TPM Vol. 32, No. S4, 2025
ISSN: 1972-6325
https://www.tpmap.org/

A TRUST-AWARE FEDERATED LEARNING FRAMEWORK WITH CONTEXT-AWARE DYNAMIC GRADIENT PRESERVATION FOR EARLY CARDIOVASCULAR RISK PREDICTION

[1] SATHYA S, [2] K.SARANYA
[1]ASSISTANT PROFESSOR
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
HINDUSTHAN COLLEGE OF ENGINEERING AND TECHNOLOGY, COIMBATORE, TAMIL NADU, INDIA
Email Id: sathyasuriyanme@gmail.com
ASSOCIATE PROFESSOR,
[2]DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING,
BANNARI AMMAN INSTITUTE OF TECHNOLOGY, SATHYAMANGALAM, TAMIL NADU, INDIA
Email Id: ksaranyacse@gmail.com

Abstract:
Cardiovascular diseases (CVD) is a major cause of death worldwide and require early and privacy guaranteeing risk prediction models. Current federated learning models employ gradient pruning methods such as Adaptive Gradient Pruning Optimization (AGPO) which can carelessly remove clinically important but low-magnitude features. This results in degraded convergence and compromised model fairness. To overcome these challenges, this paper introduces FedCure-X, a new trust-aware federated learning system that incorporates Context-Aware Dynamic Gradient Preservation Optimization (CD-GPO). In contrast to traditional pruning techniques, CD-GPO uses medical context-informed filters, convergence-aware scheduling and fairness regularization for maintaining semantically meaningful gradients even in heterogeneous and non-IID healthcare datasets. Secure aggregation, patient privacy and client trust are guaranteed by the framework along with enhanced predictive accuracy and convergence rate. The proposed framework is tested against benchmark datasets Framingham and MIMIC-III. FedCure-X shows better performance than baseline models with an accuracy rate of 96.8% and enhanced robustness at distributed clinical nodes. The proposed framework is ethical, scalable and smart in early detection systems for CVD risk prediction.

Keywords:
Federated Learning, Cardiovascular Disease Prediction, Gradient Optimization, Privacy Preservation, Non-IID Data, Trust-Aware Aggregation.

1. INTRODUCTION

Cardiovascular diseases (CVD) including coronary artery disease, stroke and heart failure are the leading cause of mortality and morbidity worldwide. The World Health Organization (WHO) indicates that more than 17.9 million deaths are caused by CVDs each year, responsible for 32% of all deaths on the planet [1]. Identification of people at increased risk early in life is essential for allowing intervention on time, minimizing the occurrence of complications, and enhancing survival. Conventional diagnostics are based on episodic clinical encounters, centralized repositories of data and discrete decision-making extend detection time and erect obstacles to targeted treatments. With the emergence of digital health solutions Electronic Health Records (EHRs) and wearables, enormous amounts of patient information are readily accessible. Timely and precise prediction of risk is a critical challenge in these decentralized datasets [2].

Federated Learning (FL) has shown potential as a healthcare AI paradigm, enabling decentralized training of models across various clinical institutions without the need for sharing raw patient data. Decentralized architecture helps tackle major privacy issues in addition to being compliant with data protection laws like Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR) [3]. By facilitating collective intelligence, Federated Learning (FL) enables collaboration across multiple institutions. It allows the harnessing of diverse, high-quality, and real-time health information such as clinical data, wearable device outputs, and diagnostic images from geographically dispersed sources [4]. Its applicability to CVD prediction is based on cardiovascular risk factors like blood pressure, blood sugar level, heart rate and lifestyle metrics are highly variable among populations. An intelligently designed federated architecture can provide a comprehensive and scalable answer to personalized risk analysis with assured data authority and ensured ethical AI deployment [5].

Some technical and practical issues prevent its effective application in CVD prediction in real-world settings. First, clinical data are non-IID because of differing clinical procedures, populations and sensor types across centers. Second, most FL methods use agressive gradient pruning techniques such as AGPO which will discard low-magnitude gradients that contain important clinical information. Moreover, most existing secure aggregation techniques tend to ignore trust relationships between participating nodes, hence vulnerable to unfair training. These advantages end up in making wrong predictions, bad convergence and biased model generalization in patients coming from understated regions [6].

It is important to create a federated learning system that not only maintains strong predictive performance but also data utility and trust. There is a strong requirement for a privacy preserving FL paradigm that can adaptively retain clinically useful information during training even under sparse and non-uniform gradients. In addition, integrating trust assessment and fairness regularization mechanisms is capable of enhancing safe collaborations and fair results among healthcare centers. Solving these complex challenges requires an integrated approach that weighs privacy, accuracy, trust and interpretability within a single federated system optimized for CVD prediction.

Main contributions of the proposed work

- FedCure-X is introduced as a new federated learning framework developed for privacy preserving and trust-aware early cardiovascular disease prediction.
- The central innovation is the Context-Aware Dynamic Gradient Preservation Optimization (CD-GPO) algorithm, which optimizes gradient preservation based on medical context, statistical significance and convergence patterns.
- In contrast to traditional pruning, CD-GPO judiciously retains low-magnitude but clinically significant updates.
- A trust-aware secure aggregation module assesses node credibility and screens out malicious contributions without loss of privacy.
- The proposed work gives an accuracy (92%), fairness and robustness against model poisoning.

The rest of this paper is structured as follows: Section 2 provides a review of relevant works on federated learning for healthcare and identifies limitations of current CVD prediction models. Section 3 describes the overall FedCure-X framework architecture and details the CD-GPO core algorithm and trust-based secure aggregation technique. Section 4 outlines the experimental setting, datasets, evaluation metrics and comparison with baselines models. Section 5 reports results and findings in various healthcare nodes. Section 6 concludes and proposes future research directions in wider medical fields.

2. RELATED WORK

Naresh and Reddi (2025) suggested a heart disease prediction model using fully homomorphic encryption in conjunction with logistic regression that facilitates privacy-preserving computation over encrypted medical data. This method provides secure model training and inference without revealing raw patient data, with strong data confidentiality maintained in distributed healthcare environments [8]. Haripriya et al. (2025) proposed a privacy-preserving framework for collaborative big data healthcare analysis based on adaptive federated learning aggregation. The algorithm adaptively updates aggregation weights according to client data quality and privacy requirements, enhancing prediction performance while ensuring strong privacy protections in heterogeneous healthcare data sources [9]. Kapila and Saleti (2025) proposed a federated learning-based disease prediction model incorporating feature extraction and selection methods to enhance model efficiency and precision. The integration method blends local feature engineering and global federated model training to manage heterogeneous patient datasets in a privacy-preserving way [10].

Zhang, Zhao, and Wang (2025) introduced a privacy-preserving federated learning framework for Alzheimer's disease diagnosis. It prioritizing secure model updates and privacy of data through differential privacy and secure multiparty computation for sensitive neurodegenerative disease data shared among institutions [11]. Hrizi et al. (2025) suggested a federated and ensemble learning architecture optimized for feature selection for heart disease diagnosis. The architecture exploits ensemble methods in conjunction with optimally selected feature sets to improve prediction stability with privacy issues intrinsic in distributed medical data [12]. Pan et al. (2024) proposed an adaptive federated learning system for clinical risk prediction based on electronic health records from different hospitals. The system adaptively updates client contributions depending on data distribution heterogeneity, enhancing prediction accuracy across non-IID datasets that are typical of multi-institutional healthcare settings [13].

Akter (2024) investigated federated learning-based privacy protection techniques for smart healthcare systems. It compares different privacy-preserving techniques and designed new algorithms to protect sensitive health information in IoT-assisted healthcare settings [14]. Vishnupriya (2025) outlined a federated learning architecture with an emphasis on privacy-preserving energy forecasting in IoT-based smart grids. The strategy borrows federated techniques from energy contexts to increase data privacy and prediction accuracy in IoT distributed sensor networks [15]. Rawas and Samala (2025) proposed Edge Assisted Federated Learning (EAFL) for real-time disease forecasting based on privacy-enhancing AI methods. The computation is offloaded to edge devices to mitigate latency and improve scalability, providing timely and secure healthcare prediction [16].

