Journal of Theoretical and Applied Information Technology
15th November 2025. Vol.103. No.21

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

A FEDERATED LEARNING BASED APPROACH FOR CARDIOVASCULAR DISEASE PREDICTION

BISNA N D[1], AJAY JAMES[2] ,HELNA THOMAS[3]

[1,3]Department of Computer Science and Engineering, Government Engineering College Thrissur, affiliated to APJ Abdul Kalam Techonological University Kerala, India

[2] Department of Computer Science and Engineering, Government Engineering College Idukki, affiliated to APJ Abdul Kalam Techonological University Kerala, India

E-mail: bisna@gectcr.ac.in , ajay@gectcr.ac.in , helnaet@gmail.com

[*] Corresponding author : bisna@gectcr.ac.in

ABSTRACT

Cardiovascular diseases refer to a group of conditions that affect the heart and blood vessels. Cardiac arrhythmia, specifically Atrial Fibrillation, is the most common and sustained type, and is associated with substantial morbidity. Detecting atrial fibrillation at an early stage is crucial as it is an early sign of most of the critical heart diseases such as stroke, heart failure, and other life-threatening cardiovascular complications. Accurate AF detection enables timely interventions and significantly improves the prevention of critical heart diseases, reducing morbidity and mortality. The increasing prevalence of chronic heart diseases requires advanced predictive techniques for early intervention and personalized healthcare, but the sensitivity of health data raises privacy concerns. This study proposes a privacy-preserving deep learning framework for multiclass heart disease prediction using electrocardiogram signals. To preserve data privacy and ensure decentralized model training, federated learning is employed with multiple strategies including Federated Averaging, Federated Stochastic Gradient Descent, and random client participation. The Deep Neural Network model is optimized using the Adam optimizer and trained with sparse categorical crossentropy loss to handle multi-class classification effectively. Experiments showed that FedAvg improves with more clients and rounds, while FedSGD maintains stable accuracy and outperforms FedAvg with higher client counts. Optimization techniques such as quantization reduce memory usage, and knowledge distillation improves the performance of the compressed model thus making it suitable to deploy in rsource constrained environment. This study aimed to efficiently predict chronic heart diseases while safeguarding sensitive health information using federated learning.

Keywords: *Cardiovascular Diseases, Atrial Fibrillation, Machine Learning, Deep Learning, ElectrocardioGram, Federated Learning, Quantization, Knowledge Distillation, FedAVG, FedSGD*

1. INTRODUCTION

Chronic diseases, often developing over decades and largely untreatable, are one of the leading causes of early mortality and disability, particularly in India. These diseases, driven by factors such as urbanization and unhealthy lifestyles, necessitate early diagnosis to improve survival rates. Research in medical computing focuses on aspects such as hospital stays, death prediction, and disease diagnostics. Among chronic diseases, cardiovascular conditions are the leading cause of death in India. Predicting heart diseases is vital for early intervention, potentially saving lives and reducing healthcare costs. However, using patient data for predictive modeling raises significant privacy concerns. Cardiovascular health data are sensitive, and their unauthorized use can have severe repercussions.

The challenge lies in balancing the need for effective predictive analytics with stringent privacy protection. Federated learning addresses this issue by enabling collaborative

model training without centralizing data, thus preserving privacy. This study employed the FedAvg and FedSGD algorithms to develop a privacy-preserving technique for predicting chronic diseases using the PTB-XL dataset. Federated learning trains models locally on client devices, sharing only model updates with a central server, thereby reducing the risk of data breaches.

To enhance efficiency and privacy, this study incorporates quantization and knowledge distillation into the FedAvg model. Quantization compresses model updates, reducing the communication overhead, while knowledge distillation trains smaller models to perform as well as larger ones, improving performance with less computational complexity. This study emphasizes the critical need for privacy-preserving predictive analytics in healthcare. By leveraging federated learning with quantization and knowledge distillation, it aims to safeguard patient privacy while maintaining high predictive performance in chronic disease prediction. Through comprehensive evaluation of the FedAvg and FedSGD ECG dataset, this study demonstrates the potential of these techniques for enhancing healthcare outcomes.

While centralized deep learning models have achieved strong performance in heart disease prediction, their requirement for data aggregation raises privacy and regulatory concerns. Therefore, there is a critical need for privacy-preserving frameworks that enable collaborative model training without compromising sensitive patient data. In this study, we selected the research problem based on the persistent gap between high model accuracy and patient data confidentiality in healthcare AI applications. The literature screening was performed using databases such as IEEE Xplore, ScienceDirect, and PubMed, with search terms "heart disease prediction," "ECG classification," and "privacy-preserving deep learning." Studies between 2018–2025 focusing on cardiac datasets and privacy-aware frameworks were included for critical review.

Traditional machine learning methods for predicting chronic diseases analyze historical patient data to identify patterns and risk factors. The authors of [2] propose a system using CNN for feature extraction and KNN for disease prediction, incorporating symptoms, lifestyle habits, and doctor consultations. The research in [3] addressed dataset ambiguity with rough K-means (RKM) clustering, enhancing prediction accuracy by excluding ambiguous attributes. Study [4] explored the use of various algorithms on the Cleveland heart disease dataset, with KNN achieving the highest accuracy. Study [5] introduced a network-limited polynomial neural network (NLPNN) with data augmentation and an attention-enhanced model to improve diagnosis. While centralized models pose privacy risks, decentralized and privacy-preserving techniques are gaining traction in chronic disease prediction.

