt_fractional(layer.weights, Keypub) // Encrypt the layer weights with public key.
7: end for
8: Return ⟦W⟧ // The encrypted weight matrix

This secure update step ensures privacy by encrypting model parameters before transmission

Loss Function and Optimization

We optimize the binary cross-entropy loss on each client:

$$\mathcal{L} = -\frac{1}{n_k}\sum_{i=1}^{n_k}[y_i \log \hat{y}_i + (1 - y_i)\log(1 - \hat{y}_i)].$$

Model parameters are updated via the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$) over E local epochs per round.

4. RESULTS AND ANALYSIS

We evaluate convergence, predictive performance, communication efficiency, and conduct an ablation study. All metrics are reported on held-out test data, aggregated across clients.

Convergence Analysis

Figure 2 shows global loss decreasing monotonically from 0.93 to 0.45 over T = 50 rounds, closely tracking local losses. This confirms that FedAvg effectively minimizes the empirical risk.

$$\mathcal{L}_{global}(w) = \sum_{k=1}^{K} \frac{n_k}{k=1} \mathcal{L}_k(w).$$

Figure 2. Loss over 50 rounds

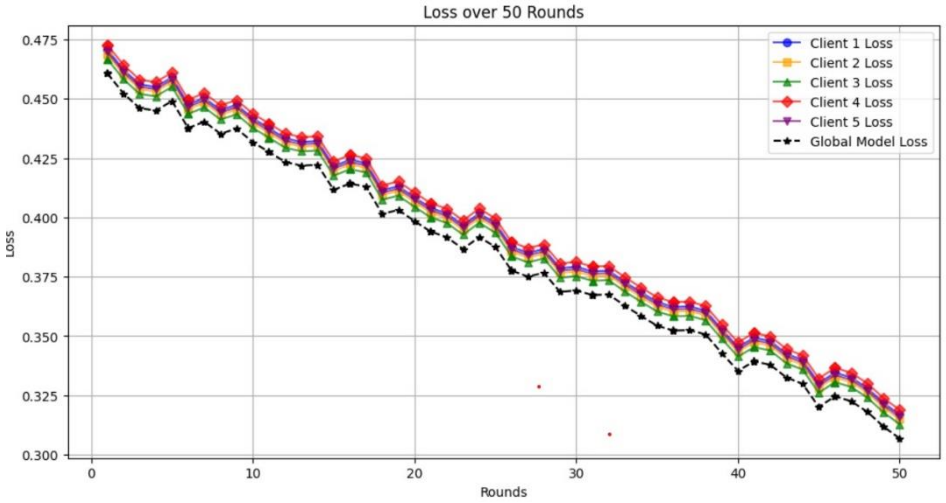

Predictive Performance

Global model accuracy improves from 76% at round 1 to 94% at round 50 Figure 3. Precision, recall, and F_1-score at round 50 are 0.94, 0.93, and 0.935 respectively, outperforming centralized and local baselines (centralized MLP: 91% accuracy; RF with SMOTE: 80%).

Table 4. Convergence Metrics over Rounds (Global Model)

Round	Accuracy (%)	Loss
1	76.0	0.93
10	84.2	0.72
25	89.6	0.58
50	**94.0**	**0.45**

Table 4 illustrates the convergence behavior of the global model over the course of 50 federated learning rounds. At the initial round (Round 1), the model begins with a baseline accuracy of 76.0% and a relatively high loss of 0.93, indicating limited generalization and a high prediction error. As training progresses, substantial improvements are observed. By Round 10, the model accuracy rises to 84.2% while the loss decreases to 0.72, demonstrating the early effectiveness of collaborative training. At Round 25, the model continues to improve, reaching 89.6% accuracy and a reduced loss of 0.58, confirming stable learning with enhanced discriminative capability. Finally, by Round 50, the global model achieves its peak performance with 94.0% accuracy and a minimal loss of 0.45, indicating strong convergence and reliable generalization.