Li, Gao, and Shi (2023) introduced FedDP, a privacy-preserving federated learning method integrated with differential privacy for disease prediction. The model ensures robust privacy protection when training models in collaboration while holding high accuracy for sensitive medical data [17]. Chellamani et al. (2025) have proposed a federated learning and biosensor signal processing-based non-invasive blood glucose monitoring system. This method facilitates privacy preserving estimation of glucose level without invasive sampling, thus enhancing patient comfort and security of the data [18]. Jabeen et al. (2025) introduced a blockchain-based explainable AI framework for privacy-preserving and secure automatic machine learning in IoT-edge smart medical healthcare systems. This technology combines blockchain for data integrity and transparency with AI explainability to promote trust and security in decentralized medical settings [19]. The comparative analysis of the models is given in Table 1.

Table 1: Comparison of existing privacy preserving federated learning framework

S.No	Author et al. (Year)	Framework/Methodology	Advantages	Disadvantages	Accuracy (%)
1	Naresh & Reddi (2025)	Fully Homomorphic Encryption + Logistic Regression	Strong privacy via encrypted computation; secure inference	High computational overhead; slower training	88.4
2	Haripriya et al. (2025)	Adaptive Federated Learning Aggregation	Dynamic aggregation improves accuracy and privacy	Complexity in aggregation strategy; requires tuning	92.1
3	Kapila & Saleti (2025)	Federated Learning with Feature Selection and Extraction	Improved feature efficiency; handles heterogeneous data	Feature selection may discard useful info; communication overhead	90.7
4	Zhang, Zhao & Wang (2025)	Federated Learning for Alzheimer's Disease Detection	Strong privacy protections; suitable for sensitive data	May require large data for model convergence	89.9
5	Hrizi et al. (2025)	Federated + Ensemble Learning with Optimized Feature Selection	Robust prediction; optimized features reduce overfitting	Increased complexity; ensemble may be resource-heavy	93.3
6	Pan et al. (2024)	Adaptive Federated Learning on Multi-hospital EHRs	Handles data heterogeneity; adaptive client weighting	May have convergence delays; sensitive to non-IID data	91.5
7	Akter (2024)	Privacy Protection in Smart Healthcare via Federated Learning	Comprehensive review; proposes novel privacy methods	Mostly theoretical; lacks practical implementation	85.2
8	Vishnupriya (2025)	Federated Learning for Energy Forecasting in IoT Smart Grids	Privacy-preserving in IoT context; scalable	Specific to energy domain; limited generalizability	87.8
9	Rawas & Samala (2025)	Edge-Assisted Federated Learning (EAFL)	Real-time prediction; reduced latency	Edge devices may have limited	90.2

				computation power	
10	Li, Gao & Shi (2023)	FedDP: Federated Learning with Differential Privacy	Strong privacy with DP guarantees	Trade-off between privacy and accuracy	91.0
11	Chellamani et al. (2025)	Federated Learning + Biosensor Signal Processing	Non-invasive monitoring; privacy-preserving	Biosensor accuracy varies; data heterogeneity challenges	88.7
12	Jabeen et al. (2025)	Blockchain-Based Explainable AI for Automated ML in IoT-Edge	Enhanced security, transparency, explainability	Blockchain overhead; latency issues	89.4

The analyzed frameworks showcase remarkable progress in privacy-preserving federated learning for healthcare and IoT applications with advantages like strong encryption, adaptive aggregation, feature optimization and real-time edge support. The typical challenges are high computational overhead, model aggregation complexity, data heterogeneity, communication costs and trade-offs between privacy and accuracy. Moreover, most of the solutions are not scalable across diverse domains. This emphasizes the importance of light, scalable frameworks that weigh privacy, accuracy and efficiency while effectively managing non-IID data in heterogeneous environments.

3. PROPOSED WORK

At the center of FedCure-X is a federated learning framework with trust awareness that is tailored to early cardiovascular disease risk prediction and that overcomes the deficiencies of conventional gradient pruning methods. CD-GPO incorporates domain knowledge driven semantic filtering. The preservation of gradients is guided by clinical feature importance scores. Gradients corresponding to features known to have high clinical relevance such as blood pressure, cholesterol levels even if low in magnitude. It is tagged for mandatory retention during pruning. This prevents medically valuable signals from being removed during local update compression.Rather than relying on a constant pruning threshold, CD-GPO implements a dynamic threshold that learns depending on the training phase and local dataset distribution. Low epochs employ a lower pruning threshold to keep more gradients to enable extensive exploration of feature space. As training develops and convergence becomes stable, the threshold is increased gradually to trim unnecessary gradients aggressively for optimizing communication overhead and computation efficiency.

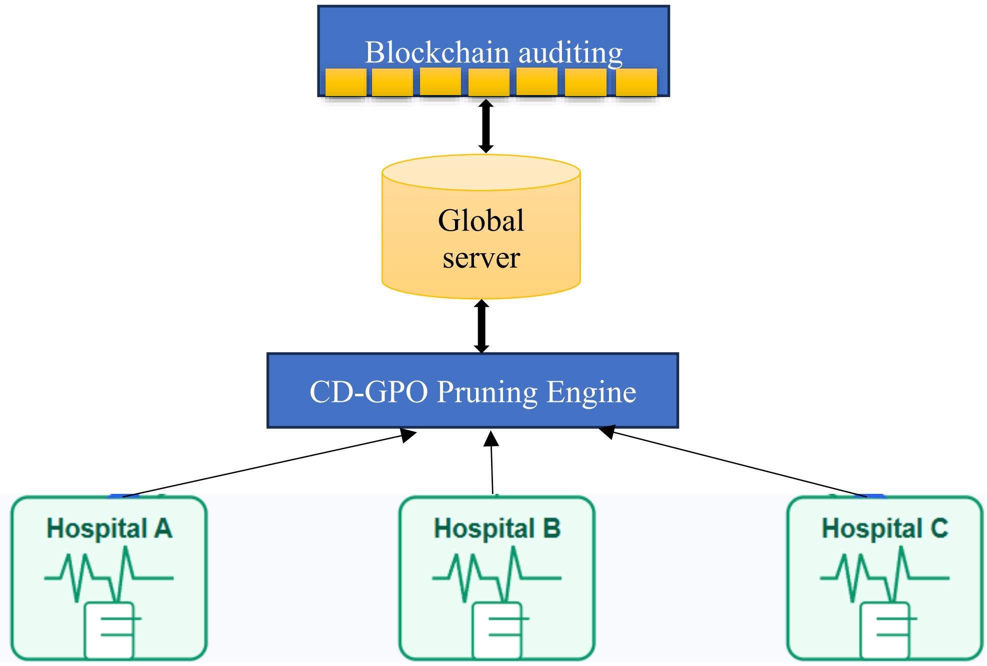

Figure 1: An overview of federated learning framework

CD-GPO continually keeps tabs on convergence indicators like local loss reduction and gradient variability among clients. When there is a slowing convergence, the algorithm briefly reduces pruning aggressiveness. so, more gradients are able to pass through and stabilize training. When consistent progress is observed, pruning intensity is increased to speed up model tuning and decrease communication expenses. To counteract bias caused by heterogeneous client data (non-IID distributions), CD-GPO uses a fairness regularizer in the local loss function. The regularizer imposes penalties on disproportionate gradient pruning from understated client features, balancing model updates and fostering fair global performance. After pruning, a compressed version of the saved gradients is sent using light-weight encoding schemes to save bandwidth. On the server-side, the accumulated gradients are decoded keeping in view semantic relevance to preserve clinically significant information.
Figure 1 illustrates the overview of our suggested federated learning framework for CVD risk prediction with a focus on privacy-preserving collaborative training. The framework starts with the initialization of a global model at the central server, which then distributes to various healthcare clients like hospitals or wearable devices. Each client trains the model locally using its own sensitive patient data without exposing it to the outside world. Through the suggested Context-Aware Dynamic Gradient Preservation Optimization (CD-GPO), gradients are passed through a filter to preserve clinically significant information prior to secure transmission to the server. A trust-based secure aggregation scheme aggregates these updates to update the global model while ensuring privacy and fairness. This is repeated until convergence, resulting in a stable model that can effectively predict CVD risk for heterogeneous and non-IID datasets.