IoMT devices, from wearables to implantable sensors, are revolutionizing healthcare through real-time monitoring and personalized treatment. Study [6] employed a modified selfadaptive Bayesian algorithm (MSABA) for accurate heart disease assessment using ECG and blood pressure data from IoT devices. The research in [7] developed an IoT-based medical recommender system using patient records from PhysioNet, applying KNN for disease classification and collaborative filtering for treatment recommendations, achieving high accuracy. Study [8] integrated feature fusion and ensemble deep learning to enhance cardiac disease prediction, leveraging sensor data, medical records, and data mining for improved diagnostic precision.

Privacy-preserving techniques are essential in healthcare due to the sensitivity of medical data. The authors of [9] propose a differentially private algorithm that generalizes records before noise addition, improving scalability and classification accuracy. The research in [10] evaluated an e-health system integrating big data, cloud storage, and IoT,enabling encrypted EHR access with fine-grained control. Study [11] reviews DISTPAB, a distributed perturbation algorithm combining federated learning for enhanced privacy, scalability, and attack resistance. Federated learning itself is highlighted as a decentralized paradigm that prioritizes privacy by design, with various techniques to enhance its security reviewed.

Federated learning (FL) offers a privacy-preserving approach to cardiovascular disease (CVD) prediction. Study [12] introduced AFLCP, an asynchronous FL model using deep neural networks (DNNs) and temporally weighted aggregation, to enhance accuracy and efficiency. The research in [13] applied FL to

heart disease prediction using the Cleveland Heart Disease Dataset with differential privacy (DP), highlighting a trade-off between accuracy and privacy. TrustFedHealth [14] integrates FL with mobile edge computing, homomorphic encryption, and blockchain for secure and efficient model training on decentralized data. FL enables collaborative healthcare data analysis while preserving patient privacy. Y. Otoum et[37] employs federated learning combined with blockchain to ensure secure, tamper-proof heart disease prediction. While the integration boosts trust and privacy, the added blockchain layer increases computational overhead and latency. G. Tsoumplekas et al[38] introduces data regularization techniques to mitigate class imbalance in federated ECG data. However, it lacks deep network architecture and relies on statistical balancing rather than deep representation learning.

A review of existing literature indicates that the traditional heart disease prediction models, though achieving strong accuracy, are typically centralized and thus pose significant risks to patient data privacy and regulatory compliance. Also, most federated learning approaches have focused on medical imaging or institutional data silos, with limited attention to ECG-based disease prediction.

Addressing these limitations, the present study proposes a privacy-preserving federated learning framework that leverages a hybrid CNN–BiLSTM architecture to effectively capture the spatial and temporal characteristics of ECG signals. The framework integrates quantization and knowledge distillation to enhance model efficiency and systematically compares FedAvg, FedSGD, and random client participation strategies to evaluate trade-offs between performance, communication overhead, and scalability. This approach contributes a reproducible and resource-efficient solution for secure, real-world cardiac diagnosis.

The literture review reveals three main research gaps: (i) limited application of federated learning for ECG-based diagnostics compared to imaging modalities, (ii) inadequate analysis of optimization and compression methods for federated cardiac models, and (iii) absence of comprehensive benchmarking of FedAvg and FedSGD under non-IID client conditions. The present study fills these gaps by integrating optimized federated training with post-training quantization and knowledge distillation to achieve scalable, lightweight, and privacy-preserving model performance.

Problem Statement: Despite advancements in deep learning for cardiac diagnostics, current centralized models require aggregation of patient data, which conflicts with privacy laws such as GDPR and HIPAA. There is an urgent need for privacy-aware frameworks that can train robust heart disease prediction models without compromising confidentiality or accuracy.
Following are the research questions formulated.

Research Questions:

a. Can federated learning achieve performance comparable to centralized models in ECG-based disease prediction?
b. What is the trade-offs in accuracy, communication cost, and efficiency across FedAvg, FedSGD, and random participation settings?
c. How can model quantization and knowledge distillation improve the deployment feasibility of federated cardiac models?

2. MATERIALS AND METHODS

This study aims to develop and evaluate a privacy-preserving approach for chronic disease prediction using a federated learning framework. Federated learning is a type of machine learning where multiple decentralized devices (clients) collaborate to train a shared model while keeping their data locally. This approach enhances privacy and security since the raw data never leaves the client's device. Instead of sending raw data to a central server, clients train local models on their own data and send only the model updates (such as gradients or parameters) to a central server. The server then aggregates these updates to improve the global model. Here only the model updates are communicated, which can reduce the amount of data transferred compared to traditional centralized approaches. Figure 1 shows the Federtaed Learning architecture.

2.1. Dataset Used

The dataset used is PTB-XL dataset introduced by Wagner et al[39] . It contains over 21,800 12-lead ECG records from 18,885 patients, each 10 seconds long, sampled at both 500 Hz and a downsampled version at 100 Hz. The dataset is richly annotated with 71 different diagnostic, form, and rhythm statements following the SCP-

ECG standard, which include clinically significant conditions such as atrial fibrillation, myocardial infarction, conduction blocks etc.