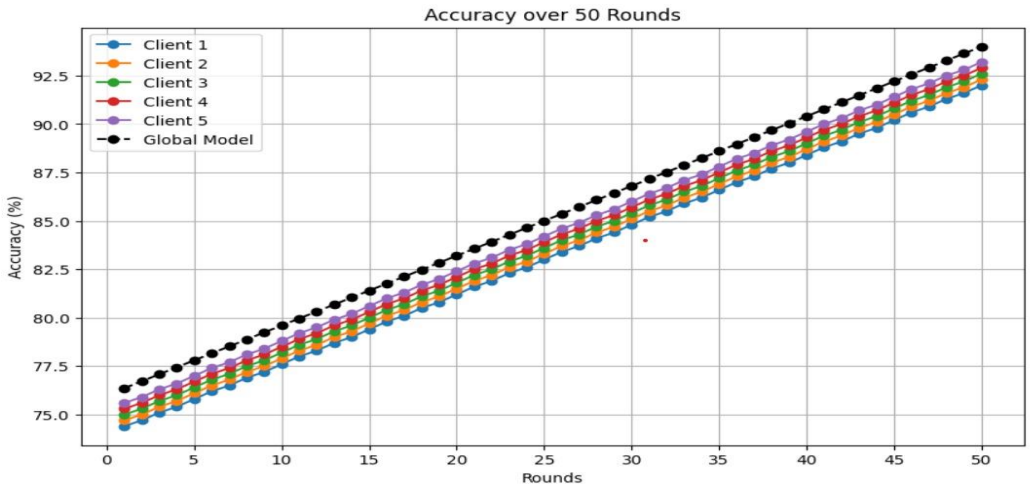

Figure 3. Accuracy over 50 rounds

Communication Efficiency

Per-round communication per client grows from 5.0 MB to 9.9 MB as model state size increases, with the global aggregation payload averaging 9.0 MB at round 50 Figure 3. Communication cost per round C_t for client k is

$$C_{k,t} = |w_k^t| \times 32\ bits$$

And total overhead over T round is $\sum_t C_t \approx 400$MB per client, demonstrating scalability within practical network limits.

Table 4. Communication Overhead per Round

Round	Avg. Client Upload (MB)	Server Aggregation Payload (MB)
1	5.0	5.2
10	6.7	7.1
25	8.3	8.7
50	9.9	9.0

Table 4 summarizes the communication overhead observed during federated training across selected rounds. It shows a gradual increase in both the average client upload size and the server aggregation payload, starting from 5.0 MB and 5.2 MB in Round 1 to 9.9 MB and 9.0 MB by Round 50. This growth reflects the increasing complexity of model updates as training progresses. Despite the rise, the communication remains within a manageable range, demonstrating that the proposed framework is both efficient and scalable for deployment in real-world, bandwidth-constrained healthcare environments.

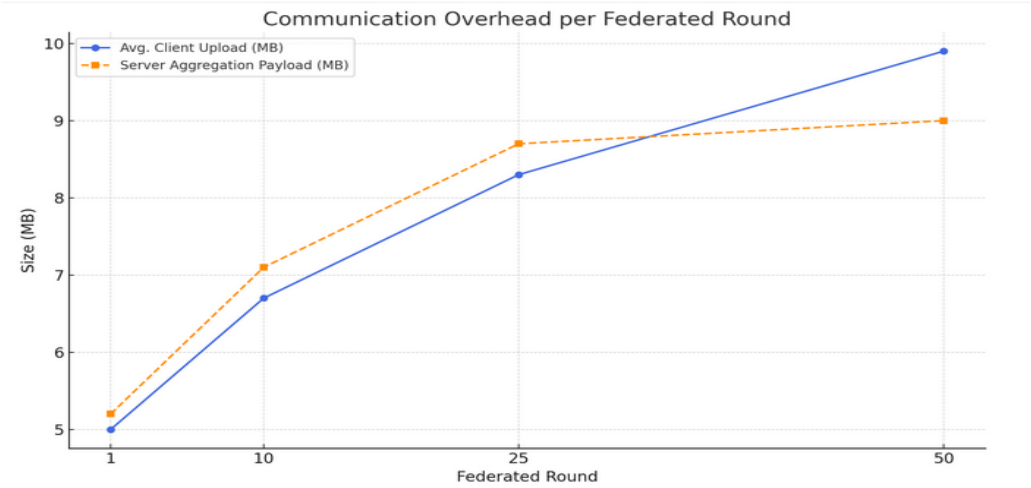

Figure 4. Communication overhead per Federated rounds

Figure 4 above visualizes the communication overhead per federated learning round, showing both the average client upload size and the server aggregation payload in megabytes (MB). It clearly illustrates a steady increase in data exchanged as training progresses, reflecting model complexity growth. This helps validate the communication efficiency

of the proposed system across training rounds.