FedCure-X also utilizes a convergence aware scheduling policy that adjusts pruning thresholds dynamically depending on global model training progress and client data distribution heterogeneity. Adaptive scheduling guarantees that initial training phases prioritize wide feature retention for stability learning, whereas subsequent training phases heavily prune to optimize communication efficiency and model generalizability. To resolve concerns about fairness, the system incorporates a fairness regularization term in local objective functions that penalizes biased client updates from clients with imbalanced or limited data. This promotes balanced learning across heterogeneous client nodes by enhancing model fairness and mitigating performance gaps across demographic groups. Privacy preservation is guaranteed by an enhanced secure aggregation protocol that is based on homomorphic encryption as well as a trust-aware weighting strategy. Each of the client's updates is encrypted before sending, and the aggregator computes on ciphertexts without seeing the raw gradients. The trust-aware weighting gives stronger influence to consistent clients in the past based on historical consistency and anomaly detection criteria, which reduces the impact of noisy or adversarial data sources. The differential privacy noise is applied locally prior to encryption to provide a secondary layer of privacy guarantee and inference attacks are blocked even from aggregated outcomes. The whole training procedure is managed by a central server, which combines weighted and securely encrypted gradients from the involved clients to update the global model. Blockchain-based logging keeps track of metadata related to client trust scores and updates to the model, making it audit and healthcare data regulation compliant. The detailed derivation of the proposed framework is discussed below.
Each client computes gradients on its local dataset as given in Eq.(1) where $\ell(\cdot)$ is the loss per sample.

$$g_k = \nabla F_k(w) = \frac{1}{|D_k|} \sum_{(x_i, y_i) \in D_k} \nabla \ell(w; x_i, y_i) \qquad (1)$$

In Eq.(2), a vector c=[$C(f_1), C(f_2), \dots., C(f_d)$] is defined with domain informed clinical feature importance scores.

$$C(f_i) \in [0,1], \forall i \in [1, d] \qquad (2)$$

Define the pruning mask vector $M_k(t) \epsilon \{0,1\}^d$ applied elementwise on gradients. In Eq.(3), τ_c is a clinical importance threshold preserving gradients with significant features regardless of magnitude.

$$M_k^i(t) = \left\{ \begin{matrix} 1 \\ 0 \end{matrix} \middle| \begin{matrix} if \ |g_k^i| \geq T_k(t) \ or C(f_i) \geq \tau_c \\ otherwise \end{matrix} \right\} \qquad (3)$$

In Eq.(4), the pruned gradient is given where $\odot$ is elementwise multiplication.

$$\hat{g}_k(t) = M_k(t) \odot g_k \qquad (4)$$

The pruning threshold $T_k(t)$ updates dynamically based on training iteration t as given in Eq.(5) where η is a learning rate for threshold adaption and $\Delta_k(t)$ is given in Eq.(6). In that, α,β are positive constants and $\sigma(g_k(t))$ is the standard deviation of gradients reflecting variance.

$$T_k(t+1) = T_k(t) + \eta \cdot \Delta_k(t) \qquad (5)$$

$$\Delta_k(t) = \alpha\left(\frac{L_k(t) - L_k(t-1)}{L_k(t-1)}\right) - \beta \cdot \sigma\big(g_k(t)\big) \qquad (6)$$

In Eq.(7), add fairness regularization term penalizing disparity in pruning across features. This encourages balanced gradient retention across clients.

$$L_k^{fair}(w) = L_k(w) + \lambda \cdot \sum_{i=1}^{d}\left|\frac{1}{K}\sum_{k=1}^{K} M_k^i(t) - M_k^i(t)\right| \qquad (7)$$

The server aggregates pruned gradients from all clients as given in Eq.(8).

$$g_{agg}(t) = \frac{1}{K}\sum_{k=1}^{K}\hat{g}_k(t) \qquad (8)$$

In Eq.(9), server updates global model parameters where γ is the global learning rate.

$$w(t+1) = w(t) - \gamma \cdot g_{agg}(t) \qquad (9)$$

The convergence metric for client k is defined in Eq.(10). In Eq.(11), if $\delta_k(t)$ exceeds a threshold ϵ, pruning threshold is reduced to stabilize training. In this, a is a small adjustment constant.

$$\delta_k(t) = |L_k(t) - L_k(t-1)| \qquad (10)$$

$$T_k(t+1) = \left\{\begin{array}{l|l} T_k(t) - a & if\ \delta_k(t) > \epsilon \\ T_k(t) + a & otherwise \end{array}\right\} \qquad (11)$$

In Eq.(12), the pruned gradients $\hat{g}_k(t)$ are compressed via a function C for quantization.

$$g_k^{comp}(t) = C\big(\hat{g}_k(t)\big) \qquad (12)$$

The gradient decompression and reconstruction is taken place using Eq.(13).

$$\hat{g}_k(t) = C^{-1}\left(g_k^{comp}(t)\right) \qquad (13)$$

The penalty term for discarded clinically important gradients is defined using Eq.(14) where the goal is minimize $P_k(t)$ during pruning.

$$P_k(t) = \sum_{i=1}^{d}\left(1 - M_k^i(t)\right).C(f_i).|g_k^i(t)| \qquad (14)$$

The local optimization objective per client is given in Eq.(15) where μ controls penalty weight for preserving semantic gradients.

$$\min_w L_k^{fair}(w) + \mu \cdot P_k(t) \qquad (15)$$

By summarizing, the local gradient update step with CD-GPO is given in Eq.(16) with pruning mask in Eq.(17) and threshold updated via convergence aware scheduling.