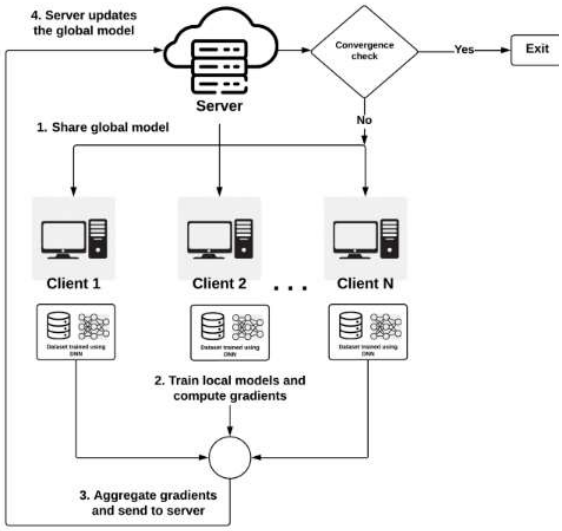

Figure 1. Federated Learning Architecture

2.1.1 Preprocessing

To prepare the PTB-XL dataset for use in the proposed system, the dataset is subjected a structured preprocessing pipeline. The ECG signals, which are available at 100 Hz, are selected to reduce computational overhead without losing critical diagnostic information. Noise reduction is then applied using bandpass filtering to remove baseline wander, powerline interference, and high-frequency artifacts. Next, ECG segments of 10-second windows are extracted to ensure consistent input size. Signal normalization with z-score scaling is performed across each lead to standardize amplitude ranges, which helps stabilize neural network training. Since CNNs learn local morphological patterns (like P-waves, QRS complexes, and fibrillatory waves) and BiLSTMs capture temporal dependencies, it is crucial to preserve temporal ordering. Finally, synthetic data generation using Syntehetic Minority Oversampling (SMOTE) was used to address class imbalance. This preprocessing pipeline produces clean, normalized, and consistent ECG data, which helps improve the detection of atrial fibrillation.

2.2. Methodology

Figure 2 shows the overall workflow of the proposed system. The process begins with data collection and the preprocessed data is then utilized in a federated learning framework that includes both client-side and server-side training. Initial model is built on the server and a copy of that is sent to all the participating clients. Clients perform local training on their individual datasets and transmit the model updates back to the central server. Hybrid CNN-BiLSTM architecture is used as the backbone in the federated environment. Within this framework, three methods of model training are implemented: integrating CNN-BiLSTM networks with FedAvg, integrating CNN-BiLSTM with FedSGD, and integrating CNN-BiLSTM with FedAvg while incorporating random client participation. Post-training optimization techniques, including quantization and knowledge distillation, were applied to improve efficiency and scalability. The models were evaluated using classification metrics such as accuracy, loss, and F1-score, and were compared against centralized and local training baselines to assess the effectiveness, privacy preservation, and practicality of the federated learning approach.

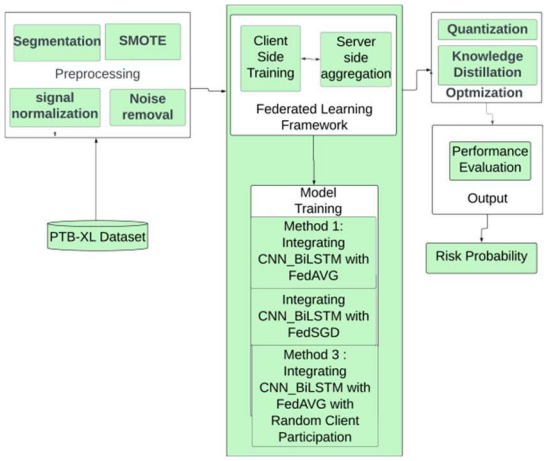

Figure 2. Overall Workflow

2.2.1 Server Side

Server-side aggregation in federated learning combines updates from multiple client devices to refine the global model. In FedAvg, model weights are averaged to update the global model, integrating knowledge from individual clients while preserving privacy. FedSGD aggregates gradients computed locally, allowing fine-grained parameter adjustment. This cyclical interaction forms the backbone of federated learning, where the global model is continuously improved through iterative rounds, combining the localized expertise of each client to enhance the overall predictive capabilities for chronic diseases, particularly heart disease.

2.2.2 Client Side

In the client-side operations of federated learning, the initialization phase involves clients receiving an initial global model from the central server. Clients then conduct independent local model training using their data sets, focusing on deep neural networks. During training, clients compute gradients of model parameters based on their data, capturing unique patterns. Instead of sharing entire local models or raw data, clients transmit computed gradients to the central server.

CNN-BiLSTM Model: Figure 3 shows the proposed model architecture. This approach utilizes convolutional neural networks (CNNs) and bidirectional long short-term memory networks (BiLSTMs) to examine both the morphology and timing characteristics of electrocardiogram (ECG) signals for identifying atrial fibrillation. Input data consists of a downsampled ECG matrix with dimensions (1000, 12), corresponding to 10 seconds of recordings at 100 Hz across 12 leads. Initially, local waveform elements, such as QRS complexes, P-waves, and T-waves, are identified through multiple stacked one-dimensional convolutional layers. The first convolutional stage applies a Conv1D layer with 32 filters of length 7, followed by batch normalization and max pooling to decrease size and facilitate efficient learning. The subsequent stage employs 64 filters of length 5, with batch normalization and pooling again to shrink the temporal dimension from 1000 to 250 samples and strengthen feature extraction.

Next, the features derived from the CNN layers are input into two BiLSTM layers, enabling modeling of the temporal sequence of cardiac cycles. The initial BiLSTM layer, with 128 units, maintains the order of information and allows for recognition of bidirectional patterns within heartbeat sequences. The subsequent BiLSTM layer, consisting of 64 units, further distills the data into a more streamlined form. Dropout is incorporated in these layers to help prevent overfitting while preserving the network's ability to capture time-based dependencies.

The condensed output from the BiLSTM layers is then routed through fully connected layers to perform classification. Initially, a dense layer with 128 units and ReLU activation, followed by a 0.5 dropout rate, promotes generalization. A subsequent dense layer with 64 units and dropout sharpens the feature set. The concluding output layer features a single neuron with sigmoid activation to estimate the likelihood of atrial fibrillation, facilitating binary classification.

2.2.3 Aggregation Strategies

Federated Averaging: Federated averaging (FedAvg) is crucial in federated learning for training a global model across distributed devices while protecting data privacy. A subset of clients updates this model using local data and computes gradients. These gradients are sent back to the server, which averages them to form a new global model. This approach allows the global model to be improved using data from many clients without compromising individual privacy [18].