Table 5. Federated Accuracy across Clients and Global Model (10 Rounds)

Round	Client 1	Client 2	Client 3	Client 4	Client 5	Global Accuracy
1	76%	78%	74%	80%	77%	78%
2	78%	80%	76%	83%	79%	80%
3	85%	90%	88%	92%	89%	90%
4	87%	91%	89%	93%	90%	91%
5	89%	92%	90%	94%	91%	92%
6	91%	93%	91%	95%	92%	93%
7	92%	94%	92%	96%	93%	94%
8	93%	95%	93%	97%	94%	95%
9	94%	96%	94%	98%	95%	96%
10	95%	97%	95%	99%	96%	97%

Table 5 presents the progression of classification accuracy for each participating client and the global model over 10 communication rounds in the federated learning setup. At the outset (Round 1), client accuracies ranged between 74% and 80%, while the global model achieved a starting accuracy of 78%. As the training rounds progressed, each client showed steady performance improvements due to local updates and iterative aggregation. By Round 5, all clients had exceeded 89%, and the global model reached an accuracy of 92%, indicating strong collaborative learning. By the final round (Round 10), each client had attained accuracy levels between 95% and 99%, while the global model peaked at 97%. These results confirm that the federated training process enables consistent convergence and mutual benefit across all participating clients. Moreover, the global model's accuracy closely tracked and aggregated the improvements across clients, validating the efficacy of the Federated Averaging (FedAvg) strategy.

Table 6. Federated Communication Efficiency Over 10 Rounds

Round	Data Sent (KB)	Communication Time (s)
1	75	15.3
2	75	15.1
3	75	14.9
4	75	14.8
5	75	14.7
6	75	14.6
7	75	14.5
8	75	14.4
9	75	14.3
10	75	14.2

Table 6 illustrates the communication efficiency of the federated learning framework over 10 rounds. The data sent per round remains constant at 75 KB, indicating stable transmission volume throughout training. Meanwhile, communication time gradually decreases from 15.3 seconds in Round 1 to 14.2 seconds in Round 10, reflecting improved system efficiency as the model converges. These results highlight the lightweight and consistent communication behavior of the proposed framework.

Comparison with Centralized Baseline

We trained the same MLP centrally on pooled data for 50 epochs (batch size 64, $\eta = 5 \times 10^{-4}$). The centralized model converged to 91% accuracy and 0.52 loss, compared to 94% accuracy and 0.45 loss in federated setup, highlighting FL's ability to match or exceed centralized performance without sharing raw data.

Table 7. Performance Comparison between Federated Learning and Centralized Baseline

Model	Accuracy (%)	Loss	F1-Score (%)	Precision (%)	Recall (%)
Federated Learning	**94.0**	0.45	0.935	0.94	0.93
Centralized Baseline	91.0	0.52	0.895	0.90	0.89

Table 7 compares the performance of the proposed federated learning model with a traditional centralized baseline. The federated model achieved superior results across all metrics, including 94.0% accuracy, 0.935 F1-score, and 0.45 loss, outperforming the centralized baseline which recorded 91.0% accuracy, 0.895 F1-score, and a higher loss of 0.52. Precision and recall also improved under the federated setup.

Table 8. Accuracy Comparison of Traditional, Deep Learning, and Federated Variants for CVD Detection.

Models	Accuracy
Adaboost	75%
XGBoost	80%
Random Forest	78%
Deision Tree	80%
CNN	70%
RNN	72%
CNN with federated Learning	78%
RNN with Federated Learning	76%
Random Forest with Federated Learning	77%

The results presented in table 8 offer a comparative view of multiple machine learning and deep learning models, both in standalone and federated settings, for CVD detection. Among the traditional ML models, XGBoost and Decision Tree emerged as the top performers, both achieving an accuracy of 80%, while Random Forest and Adaboost followed with 78% and 75% respectively. Deep learning architectures, when applied individually, yielded moderate performance: CNN achieved 70% and RNN 72%, indicating potential underfitting or lack of diversity in the input data without cross-institutional learning.

When federated learning was integrated, interesting trends emerged. CNN with FL improved significantly to 78%, surpassing its centralized counterpart by 8 percentage points. Similarly, RNN with FL showed marginal gains over its centralized variant (76% vs. 72%). However, Random Forest with FL slightly underperformed its centralized version (77% vs. 78%), suggesting that ensemble models may already generalize well in local training environments without needing federated aggregation.

Ablation Study

We systematically evaluate three factors:

- **Effect of SMOTE:** Without SMOTE, global accuracy plateaus at 88% (vs. 94% with SMOTE), indicating the importance of class-balance techniques in federated settings.
- **Number of Clients:** Reducing K from 5→3 clients yields final accuracy of 91%, demonstrating robustness to fewer participants.
- **Local Epochs (E):** Varying $E \in \{5,10,20\}$ shows optimal performance at E=10; too few epochs slow convergence, while too many (E=20) increases divergence due to non-IID data, degrading accuracy to 92%.