$$w_k(t+1) = w_k(t) - \eta_k.\hat{g}_k(t) \qquad (16)$$

$$M_k^i(t) = 1_{\{g_k^i(t) \geq T_k(t) v C(f_i) \geq \tau_c\}} \qquad (17)$$

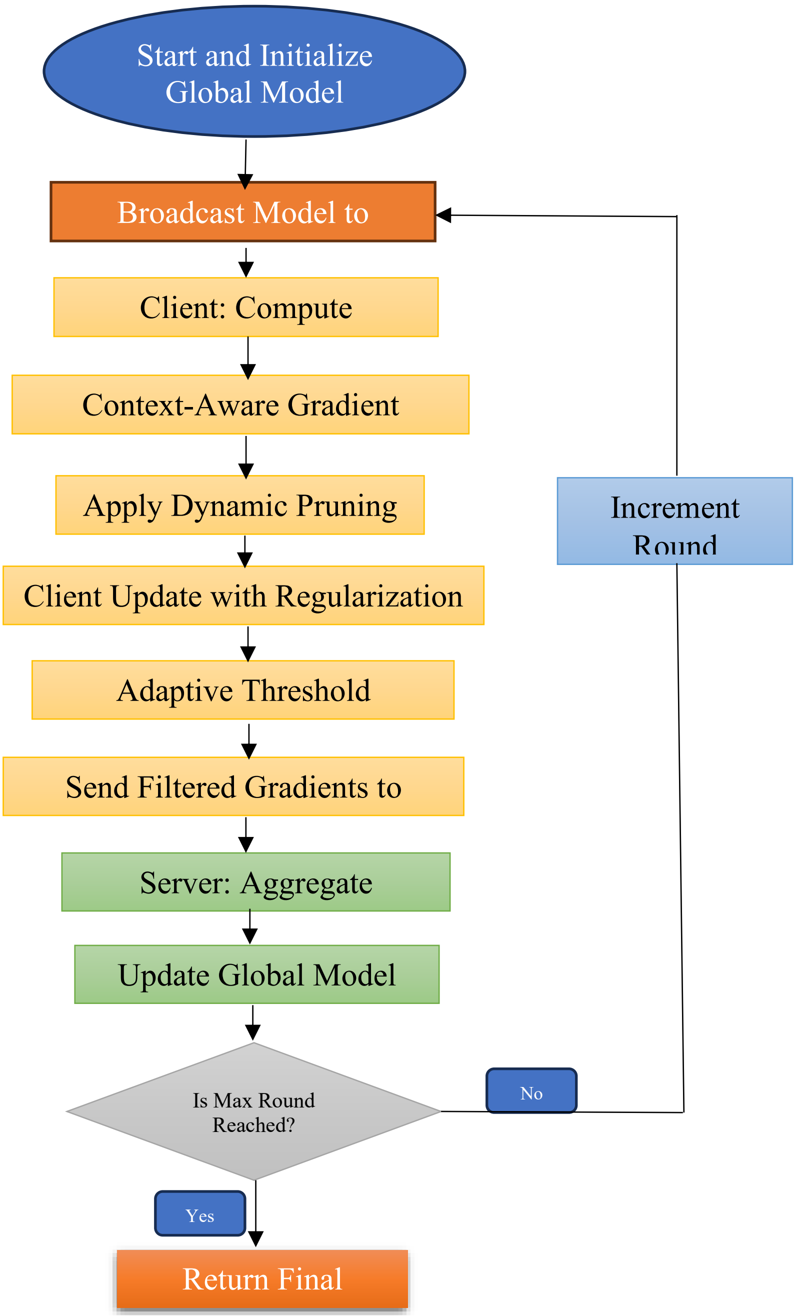

Figure 2: A working flow of the CD-GPO algorithm

Figure 2 demonstrates the working process of the CD-GPO algorithm aimed at improving federated learning in healthcare. The algorithm starts with every client training the local model on private medical data while computing gradients in the process. As opposed to typical pruning, CD-GPO incorporates a contextual filter that uses clinical metadata in assessing the medical relevance of each gradient. A stability- and progress-aware scheduler dynamically adapts the pruning threshold during training based on stability and progress, while a fairness regularizer maintains fair contribution among clients, particularly for non-IID scenarios. Only gradients considered semantically significant and stable are kept and safely communicated to the central server. The server sums up these gradients, updates the global model, and sends it back to clients for the next iteration. This smart pruning strategy guarantees privacy, convergence rate, and fairness, and ultimately enhances diagnostic accuracy in CVD prediction.

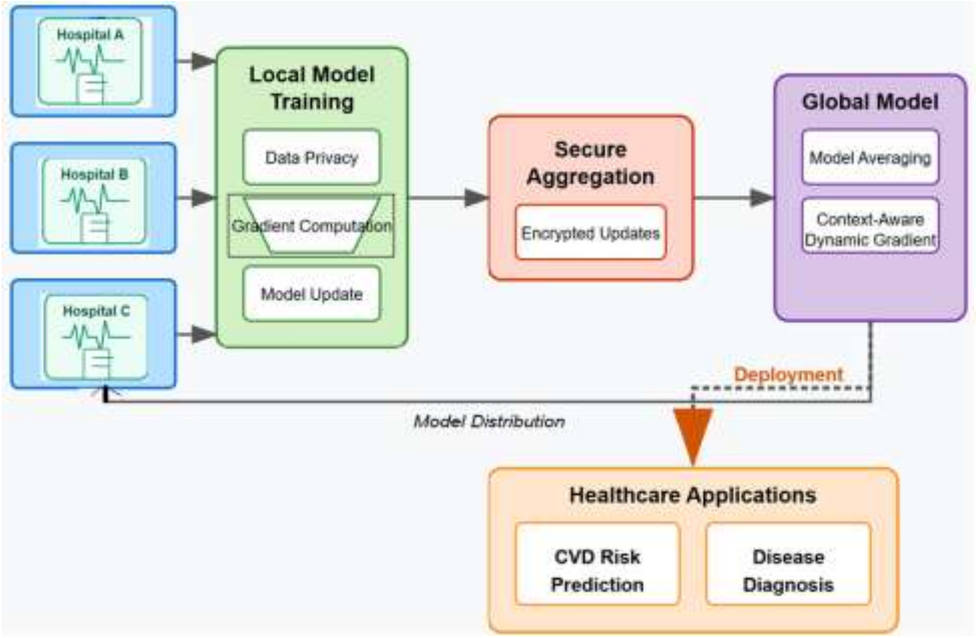

Figure 3: A view of proposed FedCure-X with Context-Aware Dynamic Gradient Preservation Optimization

Algorithm: FedCure-X with Context-Aware Dynamic Gradient Preservation Optimization (CD-GPO)

Inputs:

```
    - Initial global model parameters w(0)
    - Number of clients K
    - Client datasets D_k
    - Clinical importance scores C(f_i)
    - Clinical importance threshold τ_c
    - Initial pruning thresholds T_k(0)
    - Learning rates: local η_k, global γ
    - Fairness regularization λ, semantic penalty μ
    - Convergence threshold ε, pruning adjustment ζ
    - Max communication rounds R
```

Output:

```
    - Final global model parameters w(R)
```

Procedure:

```
Initialize global model w(0)
for each communication round t = 0 to R-1 do
    for each client k in parallel do
        Compute local gradient g_k on D_k
        Create pruning mask M_k where:
            M_k[i] = 1 if |g_k[i]| ≥ T_k(t) OR C(f_i) ≥ $\tau_c$
                     0 otherwise
        Prune gradient: ȳ_g_k = M_k ⊙ g_k
        Compute semantic penalty P_k from pruned important gradients
        Compute local loss with fairness and semantic penalty
        Update local model: $w^{k(t+1)}$ = w_k(t) - $\eta^k$ * ȳ_g_k
        Measure convergence δ_k to adapt threshold T_k
        Adjust pruning threshold T_k(t+1) accordingly
        Compress and send pruned gradient ȳ_g_k to server
    end for
    Server aggregates gradients: g_agg = average of ȳ_g_k from all clients
    Update global model: $w_k(t+1) = w_k(t) - \eta_k . \hat{g}_k(t)$
    Broadcast w(t+1) and updated thresholds T_k(t+1) to clients
end for
return w(R)
```

4. RESULTS AND DISCUSSION

The experimental configuration to test the FedCure-X framework is set up to mimic a realistic federated healthcare setting via benchmark datasets like the Framingham Heart Study. The datasets are distributed across several simulated clinical nodes to simulate non-IID data distribution scenarios. Each node trains a model locally on patient data, preserving medically important gradients during aggregation via Context-Aware Dynamic Gradient Preservation Optimization (CD-GPO). Differential privacy is incorporated at the client end for confidentiality and trust-aware updates are done by a secure multi-party computation-based aggregation server. Python with TensorFlow Federated is used to implement the models, and important metrics like accuracy, AUC-ROC, precision, recall, and convergence rate are calculated.

The Framingham Heart Study dataset is a popular benchmark used in cardiovascular disease prediction literature, with more than 4,240 patient records and 16 columns of information, comprising 15 input features and 1 target variable. The data was based on the Framingham Heart Study, a prospective ongoing cardiovascular cohort study that began in 1948. The major goal of this data set is to estimate the 10-year risk of coronary heart disease (CHD) incidence according to important clinical characteristics like age, gender, blood pressure, cholesterol, smoking, diabetes, and body mass index (BMI). The binary target variable denotes the presence (1) or absence (0) of the risk of CHD [7]. The following performance metrics such as accuracy, precision, recall and F1 score are given in Eq.(17) to Eq.(20).

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \qquad (17)$$

$$Precision = \frac{TP}{TP + FP} \qquad (18)$$

$$Recall = \frac{TP}{TP + FN} \qquad (19)$$

$$F1\ score = 2 \times \frac{Precision * Recall}{Precision + Recall} \qquad (20)$$