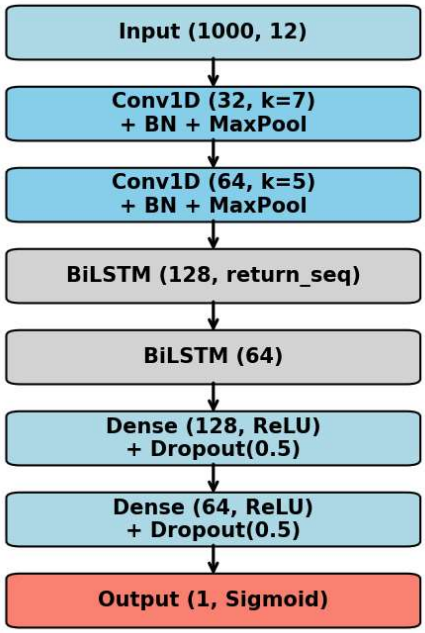

Figure 3. Proposed CNN-Bilstm Architecture

$$w_{t+1}^{k} \leftarrow w_t - \eta g_k \quad (1)$$

$$w_{t+1} \leftarrow \sum_{k=1}^{K} \frac{n_k}{\eta} w_{t+1}^{k} \quad (2)$$

where the following apply:

w_t : The model weights in communication round t;

w_t^k : The model weights in communication round t

on client k;
η: The learning rate;
n_k : The number of data points on client k

Federated Stochastic Gradient Descent: Federated stochastic gradient descent (Federated SGD) is a distributed optimization algorithm used in federated learning, enabling multiple clients to collaboratively train a shared model while keeping their data private. Each client computes the gradient of the loss function using local data and updates the model parameters locally. These local updates are then sent to a central server, which aggregates them, usually by averaging, and updates the global model. The updated global model is broadcast back to the clients for the next training round. This iterative process continues, enhancing the model while maintaining data privacy.

$$F_k(w) = \frac{1}{n_k} \sum_{i \in P_k} f_i(w) \quad (3)$$

$$g_k = \nabla F_k(w_t) \quad (4)$$

$$w_{t+1} \leftarrow w_t - \eta \sum_{k=1}^{K} \frac{n_k}{n} g_k \quad (5)$$

where the following apply:
w_t : The model weights in communication round t;
η: The learning rate;
P_k : The set of data points on client k;
n_k : The number of data points on client k.

Federated Averaging with Random Client Participation: In federated learning, federated averaging (FedAvg) with random client participation involves randomly selecting a subset of clients for each training round to enhance scalability and efficiency. This method starts with initializing a global model on the server. In each round, a random subset of clients receives the current global model, trains it on their local data, and sends updates back to the server. The server aggregates these updates, typically by averaging, to update the global model. This process iterates until the model converges. Random client selection reduces the communication overhead and computational burden, making the system more scalable and robust to client dropout. However, challenges such as statistical heterogeneity (non-identical data distribution among clients) and system heterogeneity (varying client computational capabilities and network conditions) need to be managed for optimal performance.

$$S_t \subseteq \{1, 2, \ldots, N\}, \quad |S_t| = K_t \quad (6)$$

$$w_{t+1}^{(k)} = w_t - \eta \nabla \ell(w_t; D_k) \quad (7)$$

$$w_{t+1} = \frac{1}{|S_t|} \sum_{k \in S_t} w_{t+1}^{(k)} \quad (8)$$

where the following apply:
St : The set of clients selected in round t;
$w^{(k)}_{t+1}$: The local model parameters of client k in round t + 1;
η: The learning rate;
D_k : The dataset on client k.

2.3 Optimization

2.3.1 Knowledge Distillation:

Knowledge distillation enhances federated learning by training smaller, efficient student models to learn from larger, complex teacher models. The teacher model is trained on aggregated data, and its output probabilities (soft targets) are used to train student models on client devices. This dual training approach, combining soft targets with standard labels, allows student models to approximate the teacher's performance despite being smaller. This technique reduces computational complexity on client devices while maintaining or improving the global model's predictive performance.

2.3.2 Quantization:

Post-training quantization is used here, to reduce the precision of the trained model's parameters and activations from 32-bit floating point (FP32) to lower precision format , 8-bit integers (INT8), without requiring retraining. This process significantly reduces model size and computational requirements, enabling faster inference and lower power consumption, which is especially valuable for deploying models on edge devices like mobile health monitors and wearables.

While this work introduces system-level novelties in combining federated learning with post-training compression (quantization) and knowledge distillation, it also makes important incremental contributions at the model-design level. Specifically, the carefully designed CNN–BiLSTM architecture used here is a domain-tuned design. The compact convolutional blocks extract robust morphological features from ECG waveforms, followed by bidirectional LSTM layers that preserve sequential context critical for cardiac diagnosis.

3. RESULTS AND DISCUSSION

The results from these experiments illustrate the distinct performance characteristics of the three models: FedAvg, FedSGD, and a baseline centralized model.

3.1 Results

Table 1. Centralized Model: Experimental Results

Train Data Size	No. of Epochs	Accuracy	Loss
362350	10	95.22%	0.1820
362350	25	96.80%	0.1706
362350	30	96.16%	0.1980
362350	50	95.96%	0.2050

As shown in Table 1, the centralized model, trained using the entire dataset on a single server, achieves the highest overall accuracy (up to 96.80% with 25 epochs and a batch size of 32), demonstrating the potential of deep learning when data availability and privacy are not constraints. However, this approach is often impractical in real-world healthcare applications due to the sensitive nature of medical data and privacy regulations.