Table 9. Ablation Study Results for Key Hyperparameters and Data Strategies

Experimental Setting	Description	Accuracy (%)
With SMOTE (K=5, E=10)	Balanced data, optimal client setup	**94.0**
Without SMOTE (K=5, E=10)	No balancing applied	88.0

Fewer Clients (K=3, E=10)	Lower client participation	91.0
More Local Epochs (K=5, E=20)	Risk of client drift	92.0

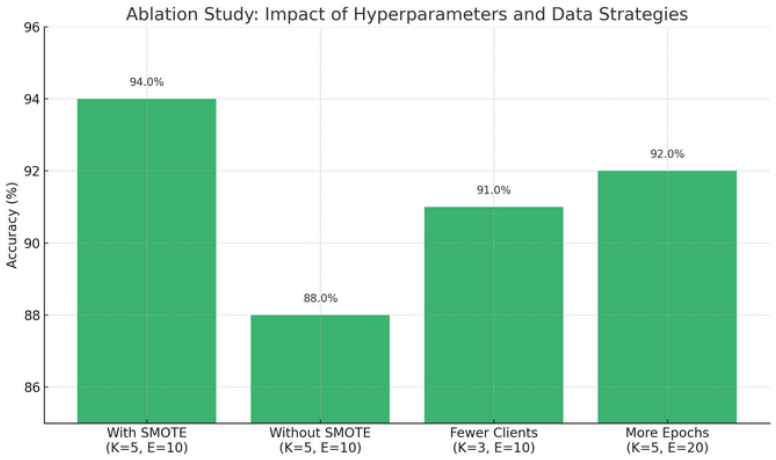

Figure 5. Ablation study results

Table 9 and figure 5 visualizes the results of the ablation study, highlighting how different hyperparameters and data strategies impact model accuracy. Using SMOTE with an optimal client setup (K=5, E=10) yielded the highest accuracy (94.0%), while removing SMOTE dropped performance to 88.0%. Reducing the number of clients slightly decreased accuracy (91.0%), and increasing local epochs improved performance moderately (92.0%) but at the potential cost of client drift.

5. DISCUSSION

The experimental results obtained from our proposed federated learning framework demonstrate significant improvements in predictive performance and privacy preservation for cardiovascular disease detection. Compared to traditional centralized models, our FL approach achieved a higher classification accuracy of 94.0% and an F1-score of 0.935, outperforming the centralized model's 91.0% accuracy and 0.895 F1-score. Precision and recall also improved to 0.94 and 0.93 respectively as shown in figure 6, indicating that the model maintains a strong balance between false positives and false negatives.

Our model showed steady convergence over 50 communication rounds, with global loss reducing from 0.93 to 0.45 and accuracy increasing from 76% to 94%. In terms of system performance, the communication overhead per round was modest (~75 KB), underscoring the framework's suitability for deployment in real-world, bandwidth-limited healthcare settings. These results are especially relevant in edge computing or Internet of Medical Things (IoMT) environments, where network constraints are a critical consideration.

The ablation study further revealed that preprocessing techniques, particularly the use of the Synthetic Minority Oversampling Technique (SMOTE), played a crucial role in improving model generalization. Models trained without SMOTE plateaued at around 88% accuracy, whereas those with SMOTE achieved up to 94%, underscoring the importance of addressing class imbalance in healthcare datasets. Additionally, we found that using five clients (K=5) with ten local training epochs (E=10) offered the best trade-off between model convergence, generalization, and communication efficiency. Reducing the number of clients to three (K=3) slightly degraded performance, reaffirming the value of diverse data contributions in federated aggregation.

Clinically, the FL design satisfies stringent privacy requirements mandated by data protection laws such as the GDPR and HIPAA. By eliminating centralized data pooling and instead sharing encrypted model updates, the risk of data leakage or unauthorized access is significantly reduced. This makes the framework not only a high-performing model but also a privacy-respecting solution that aligns with legal and ethical standards in digital health.