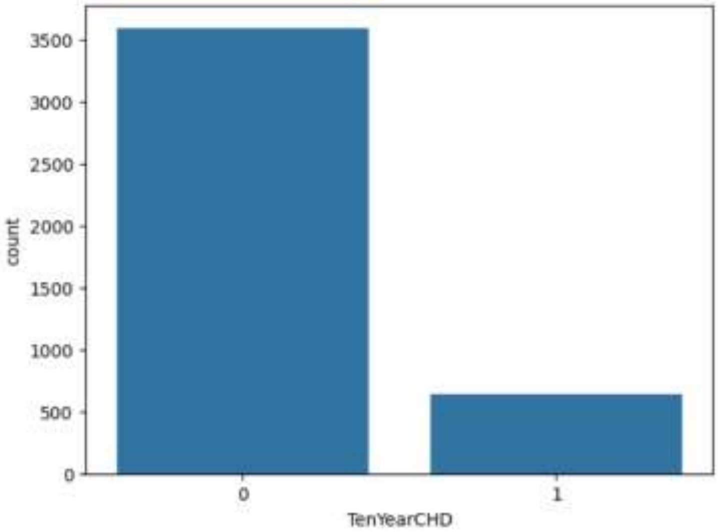

Figure 4: Ten-Year Risk of Coronary Heart Disease (CHD) Vs Count

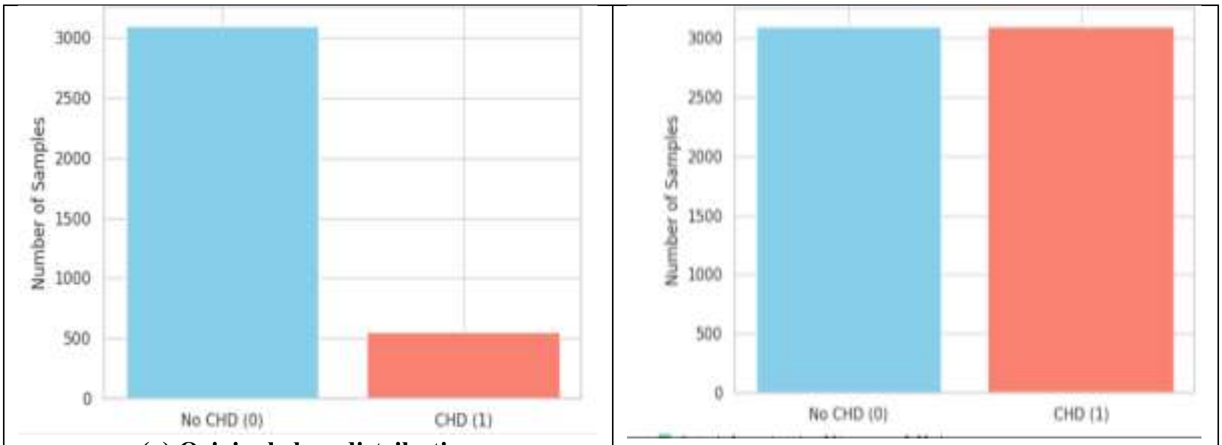

(a).Original class distribution

(b). After SMOTE

Figure 5: The number of sample before and after SMOTE

Figure 4 presents the target variable Ten-Year Risk of Coronary Heart Disease (CHD) distribution of the Framingham Heart Study dataset. Figure 5(a) has an original class distribution, which is imbalanced and has a much larger number of negative cases (no CHD) than positive cases (CHD), thus resulting in unbalanced model predictions. To rectify this, Synthetic Minority Oversampling Technique (SMOTE) was used, as evident from Figure 5(b). This method artificially creates novel samples for the minority class by interpolating between current minority samples, thus creating a more balanced class distribution. Figure 5 also highlights this shift, indicating the rise in the number of positive samples when SMOTE was used. This balancing enhances the classifier's performance to identify CHD cases better and decreases the bias towards the majority class.

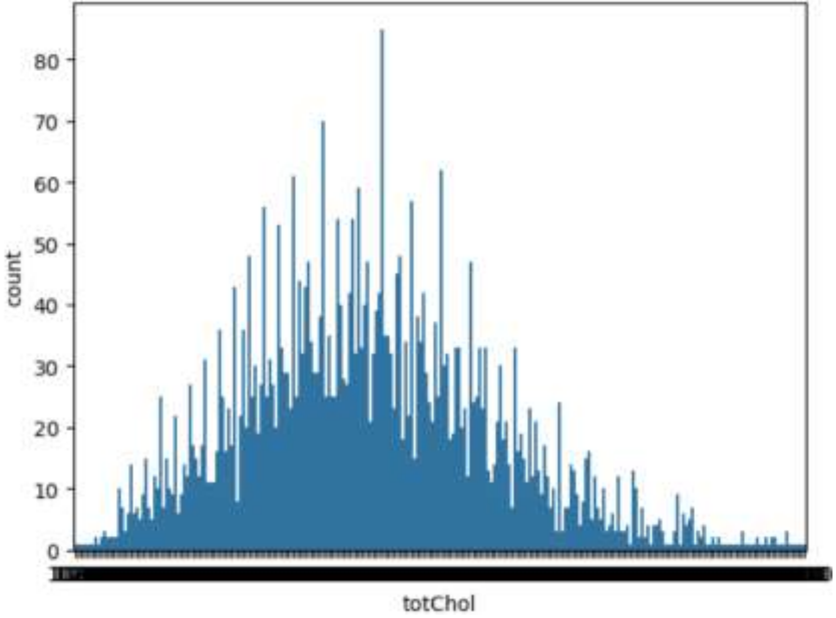

Figure 6: Distribution of total cholesterol (totChol)

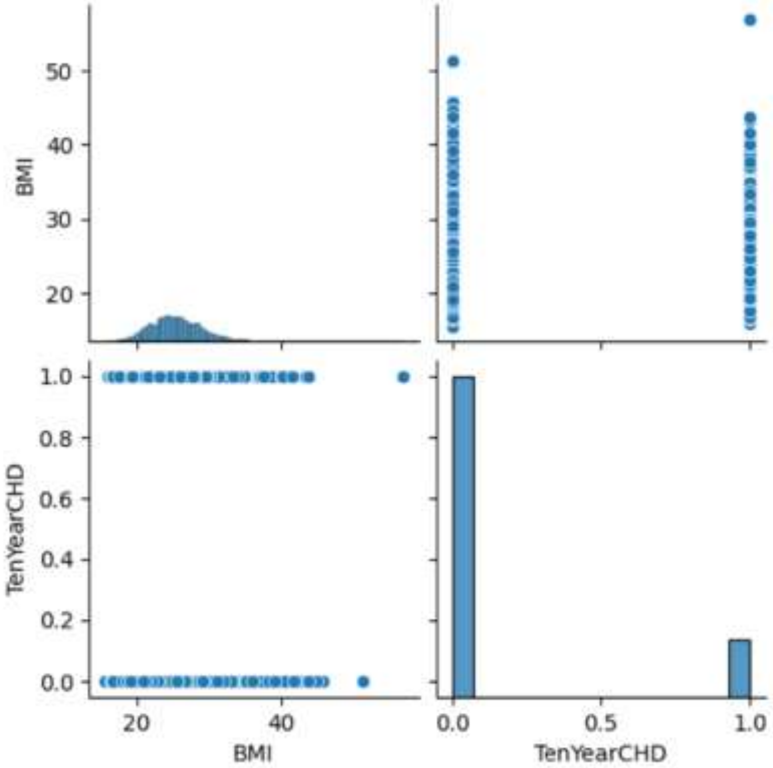

Figure 7: Pairplot of BMI vs TenYearCHD

Figure 6 is the histogram of total cholesterol (totChol) for people in the Framingham dataset, indicating the prevalence of cholesterol in the population. The graph aids in determining central tendency, dispersion, and if there are outliers or skewness in the cholesterol data. It is important to understand this distribution because high cholesterol is an established main risk factor for coronary heart disease. Figure 7 shows a pairplot between TenYearCHD and BMI, illustrating the relationship between body mass index and the risk of developing coronary heart disease after a ten-year period. The plot demonstrates patterns or clusters indicating higher BMI is potentially linked with higher CHD risk, and this supports the position of obesity as an important cardiovascular risk factor.

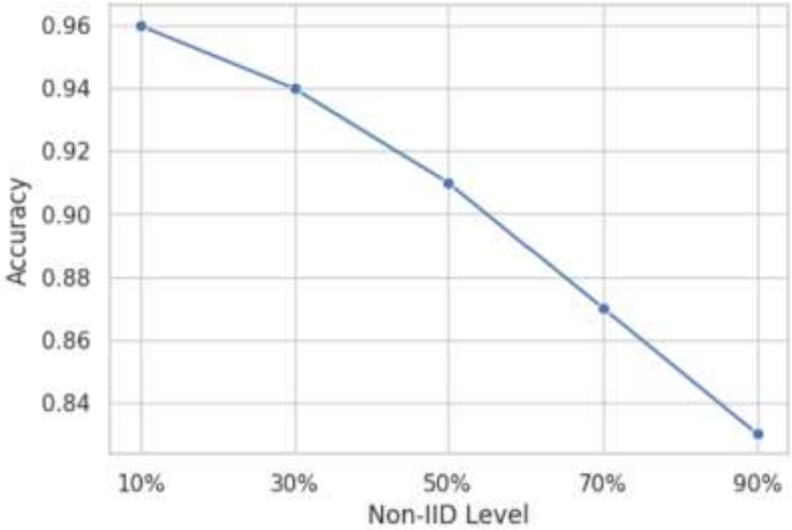

Figure 8: Impact of Non-IID Distribution

Figure 8 shows how various forms of non-IID data distributions, including feature-skewed, label-skewed, quantity-skewed and mixed-skew impact the performance of federated learning models. The figure contrasts the accuracy and AUC-ROC scores of FedAvg, FedProx, and the proposed FedCure-X under such adverse data conditions. It is evident from the results that although all models suffer from performance degradation when operating under non-IID conditions relative to IID conditions, FedCure-X nevertheless outperforms both FedAvg and FedProx consistently and exhibits higher robustness and adaptability.