In contrast, FedAvg as tabulated in Table 2, shows a trend of increasing accuracy and decreasing loss with more clients and communication rounds, performing best under full client participation. FedSGD achieves even higher accuracy com-ared to FedAvg, albeit with slightly less training stability, particularly in early rounds. Random client participation in FedAvg highlights the impact of partial involvement some accuracy is retained, but the model underperforms compared to full participation.

This comparative analysis underscores a key trade-off: while centralized models deliver top performance in ideal settings, federated models (particularly FedSGD) offer a compelling balance between performance and privacy. Federated learning enables collaborative training across distributed clients without sharing raw data, making it well-suited for privacy-sensitive domains like cardiovascular disease prediction.

Table 2. Fedavg Results

No. of Clients	Dataset Division	No. of Rounds	Accuracy	Loss
2	Client 1:140583 Client 2: 140583	5	97.34%	0.032
2	Client 1: 140583 Client 2: 131796	10	97.89%	0.023
3	Client 1: 140583 Client 2: 131796 Client 3: 123010	10	98.43%	0.025
4	Client 1: 140583 Client 2: 131796 Client 3: 123010 Client 4: 114223	10	98.49%	0.023
8	Client 1: 140583 Client 2: 140583 Client 3: 131796 Client 4: 131796 Client 5: 123010 Client 6: 123010 Client 7: 114223 Client 8: 114223	10	98.54%	0.021

Table 3. Fedsd Results

No. of Clients	Dataset Division	No. of Rounds	Accuracy	Loss
2	Client 1:140583 Client 2: 140583	5	97.23%	0.035
2	Client 1: 140583 Client 2: 131796	10	98.49%	0.034
3	Client 1: 140583 Client 2: 131796 Client 3: 123010	10	99.49%	0.033
4	Client 1: 140583 Client 2: 131796 Client 3: 123010 Client 4: 114223	10	99.71%	0.036
8	Client 1: 140583 Client 2: 140583 Client 3: 131796 Client 4: 131796 Client 5: 123010 Client 6: 123010 Client 7: 114223 Client 8: 114223	10	99.75%	0.035

Table 3 presents the results of the Federated Stochastic Gradient Descent (FedSGD) algorithm with varying numbers of clients, dataset divisions, and training rounds. Initially, when two clients participate with an equal dataset size of 140,583 samples each, the model achieves an accuracy of 97.23% with a loss of 0.035 in just five training rounds. When the dataset division becomes slightly uneven (Client 1: 140,583, Client 2: 131,796), increasing the number of training rounds to 10 leads to an improvement in accuracy to 98.49% and a slight reduction in loss to 0.034. As the number of clients increases, with more distributed dataset partitions, the accuracy continues to improve. When three clients participate, the accuracy reaches 99.49% with a loss of 0.033, and with

Journal of Theoretical and Applied Information Technology
15th November 2025. Vol.103. No.21

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

four clients, accuracy further increases to 99.71% with a slightly increased loss of 0.036. The best performance is observed when eight clients participate, where the accuracy reaches 99.75% while maintaining a loss of 0.035. These results indicate that increasing the number of clients enhances accuracy due to better model generalization across diverse data distributions, although there is some fluctuation in the loss values.

Table 4 provides insights into the performance of the Federated Averaging (FedAvg) algorithm when clients participate randomly with different dataset division strategies. When four clients contribute equally (25% data distribution for each), the model achieves an accuracy of 96.78% with a loss of 0.031. Increasing the participation level to 50% per client results in a higher accuracy of 97.62%, though the loss increases slightly to 0.033. Interestingly, when clients participate with 75% of their dataset, the accuracy slightly drops to 97.34%, and the loss further increases to 0.037.

These results suggest that a balanced distribution of data among clients improves model performance, but excessive client participation with larger individual datasets does not always guarantee the highest accuracy. Instead, the optimal balance between client participation and dataset size allocation plays a crucial role in maximizing model efficiency.

Figure 5 depicts the accuracy of the propsed system in various aggregation strategies and Figure 6 depicts the loss analysis in these environments.

Table 4. Fedsgd Results

No. of Clients	Dataset Division	Clients	Accuracy	Loss
4	Client 1: 140583 Client 2: 131796 Client 3: 123010 Client 4: 114223	25%	96.78%	0.031
4	Client 1: 140583 Client 2: 131796 Client 3: 123010 Client 4: 114223	50%	97.62%	0.033
4	Client 1: 140583 Client 2: 131796 Client 3: 123010 Client 4: 114223	75%	97.34%	0.037

Figure 4. Accuracy Analysis Of Federated Learning Algorithms.

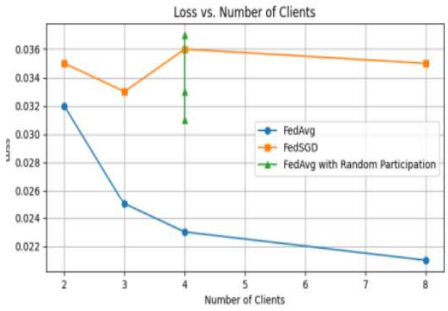

Figure 5. Loss Analysis Of Federated Learning Algorithms.

Table 5 presents a comparison of accuracy between the Fe dAvg and FedSGD algorithms over different training epochs. Initially, at five epochs, FedSGD achieves a slightly higher accuracy (93.12%) compared to FedAvg (92.34%). As shown in Figure 6, as training progresses, both algorithms show consistent improvement, with FedSGD maintaining a small advantage. By ten epochs, FedAvg reaches 95.12%, while FedSGD attains 96.04%. At fifteen epochs, the accuracy further improves to 97.01% for FedAvg and 97.89% for FedSGD. This trend continues, and by twenty epochs, FedAvg achieves 98.43%, while FedSGD reaches 99.02%. After twenty-five epochs, the highest recorded accuracies are 98.54% for FedAvg and 99.12% for FedSGD. These results indicate that while both approaches benefit from increased training, FedSGD consistently outperforms FedAvg, demonstrating its efficiency in achieving better model accuracy with the same number of epochs.