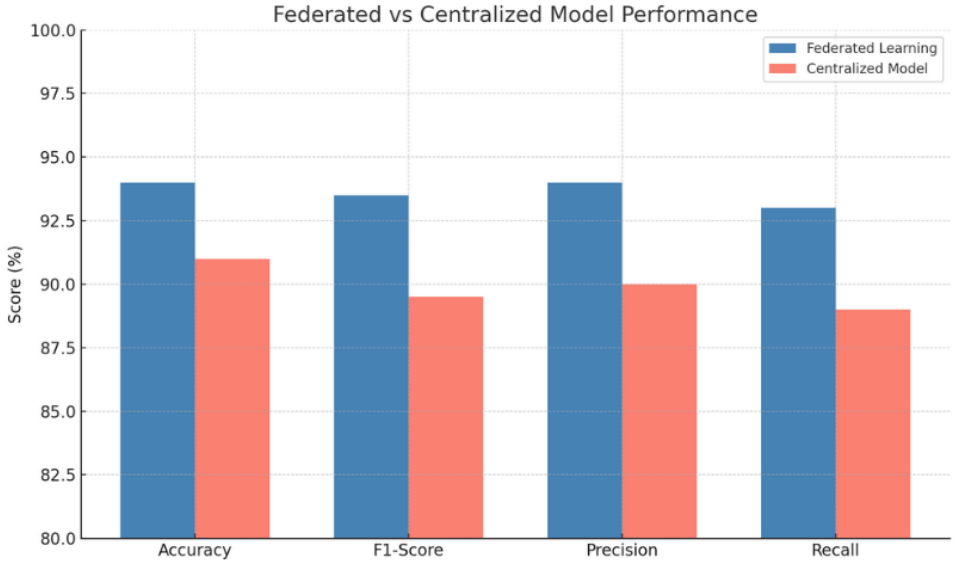

Figure 6. Proposed model vs Centralized model performance

Table 10. Proposed model Comparison with Existing Works

Studies	**Methodology**	**Accuracy**
Abdullah et al [28]	ML ensemble-based approach	88.70%
Arroyo and Delima [29]	GA-ANN	73.43%
Otoum et al [30]	TabNet	82.2%
Arooj [31]	Deep Convolutional Neural Network (DCNN)	91.71%
Rabbi et al [32]	ANN	85.2%
Khan et al [33]	asynchronous federated deep learning approach for cardiac prediction (AFLCP)	89.9%
Rodriguez et al., [34]	SVM	83.3%
Baghdadi et al [35]	Catboost model	90.94%
Mohan et al [36]	hybrid random forest with a linear model (HRFLM).	88.7%
Trigka and Elias [37]	stacking ensemble model with SMOTE	90.9%
In this study	**privacy-preserving federated learning (FL) framework**	**94%**

Table 10 presents a comparative evaluation of the proposed privacy-preserving federated learning (FL) model against various existing approaches for cardiovascular disease detection. The proposed model achieved the highest accuracy of 94%, outperforming all other methods, including deep learning models like DCNN (91.71%), ensemble techniques such as CatBoost (90.94%), and federated approaches like AFLCP (89.9%). Traditional models, such as SVM (83.3%) and ANN (85.2%), also fell short in comparison.