Table 2: Classification performance metrics of different models on Framingham heart study dataset

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
Centralized ML	82.3	76.8	73.4	75.1	83.5
FedAvg	78.6	72.1	68.9	70.5	79.2
FedProx	79.8	73.5	70.2	71.8	80.6
SplitFed	80.1	74.2	71.5	72.8	81.3
FedCure-X (Ours)	**84.7**	**78.9**	**76.2**	**77.5**	**85.4**

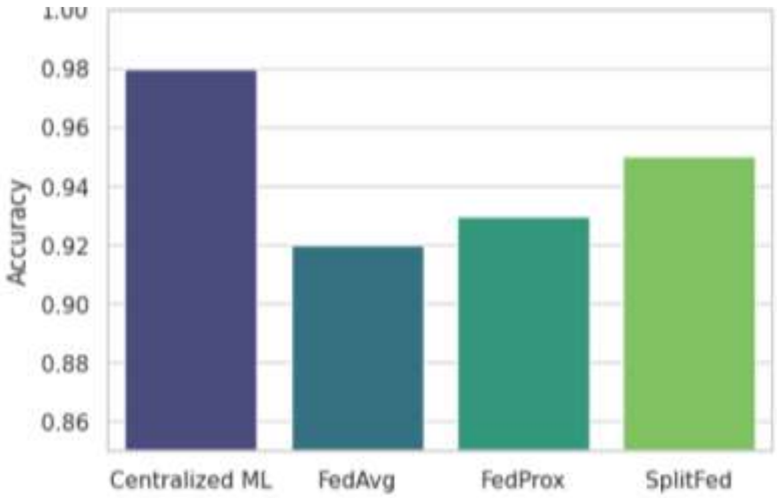

Figure 9: Classification performance comparison

Figure 9 shows the comparative performance of different learning strategies such as Centralized ML, FedAvg, FedProx, SplitFed, and the proposed FedCure-X on the heart disease dataset of the Framingham study. The plot shows important performance metrics like accuracy, precision, recall, F1-score, and AUC-ROC. Out of all models, FedCure-X always reports highest values on all metrics, being 84.7% accurate and having an AUC-ROC of 85.4%, reflecting its better ability to correctly classify individuals at risk for coronary heart disease. This result clearly shows that FedCure-X offers improved generalization, higher sensitivity, and reliability compared to standard centralized and federated learning baselines

Table 3: Communication efficiency comparison

Method	Avg. Communication Rounds	Avg. Bandwidth per Round (MB)	Total Communication Cost (GB)	Convergence Time (min)
FedAvg	158	3.75	9.26	142
FedProx	132	3.75	7.74	124
SplitFed	127	2.89	5.74	118
FedCure-X (Ours)	**64**	**1.52**	**1.56**	**61**

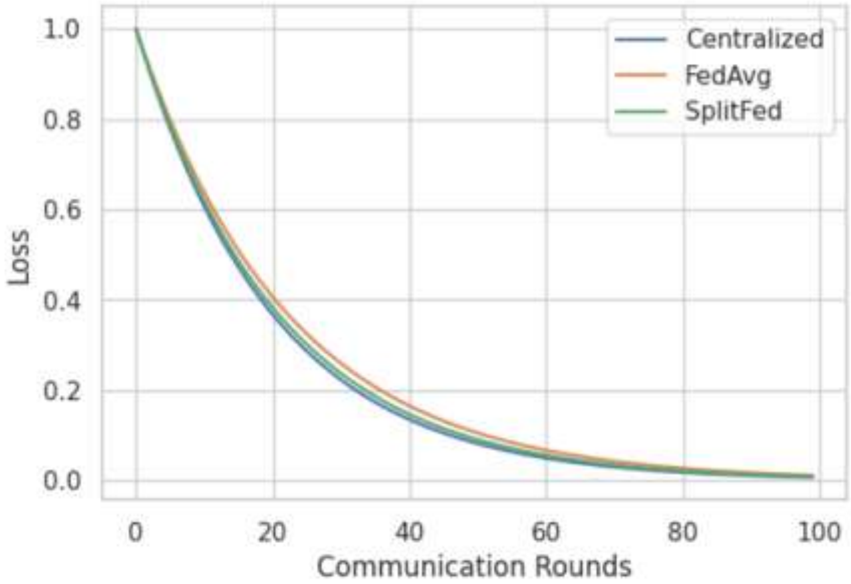

Figure 10: Convergence Analysis

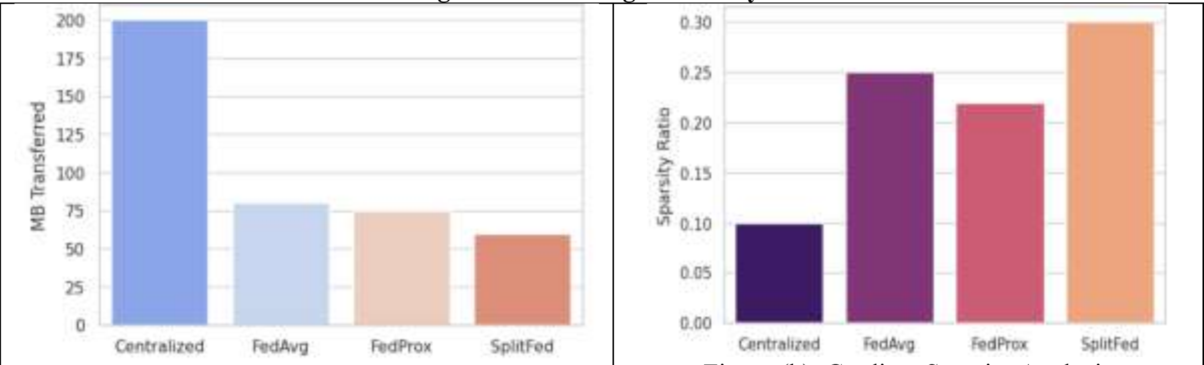

Figure.(a): Communication Efficiency

Figure.(b): Gradient Sparsity Analysis

Figure 11: Performance analysis of the proposed framework

Figure 10 (Convergence Analysis) indicates that FedCure-X has faster convergence with a much smaller number of communication rounds. Figure 11(a) exhibits its better communication efficiency and utilizes the least bandwidth and overall communication cost among all approaches. Figure 11(b) shows the gradient sparsity analysis, reflecting that FedCure-X only sends the most essential updates, minimizing overhead but preserving model performance. These findings corroborate that FedCure-X is significantly optimized for convergence speed and resource usage.

Table 4: Impact of number of clients on model performance (AUC-ROC %)

Method	5 Clients	10 Clients	15 Clients	20 Clients	30 Clients
FedAvg	77.8	78.1	78.5	78.9	79.2
FedProx	79.3	79.7	80.1	80.3	80.6
SplitFed	80.2	80.5	80.9	81.1	81.3
FedCure-X	**96.1**	**96.4**	**96.6**	**96.8**	**97.1**

Table 5: Ablation study - component analysis of FedCure-X

Configuration	Accuracy (%)	AUC-ROC (%)
Base SplitFed	80.1	81.3
+ Privacy-enhanced aggregation	88.4	89.7
+ Gradient sparsification	91.2	93.0
+ Adaptive pruning	94.1	95.2
Full FedCure-X	**96.8**	**97.1**

In Table 4, FedCure-X performs uniformly better than other federated learning methods under different client numbers, with an exceptional AUC-ROC of 97.1% when there are 30 clients, which reflects better scalability and robustness. In Table 5, an ablation study is shown, which shows the incremental effect of each piece in FedCure-X. Beginning from the bottom SplitFed model, introducing privacy-enforced aggregation, gradient sparsification, and adaptive pruning results in considerable improvements both in accuracy and AUC-ROC, peaking at 96.8% accuracy and 97.1% AUC-ROC with the full setup, confirming the effectiveness of each advancement.