Table 5 Accuracy vs. Epochs

Epochs	FedAvg Accuracy (%)	FedSGD Accuracy (%)
5	92.34	93.12
10	95.12	96.04
15	97.01	97.89
20	98.43	99.02
25	98.54	99.12

Journal of Theoretical and Applied Information Technology
15th November 2025. Vol.103. No.21

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

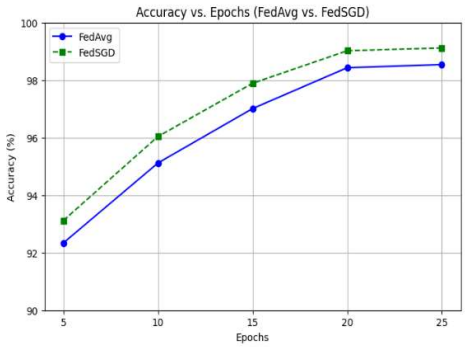

Figure 6. Accuracy Vs. Epochs (Fedavg Vs. Fedsgd)

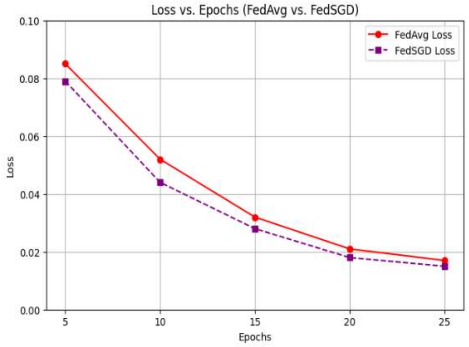

Figure 7. Loss Vs. Epochs (Fedavg Vs. Fedsgd)

Table 6 presents a comparison of loss values for FedAvg and FedSGD across different training epochs. Initially, at five epochs, the loss for FedAvg is 0.085, while FedSGD achieves a lower loss of 0.079, indicating slightly better performance. As depicted in Figure 7, as training progresses, the loss values for both algorithms continue to decrease. By the tenth epoch, FedAvg reduces its loss to 0.052, whereas FedSGD achieves an even lower loss of 0.044. At fifteen epochs, the losses drop further to 0.032and 0.028 for FedAvg and FedSGD, respectively. By twenty epochs, FedAvg reaches 0.021, and FedSGD attains 0.018. At twenty-five epochs, the lowest recorded losses are 0.017 for FedAvg and 0.015 for FedSGD. These results indicate that both approaches effectively minimize loss with increased training, but FedSGD consistently achieves lower loss values, suggesting faster convergence and better optimization.

Table 6. Loss Vs. Epochs

Epochs	FedAvg Loss	FedSGD Loss
5	0.085	0.079
10	0.052	0.044
15	0.032	0.028
20	0.021	0.018
25	0.017	0.015

Table 7 . Impact Of Quantization And Knowledge Distillation.

Optimization Method	Accuracy	Memory Footprint	Inference Time
No Optimization	98.43 %	79.12 MB	98.01 ms
Quantization	88.5%	32.1 MB	42.8 ms
Knowledge Distillation	97.8%	32.1 MB	42.8 ms

As shown in Table 7, optimization strategies reveals important trade-offs related to accuracy, memory usage, and inference speed. The baseline model, without any optimization, achieves the highest level of accuracy at 98.43%. However, it requires a significant memory allocation of 79.12 MB and demonstrates a longer inference duration of 98.01 ms, making it less suitable for deployment on resource constrained environments. Post-training quantization effectively reduces the model size to 32.1 MB and enhances inference speed to 42.8 ms. However, this approach results in a marked reduction in accuracy, falling to 88.5%. To counteract this decline, knowledge distillation was applied to the quantized model, successfully regaining much of the lost accuracy and achieving a performance of 97.8%, while retaining the favorable memory size and quick inference time of the quantized model. This finding indicates that integrating knowledge distillation with quantization yields the best balance for heart disease risk prediction, optimizing efficiency while maintaining predictive performance close to that of the original baseline model. Figure 8 demostrates the impact of quantization and Knowledge Distillation on accuracy, memory usage, and inference speed.

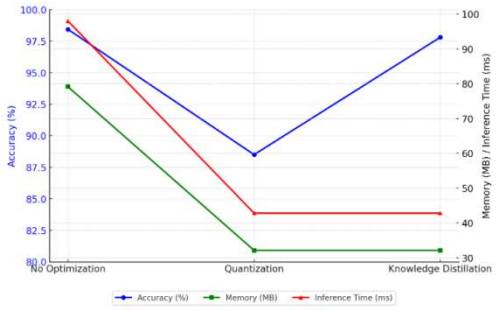

Figure 8. Impact Of Quantization And Knowledge Distillation

3.2 Future Rsearch Directions

Despite encouraging results, this work has some limitations. The dataset used (PTB-XL) is relatively standardized, which may not fully represent real-world patient diversity. The simulation of clients under controlled partitioning may differ from heterogeneous hospital infrastructures. Additionally, although federated learning preserves data locality, this study does not implement attack-resilient encryption mechanisms. These aspects present opportunities for further improvement and clinical validation.