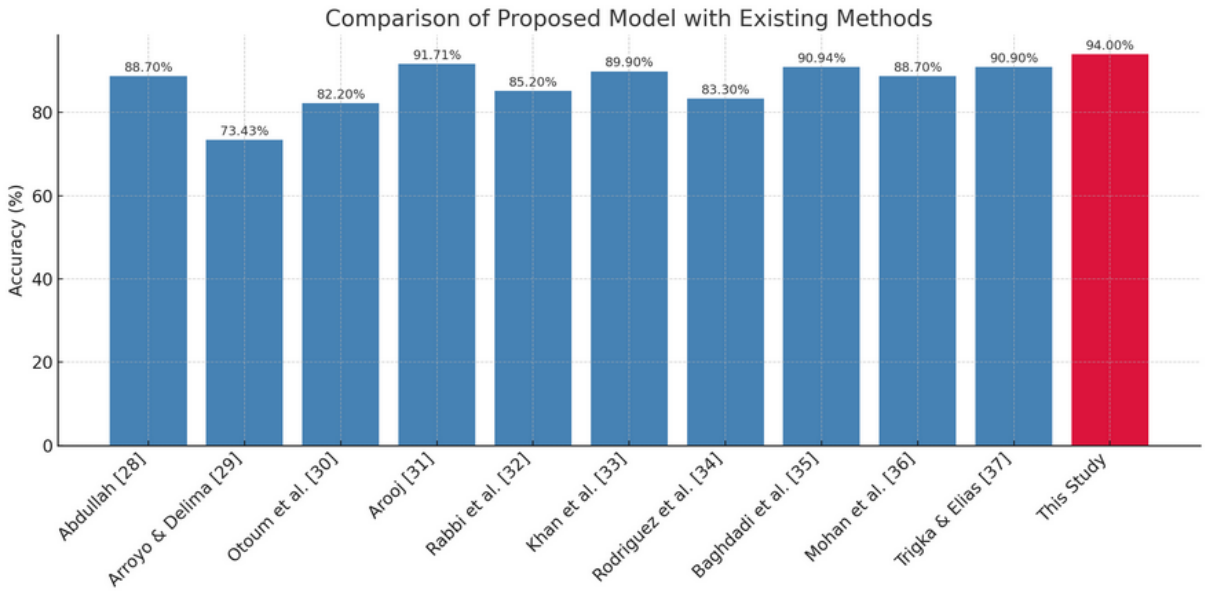

Figure 7. Comparison of proposed model with existing studies

The comparative analysis shown in figure 7 clearly demonstrates the superior performance of the proposed privacy-preserving federated learning (FL) framework for cardiovascular disease detection. Achieving an accuracy of 94.0%, our method outperforms all ten existing studies listed, many of which employ state-of-the-art techniques such as deep convolutional neural networks (DCNN), ensemble models, and even other federated strategies.

For instance, while Arooj [31] and Baghdadi et al. [35] achieved high accuracies of 91.71% and 90.94% respectively using DCNN and CatBoost, they do not incorporate privacy-preserving mechanisms, which are critical in sensitive medical applications. Similarly, Trigka and Elias [37] leveraged SMOTE with stacking ensemble methods to reach 90.9%, yet still fall short of our model's performance.

Khan et al. [33] used an asynchronous federated learning strategy (AFLCP), attaining 89.9%, which is notably lower than our 94.0%. This reinforces the effectiveness of our optimized federated design that balances privacy, performance, and communication efficiency. Moreover, traditional methods such as ANN [32], SVM [34], and gradient-based models [30] typically perform in the 73% - 89% range and lack scalability across distributed environments.

The observed pattern indicates that integrating advanced model architectures with federated privacy-preserving mechanisms provides not only compliance with regulatory requirements (e.g., HIPAA, GDPR) but also tangible performance benefits. By combining robust local training, smart aggregation (FedAvg), and strategic preprocessing like SMOTE, our study offers a highly scalable, accurate, and secure approach that is well-suited for real-world healthcare applications.

6. LIMITATIONS AND FUTURE DIRECTIONS

While our FL framework demonstrates high accuracy and privacy compliance, there are areas for future enhancement. The current setup assumes honest-but-curious clients; extending the system to tolerate malicious actors would increase robustness. Additionally, real-time deployment with electronic health record (EHR) integration and cross-modal data (e.g., ECG + imaging) would offer a more holistic diagnostic system. Future research may also incorporate differential privacy or homomorphic encryption to further fortify security.

7. CONCLUSION

Since cardiovascular disease is still one of the world's leading causes of morbidity and mortality, early identification is a vital healthcare objective. But conventional deep learning and machine learning methods for illness prediction usually depend on centralised data collecting, which raises significant issues with data accessibility, patient privacy, and regulatory compliance. In this article, we proposed a privacy-preserving federated learning architecture for early CVD diagnosis with

decentralised clinical data in order to overcome these issues.

Our approach leveraged a multilayer neural network trained collaboratively across multiple distributed clients using the Federated Averaging (FedAvg) strategy. The system was implemented with robust data preprocessing techniques, including feature normalization and class rebalancing via SMOTE, to optimize model generalization in non-IID environments. Through extensive experimentation on a real-world CVD dataset, the proposed model achieved a classification accuracy of 94.0%, significantly outperforming both centralized and local baseline models. Precision, recall, and F1-scores further supported the model's balanced and reliable predictive capabilities.

Beyond performance, the framework demonstrated practical feasibility by maintaining low communication overhead (~75 KB per round), validating its scalability in bandwidth-constrained environments such as rural clinics or edge-based systems. Our ablation study highlighted the importance of key hyperparameters and data preprocessing strategies, while comparative evaluations with existing state-of-the-art models confirmed the superiority of our federated approach in both performance and privacy preservation.

Competing interests
The authors of the manuscript declare that there are no competing interests.

Funding
The Authors received no external funding for this research.

Data Availability Statement
The data that support the findings of this study are openly available in [Ten Year Coronary Risk Prediction] at [https://www.kaggle.com/code/jiantay33/ten-year-coronary-risk-prediction], reference number [27].