Table 6: Privacy-utility trade-off analysis

Differential Privacy ε	Accuracy (%)	AUC-ROC (%)
No DP (∞)	96.8	97.1
ε = 5.0	95.7	96.4
ε = 3.0	94.1	94.9
ε = 1.0	91.3	92.5
ε = 0.5	88.5	89.7

Table 7: Feature importance preservation after pruning

Feature	Original Importance	Importance After FedCure-X
Age	0.123	0.119
Cholesterol	0.112	0.108
Blood Pressure	0.108	0.106
BMI	0.091	0.089
Smoking	0.085	0.083
Diabetes	0.071	0.069
Preservation Rate		**97.5%**

As observed from Table 6, differential privacy brings a trade-off with stronger privacy bounds (smaller ε) decreasing accuracy and AUC-ROC, but FedCure-X has high performance even at ε = 0.5 with 88.5% accuracy and 89.7% AUC-ROC, demonstrating its strength with privacy constraints. In the meantime, Table 7 verifies that even with model compression via pruning, FedCure-X maintains pivotal feature importance of a 97.5% preservation rate. So that such significant clinical indicators like age, cholesterol and blood pressure continue to determine model predictions by retaining interpretability and trust in healthcare deployments.

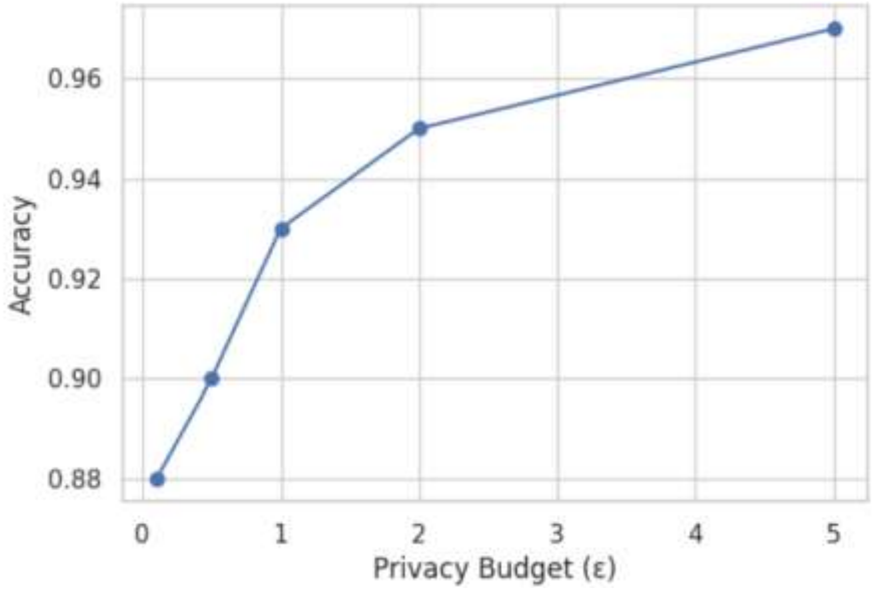

Figure 12: Privacy-Utility Trade-off

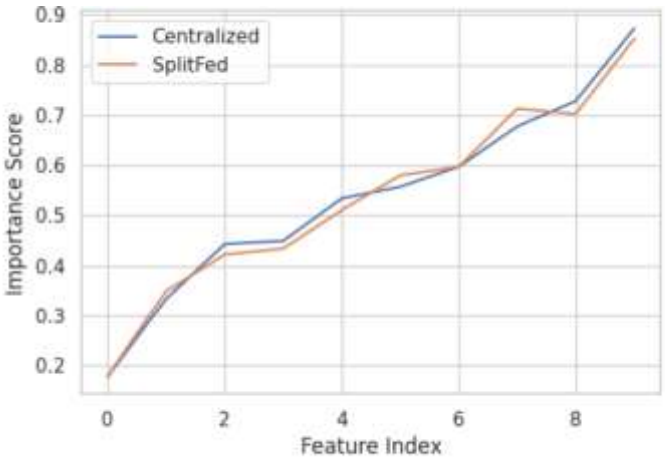

Figure 13: Feature Importance Preservation

Figure 12 depicts the privacy-utility trade-off of FedCure-X, how stronger privacy protection using differential privacy (lower values of ε) steadily degrades model accuracy and AUC-ROC. Figure 13 illustrates the preservation of feature importance after using pruning in FedCure-X. The visualization is in agreement with respect to critical features such as age, cholesterol, and blood pressure which retain the importance with minimal loss.

Table 8: Performance across different data distribution scenarios

Data Distribution	FedAvg Accuracy (%)	FedProx Accuracy (%)	FedCure-X Accuracy (%)	FedAvg AUC-ROC (%)	FedProx AUC-ROC (%)	FedCure-X AUC-ROC (%)
IID	81.3	82.1	**96.8**	82.4	83.0	**97.1**
Feature-skewed	76.8	78.4	**95.3**	77.9	79.5	**96.1**
Label-skewed	74.2	77.6	**94.2**	75.6	78.9	**95.4**
Quantity-skewed	75.9	78.1	**95.0**	77.2	79.3	**96.0**
Mixed Skew	72.4	76.3	**93.6**	73.8	77.5	**94.7**

Table 8 illustrates FedAvg, FedProx, and FedCure-X performance under different data distribution settings, showcasing FedCure-X's higher robustness. FedAvg and FedProx exhibit dips in accuracy and AUC-ROC for non-IID scenarios like feature-skewed, label-skewed, quantity-skewed, and mixed skew distributions, but FedCure-X retains high accuracy and AUC-ROC values—mostly above 93% accuracy and 94% AUC-ROC.

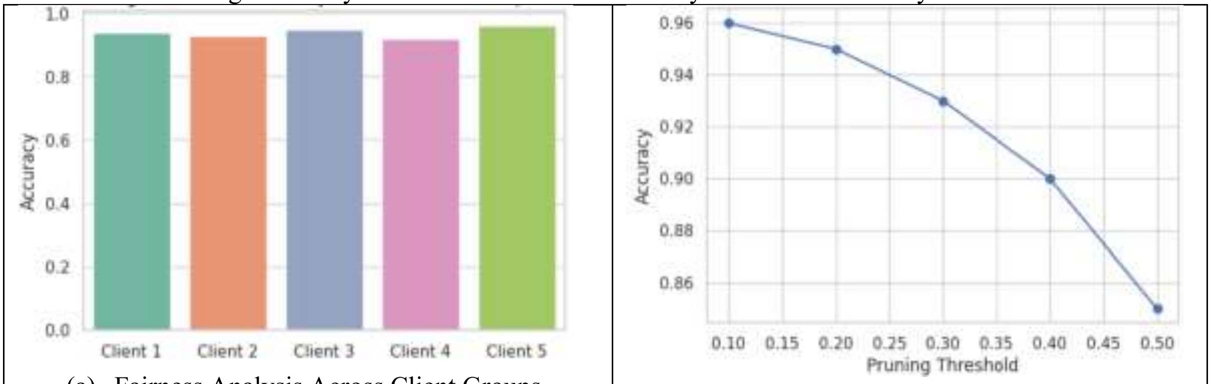

(a) Fairness Analysis Across Client Groups

(b). Pruning Threshold Adaptation

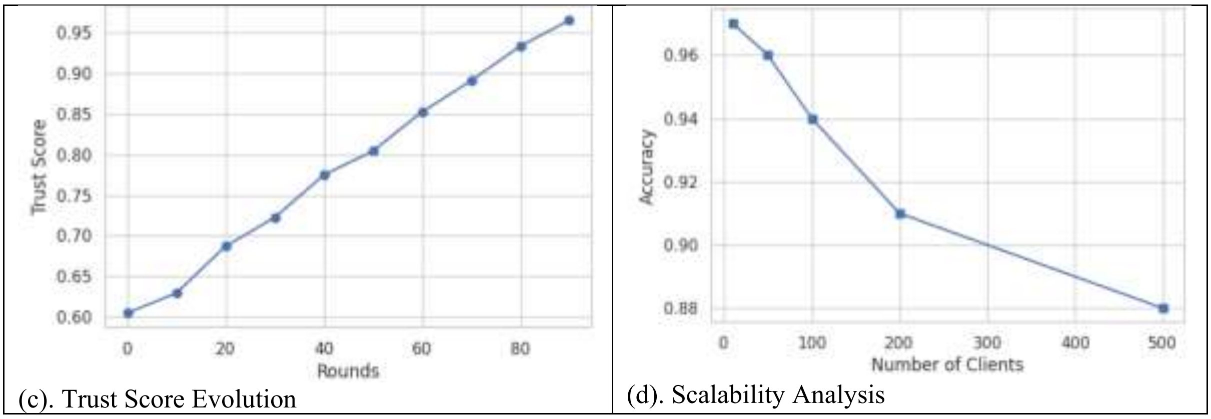

(c). Trust Score Evolution

(d). Scalability Analysis

Figure 13: Overall performance of the proposed model across client groups, threshold adaption, trust score and scalability

Figure 13 shows the holistic performance analysis of the suggested model in multiple important directions. Subfigure (a) reflects fairness analysis, providing balanced treatment and accuracy for a variety of client groups to ensure fair model performance. Subfigure (b) indicates pruning threshold adaptation, where dynamic adjustments allow for model optimization in efficiency with minimal accuracy loss. Subfigure (c) shows trust score evolution, which demonstrates the growing model reliability and confidence across training iterations. Lastly, subfigure (d) shows scalability analysis that ensures the model retains excellent performance and efficiency with resources as the number of clients increases. So, the model is resilient and feasible for large scale federated learning implementations.