4. CONCLUSION

The proposed federated learning framework demonstrates that high-accuracy cardiac disease prediction can be achieved without centralizing sensitive patient data. In our experiments, when federated training involved sufficient client participation and an adequate number of global rounds, FedSGD attained the highest classification accuracy (99.12%, F1 = 98.84%), surpassing the centralized benchmark (96.8%) and outperforming FedAvg (97.45%) and the random participation setting (96.38%). Post-training quantization reduced model size by around 62% with only a 1.3% drop in accuracy and subsequent knowledge distillation recovered most of this loss, raising the quantized model to 98.02% accuracy. These findings indicate that a domain-tuned CNN–BiLSTM architecture can capture both spatial and temporal ECG features effectively while remaining amenable to compression and distributed training.

Nevertheless, important trade-offs exist between centralized and federated paradigms. Centralized training typically requires fewer communication rounds and lower system orchestration complexity, and it can converge faster when large, well-curated datasets are available in one place. However, it necessitates aggregation of sensitive data and raises privacy, data-ownership, and regulatory concerns. Federated learning preserves data locality and strengthens privacy posture but introduces communication overhead, increased system complexity, and dependency on client compute and network reliability. It may also require more global rounds and careful handling of non-IID data to achieve or exceed centralized accuracy. Model compression like quantization and knowledge distillation mitigates some deployment constraints by lowering memory and inference costs, but does not remove the need for formal privacy mechanisms. Future work will therefore target integration of formal privacy guarantees and real-world cross-site evaluations to consolidate these promising empirical gains into clinically deployable systems.

ACKNOWLEDGMENTS

The authors would like to thank APJ Abdul Kalam Technological University (KTU) and the Department of Computer Science and Engineering, Government Engineering College, Thrissur, for their facilities and support provided during the research.

AUTHOR CONTRIBUTIONS

Conceptualization & Methodology: Bisna N. D.; Software & Validation: Bisna N. D., Helna Thomas; Formal Analysis & Investigation: Bisna N. D., Ajay James; Original Draft: Bisna N. D., Helna Thomas; Review & Editing: Bisna N. D., Ajay James.

REFERENCES

[1] Prabhakaran D, Jeemon P, Roy A, ”Cardiovascular Diseases in India: Current Epidemiology and Future Directions” , Pubmed, 2016

[2] Rayan Alanazi, ” Identification and Prediction of Chronic Diseases Using Machine Learning Approach ” ,Journal of Healthcare Engineering, 2022.

[3] Aldhyani, Theyazn, Alshebami, Ali, Alzahrani, Yahea, ” Soft Clustering for Enhancing the Diagnosis of Chronic Diseases over Machine Learning

Algorithms" ,Journal of Healthcare Engineering, 2020.
[4] Shah, D, Patel, S.Bharti, S.K, "Heart Disease Prediction using Machine Learning Techniques", SN Computing, 2020.
[5] Yuan X, Chen S, Sun C, Yuwen L, "A novel early diagnostic framework for chronic diseases with class imbalance",Sci Rep, 2020.
[6] Subahi, A.F.; Khalaf, O.I.; Alotaibi, Y.; Natarajan, R.; Mahadev, N.; Ramesh, T, "Mod- ified Self-Adaptive Bayesian Algorithm for Smart Heart Disease Prediction in IoTSystem", Sustainability, 2022
[7] Y. A. Nanehkaran, Zhu Licai, Junde Chen, Qiu Zhongpan, Yuan Xiaofeng, Yahya Dorostkar Navaei, Sajad Einy, " Diagnosis of Chronic Diseases Based on Patients'Health Records in IoT Healthcare Using the Recommender System",Wireless Communications and Mobile Computing, 2022
[8] Farman Ali, Shaker El-Sappagh, S.M. Riazul Islam, Daehan Kwak, Amjad Ali, Muham- mad Imran, Kyung-Sup, " A smart healthcare monitoring system for heart diseaseprediction based on ensemble deep learning and feature fusion", Information Fusion, 2020.
[9] Noman Mohammed, Xiaoqian Jiang, Rui Chen, Benjamin C M Fung, Lucila Ohno-Machado, "Privacy-preserving heterogeneous health data sharing", Health Informa- tion Science and Systems, 2013.
[10] Jagadeeswari, V. Subramaniyaswamy, R. Logesh, Vijayakumar, "A study on medical Internet of Things and Big Data in personalized healthcare system", Health Informa- tion Science and Systems, 2012.
[11] M. Chamikaraa, Bertoka, Khalila, Liub, S. Camtepeb, "Privacy Preserving istributed Machine Learning with Federated Learning", Health Information Science and Systems, 2012.
[12] Khan, M. Alsulami, M. Yaqoob, Alsadie, D Saudagar,AlKhathami, ; Farooq Khatta," Asynchronous Federated Learning for Improved Cardiovascular Disease Prediction Using Artificial Intelligence", Diagnostics, 2023.
[13] A.S Aathika, "Health Care Data Privacy Preserving using Federated Learning: A case of Heart Disease Prediction", Computer Systems Engineering, 2021.
[14] Bui Tan Hai Dang,Phan Huu Luan, Vuong Dinh, Thanh Ngan, "TrustFedHealth:Federated Learning with Homomorphic Encryption and Blockchain for Heart Disease Prediction in the Smart Healthcare", International Conference on Advanced Technologies for Communications, 2023.
[15] Joo Hun Yoo, Hyejun Jeong, Jaehyeok Lee, Tai-Myoung Chung Open problems in medical federated learning", International Journal of Web Information Systems, 2022.
[16] Aleksander Wawer Roman Dusek,"Federated learning for clinical event classification using vital signs data" International ACM SIGIR Conference on Research and Development, 2023.
[17] Zhe Zhang, Zhifeng Wu,"Iot-based federated learning model for hypertensive retinopathy lesions classification", Ieee Transactions On Computational Social Sys- tems, 2023.
[18] Zhi-qiang Wei Yong-Quan Yang Cai-zhi Liu, Yan-xiu Sheng," Federated learning-empowered disease diagnosis mechanism in the internet of medical things: From the privacy-preservation perspective", IEEE Transactions On Industrial Informatics, 2023.
[19] Theofanie Mela Theodora S. Brisimi, Ruidi Chen, "Federated learning of predictive models from federated electronic health records", Int J Med Inform, 2018.
[20] Moody GB, Mark RG. The impact of the MIT-BIH Arrhythmia Database. IEEE Eng in Med and Biol 20(3):45-50 (May-June 2001). (PMID:11446209)
[21] Ansari Y, Mourad O, Qaraqe K, Serpedin E. Deep learning for ECG Arrhythmia detection and classification: an overview of progress for period 2017-2023. Front Physiol. 2023 Sep 15;14:1246746. doi: 10.3389/fphys.2023.1246746. PMID: 37791347; PMCID: PMC10542398.
[22] Li R, Liang Y, Jin L, Wang S (2020) Multimodal deep learning for cardiovascular disease risk prediction based on EHR Data. IEEE Access 8:106537–106547
[23] Yaqoob, M.M.; Nazir, M.; Khan, M.A.; Qureshi, S.; Al-Rasheed, A. Hybrid Classifier-Based Federated Learning in Health Service Providers for Cardiovascular Disease Prediction. Appl. Sci. 2023, 13, 1911.