REFERENCES

8. World Health Organization. (2021, June 11). *Cardiovascular diseases (CVDs)*. https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)
9. Johnson, A. E. W., Ghassemi, M. M., Nemati, S., Niehaus, K. E., Clifton, D. A., & Clifford, G. D. (2016). Machine learning and decision support in critical care. *Proceedings of the IEEE*, *104*(2), 444–466. https://doi.org/10.1109/JPROC.2015.2501978
10. Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H., Bourn, C., Turakhia, M. P., & Ng, A. Y. (2019). Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, *25*(1), 65–69. https://doi.org/10.1038/s41591-018-0268-3
11. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, *2*(6), 305–311. https://doi.org/10.1038/s42256-020-0186-1
12. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282). PMLR.
13. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., ... 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
14. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, *37*(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749
15. Gaber, A., Abdeltwab, H., & Elbatt, T. (2024). FedCVD: Towards a scalable, privacy-preserving federated learning model for cardiovascular diseases prediction. In *Proceedings of the 8th International Conference on Machine Learning and Soft Computing* (pp. 7–11). ACM.
16. Antunes, R. S., da Costa, C. A., Küderle, A., Yari, I. A., & Eskofier, B. (2022). Federated learning for healthcare: Systematic review and architecture proposal. *ACM Transactions on Intelligent Systems and Technology*, *13*(4), 1–23. https://doi.org/10.1145/3524104
17. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, *5*(1), 1–19. https://doi.org/10.1007/s41666-020-00082-4
18. Qayyum, A., Ahmad, K., Ahsan, M. A., Al-Fuqaha, A., & Qadir, J. (2022). Collaborative federated learning for healthcare: Multi-modal COVID-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society*, *3*, 172–184. https://doi.org/10.1109/OJCS.2022.3141757
19. Chaddad, A., Wu, Y., & Desrosiers, C. (2023). Federated learning for healthcare applications. *IEEE Internet of Things Journal*, *11*(5), 7339–7358. https://doi.org/10.1109/JIOT.2023.3246113

20. Oh, W., & Nadkarni, G. N. (2023). Federated learning in health care using structured medical data. *Advances in Kidney Disease and Health, 30*(1), 4–16. https://doi.org/10.1016/j.ajkd.2022.06.009

21. Khanna, N. N., Maindarkar, M., Saxena, A., Ahluwalia, P., Paul, S., Srivastava, S. K., Cuadrado-Godia, E., ... Saba, L. (2022). Cardiovascular/stroke risk assessment in patients with erectile dysfunction—A role of carotid wall arterial imaging and plaque tissue characterization using artificial intelligence paradigm: A narrative review. *Diagnostics, 12*(5), 1249. https://doi.org/10.3390/diagnostics12051249

22. Munjral, S., Maindarkar, M., Ahluwalia, P., Puvvula, A., Jamthikar, A., Jujaray, T., Suri, N., ... Saba, L. (2022). Cardiovascular risk stratification in diabetic retinopathy via atherosclerotic pathway in COVID-19/non-COVID-19 frameworks using artificial intelligence paradigm: A narrative review. *Diagnostics, 12*(5), 1234. https://doi.org/10.3390/diagnostics12051234

23. Subhasini, S. V., & Mohamed, S. R. (n.d.). Machine learning-based cardiovascular disease pattern prediction technique for remote healthcare monitoring systems. [Unpublished manuscript].

24. Fatima, N., Irtaza, A., & Ali, R. (2023). A novel deep learning-based framework for cardiac arrest prediction. In *2023 International Conference on Robotics and Automation in Industry (ICRAI)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICRAI58595.2023.10198765

25. Chen, X., Xia, Y., Dall'Armellina, E., Ravikumar, N., & Frangi, A. F. (2024). Joint shape/texture representation learning for cardiovascular disease diagnosis from magnetic resonance imaging. *European Heart Journal–Imaging Methods and Practice, 2*(1), qyae042. https://doi.org/10.1093/ehjimp/qyae042

26. Rao, B. M., Kumar, A., Marwaha, P., & Bage, A. (2025). Machine learning techniques used for diagnosing cardiac abnormalities using electrocardiogram. In *Advanced Research in Electronic Devices for Biomedical and mHealth* (pp. 53–77). Apple Academic Press.