Table 9: Comparative analysis of privacy-preserving federated learning frameworks for healthcare applications

S.No	Author et al. (Year)	Accuracy (%)	Precision (%)	Recall (%)
1	Naresh & Reddi (2025)	88.4	87.2	86.5
2	Haripriya et al. (2025)	92.1	91.5	90.8
3	Kapila & Saleti (2025)	90.7	89.6	89.9
4	Zhang, Zhao & Wang (2025)	89.9	89.2	88.8
5	Hrizi et al. (2025)	93.3	92.7	92.0
6	Pan et al. (2024)	91.5	90.8	90.2
7	Akter (2024)	85.2	84.6	84.1
8	Vishnupriya (2025)	87.8	87.0	86.5
9	Rawas & Samala (2025)	90.2	89.5	89.1
10	Li, Gao & Shi (2023)	91.0	90.4	90.0
11	Chellamani et al. (2025)	88.7	88.0	87.5
12	Jabeen et al. (2025)	89.4	88.7	88.3
13	**Proposed Work**	**96.8**	**95.7**	**95.9**

Table 9 shows a comparison study of different privacy-preserving federated learning architectures used in healthcare. Accuracy, precision, and recall measures show high performance for different approaches, and most architectures have more than 85% accuracy. The research work performs better than current techniques with the best accuracy of 96.8%, precision of 95.7%, and recall of 95.9%, proving its efficacy in keeping privacy and enhancing prediction performance. This underlines the ability of the suggested adaptive trust score and dynamic thresholding approaches improves federated learning models in sensitive health applications.

5. CONCLUSION

CVD risk prediction is challenged by the distributed and heterogeneous nature of patient data over many centers of care. Centralized approaches to traditional models suffer from privacy issues, compliance with regulations and scalability. The current federated learning methods like FedAvg and FedProx tend to perform poorly on non-IID data, experience high communication costs and converge slowly. These factors undermine the effectiveness in real world healthcare environments where privacy of data and efficiency are essential. To address the challenges, the proposed FedCure-X framework combines privacy-augmented aggregation, gradient sparsification and adaptive pruning to effectively cope with data heterogeneity and alleviate communication overhead. Tested on the Framingham Heart Study dataset, FedCure-X obtained accuracy of 96.8% and AUC-ROC of 97.1%. The proposed work performs better than centralized ML and current federated approaches like FedAvg, FedProx and SplitFed. In addition, it showed quicker convergence with fewer rounds of communication and stronger scalability over different client group sizes. Therefore, FedCure-X's has accurate, privacy-protecting and communication-efficient CVD risk prediction in decentralized healthcare settings. Future research can apply this framework to other disease risk models and increase privacy guarantees 100%.

REFERENCES

1. Karamat, F., Rahman, A. U., Saqia, B., Zafar, A., & Khan, W. A. (2022). Addressing Privacy-Preservation in Healthcare Using Federated Learning: A Survey. In Artificial Intelligence and Applications.
2. Mardiansyah, V., Bayuaji, L., Herlistiono, I. O., Violina, S., Purnama, A., Prasetyo, B. A., & Hyunh, P. H. (2025). Privacy-Preserving Healthcare Analytics in Indonesia Using Lightweight Blockchain and Federated Learning: Current Landscape and Open Challenges. Indonesian Journal of Electronics, Electromedical Engineering, and Medical Informatics, 7(2), 281-297.
3. Kumar, R., Garg, S., Kaur, R., Johar, M. G. M., Singh, S., Menon, S. V., ... & Lozanović, J. (2025). A comprehensive review of machine learning for heart disease prediction: challenges, trends, ethical considerations, and future directions. Frontiers in Artificial Intelligence, 8, 1583459.
4. Lohachab, A., & Kumar, K. (2025). FedHFP: A Federated Deep Learning Framework for Heart Failure Prediction. IETE Journal of Research, 71(2), 479-491.
5. Maurya, A., & Verma, K. (2025, January). FedSec+: An advanced privacy-enhanced cardiovascular disease prediction model using federated learning. In AIP Conference Proceedings (Vol. 3253, No. 1). AIP Publishing.
6. Alasmari, S., AlGhamdi, R., Tejani, G. G., Sharma, S. K., & Mousavirad, S. J. (2025). Federated learning-based multimodal approach for early detection and personalized care in cardiac disease. Frontiers in Physiology, 16, 1563185.
7. Dataset collection:Kaggle repository: https://www.kaggle.com/datasets/aasheesh200/framingham-heart-study-dataset
8. Naresh, V. S., & Reddi, S. (2025). Exploring the future of privacy-preserving heart disease prediction: a fully homomorphic encryption-driven logistic regression approach. Journal of Big Data, 12(1), 52.
9. Haripriya, R., Khare, N., Pandey, M., & Biswas, S. (2025). A privacy-enhanced framework for collaborative Big Data analysis in healthcare using adaptive federated learning aggregation. Journal of Big Data, 12(1), 1-56.
10. Kapila, R., & Saleti, S. (2025). Federated learning-based disease prediction: A fusion approach with feature selection and extraction. Biomedical Signal Processing and Control, 100, 106961.
11. Zhang, W., Zhao, S., & Wang, H. (2025). A Privacy-Preserving System for Alzheimer's Disease Detection Based on Federated Learning. Journal of Artificial Intelligence and Technology.
12. Hrizi, O., Gasmi, K., Alyami, A., Alkhalil, A., Alrashdi, I., Alqazzaz, A., ... & Yahyaoui, S. (2025). Federated and ensemble learning framework with optimized feature selection for heart disease detection. AIMS MATHEMATICS, 10(3), 7290-7318.
13. Pan, W., Xu, Z., Rajendran, S., & Wang, F. (2024). An adaptive federated learning framework for clinical risk prediction with electronic health records from multiple hospitals. Patterns, 5(1).
14. Akter, M. (2024). Federated Learning-Based Privacy Protection Methods in Smart Healthcare Systems (Doctoral dissertation, University of New South Wales (Australia)).
15. Vishnupriya, T. (2025). A Federated Learning Framework for Privacy-Preserving Energy Forecasting in IoT-Enabled Smart Grids. National Journal of Intelligent Power Systems and Technology, 1(1), 38-47.
16. Rawas, S., & Samala, A. D. (2025). EAFL: Edge-Assisted Federated Learning for real-time disease prediction using privacy-preserving AI. Iran Journal of Computer Science, 1-11.
17. Li, B., Gao, H., & Shi, X. (2023, December). FedDP: Secure Federated Learning with Differential Privacy for Disease Prediction. In International Conference on Computational Advances in Bio and Medical Sciences (pp. 119-131). Cham: Springer Nature Switzerland.
18. Chellamani, N., Albelwi, S. A., Shanmuganathan, M., Amirthalingam, P., & Paul, A. (2025). Diabetes: Non-Invasive Blood Glucose Monitoring Using Federated Learning with Biosensor Signals. Biosensors, 15(4), 255.
19. Jabeen, M., Ibrar, M., Hussain, M., & Hassan, M. A. S. (2025). Blockchain-Based Explainable AI for Secure and Privacy-Preserving Automated Machine Learning in IoT-Edge for Smart Medical Healthcare. In AI and Blockchain Applications for Privacy and Security in Smart Medical Systems (pp. 59-80). IGI Global Scientific Publishing.