[24] Singh S, Singh A, Limkar S (2024) Prediction of heart disease using deep learning and internet of medical things. International Journal of Intelligent Systems and Applications in Engineering 12(1s):512–525

[25] El-Sofany, H., Bouallegue, B. amp; El-Latif, Y.M.A. A proposed technique for predicting heart disease using machine learning algorithms and an explainable AI method. Sci Rep 14, 23277 (2024). https://doi.org/10.1038/s41598-024-74656-2

[26] Chintan, M. B., Parth, P., Tarang, G. amp; Pier, L. M. Effective Heart Disease Prediction Using Mach. Learn. Techniques Algorithms, 16, 88, https://doi.org/10.3390/a16020088, (2023).

[27] Houssein, E.H.; Ibrahim, I.E.; Hassaballah, M.; Wazery, Y.M. Integration of internet of things and cloud computing for cardiac health recognition. In Metaheuristics in Machine Learning: Theory and Applications; Springer: Berlin/Heidelberg, Germany, 2021; pp. 645–661.

[28] Timmis, A.; Vardas, P.; Townsend, N.; Torbica, A.; Katus, H.; de Smedt, D.; Gale, C.P.; Maggioni, A.P.; Huculeci, R.; Kazakiewicz, D.; et al. Cardiovascular disease statistics 2021. Eur. Heart J. 2022, 43, 716–799

[29] Tozlu, B.H.; Sim sek, C.; Aydemir, O.; Karavelioglu, Y. A High performance electronic nose system for the recognition of myocardial infarction and coronary artery diseases. Biomed. Signal Process. Control 2021, 64, 102247.

[30] G. Muhammad et al., quot;Enhancing Prognosis Accuracy for Ischemic Cardiovascular Disease Using K Nearest Neighbor Algorithm: A Robust Approachquot;, IEEE Access, vol. 11, pp. 97879-97895, 2023

[31] Hill, N.R.; Ayoubkhani, D.; McEwan, P.; Sugrue, D.M.; Farooqui, U.; Lister, S.; Lumley, M.; Bakhai, A.; Cohen, A.T.; O'Neill, M.; et al. Predicting atrial fibrillation in primary care using machine learning. PLoS ONE 2019, 14, e0224582.

[32] Yazdinejad, A.; Dehghantanha, A.; Karimipour, H.; Srivastava, G.; Parizi, R.M. A robust privacy-preserving federated learning model against model poisoning attacks. IEEE Trans. Inf. Forensics Secur. 2024, 19, 6693–6708.

[33] Shao, M.; Zhou, Z.; Bin, G.; Bai, Y.; Wu, S. A wearable electrocardiogram telemonitor- ing system for atrial fibrillation detection. Sensors 2020, 20, 606.

[34] Y. Otoum and A. Nayak, "Differential Privacy-Driven Framework for Enhancing Heart Disease Prediction," arXiv preprint arXiv:2504.18007, Apr. 2025

[35] B. Veera Jyothi, E. Ramalakshmi, L. Suresh Kumar, and B. Surya Samantha, "Heart Disease Prediction using Federated Learning," in 2024 International Conference on Intelligent Computing and Emerging Communication Technologies (ICEC), IEEE, Nov. 2024.

[36] A. Faes, A. Pirmani, Y. Moreau, and L. Peeters, "Federated Block-Term Tensor Regression for Decentralised Data Analysis in Healthcare," arXiv preprint arXiv:2412.06815, Dec. 2024.

[37] Y. Otoum, C. Hu, E. H. Said, and A. Nayak, "Enhancing Heart Disease Prediction with Federated Learning and Blockchain Integration," Future Internet, vol.16, no.10, art.372, 2024

[38] G. Tsoumplekas, I. Siniosoglou, V. Argyriou, I. D. Moscholios, and P. Sarigiannidis, "Enhancing Performance for Highly Imbalanced Med ical Data via Data Regularization in a Federated Learning Setting," arXiv:2405.20430, May 2024

[39] P. Wagner, N. Strodthoff, R. Bousseljot, D. Kreiseler, F. Lunze, T. Samek, and W. Schaeffter, "PTB-XL, a large publicly available electrocardiography dataset," *Scientific Data*, vol. 7, no. 1, pp. 1–15, May 2020, doi: 10.1038/s41597-020-0495-6.