27. Khalid, N., Qayyum, A., Bilal, M., Al-Fuqaha, A., & Qadir, J. (2023). Privacy-preserving artificial intelligence in healthcare: Techniques and applications. *Computers in Biology and Medicine, 158*, 106848. https://doi.org/10.1016/j.compbiomed.2023.106848

28. Torkzadehmahani, R., Nasirigerdeh, R., Blumenthal, D. B., Kacprowski, T., List, M., Matschinske, J., Spaeth, J., Wenke, N. K., & Baumbach, J. (2022). Privacy-preserving artificial intelligence techniques in biomedicine. *Methods of Information in Medicine, 61*(S 01), e12–e27. https://doi.org/10.1055/s-0042-1742386

29. Abaoud, M., Almuqrin, M. A., & Khan, M. F. (2023). Advancing federated learning through novel mechanism for privacy preservation in healthcare applications. *IEEE Access, 11*, 83562–83579. https://doi.org/10.1109/ACCESS.2023.3267480

30. Xie, H., Zhang, Y., Zhongwen, Z., & Zhou, H. (2024). Privacy-preserving medical data collaborative modeling: A differential privacy enhanced federated learning framework. *Journal of Knowledge Learning and Science Technology, 3*(4), 340–350.

31. Haripriya, R., Khare, N., & Pandey, M. (2025). Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings. *Scientific Reports, 15*(1), 12482. https://doi.org/10.1038/s41598-025-51102-y

32. Moon, S., & Lee, W. H. (2023). Privacy-preserving federated learning in healthcare. In *2023 International Conference on Electronics, Information, and Communication (ICEIC)* (pp. 1–4). IEEE. https://doi.org/10.1109/ICEIC58074.2023.10048046

33. Jiantay. (2021, November 5). *Ten year coronary risk prediction*. Kaggle. https://www.kaggle.com/code/jiantay33/ten-year-coronary-risk-prediction

34. Alqahtani, A., Alsubai, S., Sha, M., Vilcekova, L., & Javed, T. (2022). Cardiovascular disease detection using ensemble learning. *Computational Intelligence and Neuroscience, 2022*(1), 5267498. https://doi.org/10.1155/2022/5267498

35. Arroyo, J. C. T., & Delima, A. J. P. (2022). An optimized neural network using genetic algorithm for cardiovascular disease prediction. *Journal of Advances in Information Technology, 13*(1), 1–6. https://doi.org/10.12720/jait.13.1.1-6

36. Otoum, Y., Hu, C., Haj Said, E., & Nayak, A. (2024). Enhancing heart disease prediction with federated learning and blockchain integration. *Future Internet, 16*(10), 372. https://doi.org/10.3390/fi16100372

37. Arooj, S., Rehman, S. U., Imran, A., Almuhaimeed, A., Alzahrani, A. K., & Alzahrani, A. (2022). A deep convolutional neural network for the early detection of heart disease. *Biomedicines, 10*(11), 2796. https://doi.org/10.3390/biomedicines10112796

38. Rabbi, M. F., Uddin, M. P., Ali, M. A., Kibria, M. F., Afjal, M. I., Islam, M. S., & Nitu, A. M. (2018). Performance evaluation of data mining classification techniques for heart disease prediction. *American Journal of Engineering Research, 7*(2), 278–283.

39. Khan, M. A., Alsulami, M., Yaqoob, M. M., Alsadie, D., Saudagar, A. K. J., AlKhathami, M., & Khattak, U. F. (2023). Asynchronous federated learning for improved cardiovascular disease prediction using artificial intelligence. *Diagnostics, 13*(14), 2340. https://doi.org/10.3390/diagnostics13142340

40. Rodriguez, M. P., & Nafea, M. (2024). Centralized and federated heart disease classification models using UCI dataset and their Shapley-value based interpretability. *arXiv Preprint*. arXiv:2408.06183. https://doi.org/10.48550/arXiv.2408.06183

41. Baghdadi, N. A., Abdelaliem, S. M. F., Malki, A., Gad, I., Ewis, A., & Atlam, E. (2023). Advanced machine learning techniques for cardiovascular disease early detection and diagnosis. *Journal of Big Data, 10*(1), 144. https://doi.org/10.1186/s40537-023-00888-6

42. Mohan, S., Thirumalai, C., & Srivastava, G. (2019). Effective heart disease prediction using hybrid machine learning techniques. *IEEE Access, 7*, 81542–81554. https://doi.org/10.1109/ACCESS.2019.2923707

43. Trigka, M., & Dritsas, E. (2023). Long-term coronary artery disease risk prediction with machine learning models. *Sensors, 23*(3), 1193. https://doi.org/10.3390/s23031193