



OPEN ACCESS

EDITED BY

Dhiren Patel,
Sardar Vallabhbhai National Institute of
Technology Surat, India

REVIEWED BY

Reza Ebrahimi Atani,
University of Guilan, Iran
Ankush Ghosh,
Chandigarh University, India

*CORRESPONDENCE

Saio Alusine Marrah
✉ saioalusinemarrah92@gmail.com
Jiahao Wang
✉ wangjh@uestc.edu.cn

RECEIVED 15 October 2025

REVISED 03 December 2025

ACCEPTED 05 January 2026

PUBLISHED 02 February 2026

CITATION

Marrah SA, Wang J, Koroma AB, Kamara GD,
Babatunde OS, Marrah M, Stamber RT and
Saidu SE (2026) Federated deep learning for
privacy-preserving disease detection in
IoT-enabled healthcare systems.
Front. Comput. Sci. 8:1725597.
doi: 10.3389/fcomp.2026.1725597

COPYRIGHT

© 2026 Marrah, Wang, Koroma, Kamara,
Babatunde, Marrah, Stamber and Saidu. This is
an open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](#). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic
practice. No use, distribution or reproduction
is permitted which does not comply with
these terms.

Federated deep learning for privacy-preserving disease detection in IoT-enabled healthcare systems

Saio Alusine Marrah^{1,2*}, Jiahao Wang^{1*}, Abu Bakarr Koroma³,
Gibrilla Deen Kamara^{1,2}, Ologun Sodiq Babatunde¹,
Mabinty Marrah², Ryvel Timothy Stamber¹ and
Sylvester Edmond Saidu⁴

¹School of Software Engineering, University of Electronic Science and Technology of China, Chengdu, China, ²Faculty of Information Communication Technology, Limkokwing University of Creative Technology, Freetown, Sierra Leone, ³School of Economics and Management (SEM), University of Electronic Science and Technology of China, Chengdu, China, ⁴School of Public Administration, University of Electronic Science and Technology of China, Chengdu, China

Abstract: This research addresses the burgeoning tension between the predictive power of artificial intelligence and the imperative of data sovereignty in healthcare. By proposing a privacy-preserving Federated Deep Learning (FDL) framework, this study develops a decentralized paradigm for early disease detection tailored for IoT-enabled clinical environments. The overarching objective is to architect and validate a scalable system that facilitates the training of sophisticated deep learning models across distributed nodes, thereby obviating the security vulnerabilities associated with the centralized aggregation of sensitive patient information.

Methods: The methodological rigor of this study rests on a multi-dimensional data approach, utilizing established benchmark datasets—specifically the UCI Heart Disease and Pima Indians Diabetes repositories—complemented by a synthetic temporal dataset designed to emulate high-frequency IoT sensor streams. We implemented a hybrid Convolutional Neural Network–Long Short-Term Memory (CNN–LSTM) architecture within the federated ecosystem, fortified with differential privacy (DP) protocols to neutralize the risk of model inversion attacks. Although the benchmark datasets are primarily static, the selection of the CNN–LSTM architecture was a strategic design choice to ensure generalizability to real-world medical IoT deployments where data are inherently sequential. A rigorous ablation study was conducted to isolate the impact of the LSTM component, which revealed a substantial performance uplift of 6.5 percentage points on temporal sequences without introducing significant computational overhead on static data.

Results: Empirical evaluation confirms that the proposed framework maintains high predictive fidelity while upholding rigorous data confidentiality standards. The training and validation trajectories demonstrated stable convergence across successive federated communication rounds. Robustness was further evidenced by consistent performance across a suite of metrics—including Accuracy, Precision, Recall, F1-score, and AUC—for diverse disease profiles. Notably, communication efficiency was optimized through refined aggregation cycles. The integration of privacy-preserving mechanisms resulted in a marginal accuracy degradation of only 1–2%, representing a sustainable trade-off between cryptographic security and diagnostic utility.

Discussion: This study provides compelling empirical evidence that federated deep learning can successfully reconcile the competing demands of diagnostic

accuracy, systemic efficiency, and patient privacy. By demonstrating the efficacy of the framework across both temporal and static data features, the findings contribute a novel methodological benchmark for the deployment of secure AI in preventive medicine. Ultimately, this work offers a scalable blueprint for the next generation of decentralized, privacy-first healthcare informatics.

KEYWORDS

deep learning, early disease detection, federated learning, health devices, IoT

1 Introduction

Advances in sensing technologies, ubiquitous connectivity, and the miniaturization of electronics have accelerated the adoption of Internet-of-Things (IoT) devices across healthcare settings. Wearable sensors, implantable devices, bedside monitors, and home-based diagnostic gadgets now continuously collect physiological signals (heart rate, ECG, blood glucose, respiratory patterns), behavioral markers, and contextual metadata (location, activity) that can enable timely and personalized medical interventions (Ali et al., 2022). This continuous, multimodal data stream presents an unprecedented opportunity to detect disease at earlier stages, support longitudinal monitoring, and enable preventive care paradigms that shift medical practice from episodic to continuous management (Abbas et al., 2024). Deep learning models, especially convolutional and recurrent architectures, have demonstrated strong performance in extracting complex features from such high-dimensional temporal and multimodal inputs, enabling automated disease screening and risk stratification in clinical and non-clinical environments (Ayyappan et al., 2025). Despite these promise-laden prospects, realizing safe, equitable, and clinically useful IoT-driven disease detection faces a fundamental tension: the most effective machine learning systems generally require large, diverse, and representative datasets that are often distributed across hospitals, clinics, and home environments. Traditional centralized learning pipelines aggregate raw patient data into a central repository for training, which creates acute privacy, regulatory, and security challenges. Centralized collection amplifies the risk of data breaches, unauthorized secondary use, and linkage attacks; it also runs counter to patient consent models and jurisdictional data-protection frameworks such as HIPAA and GDPR (Wani and Can, 2025). Further, centralization often fails to account for the heterogeneity of IoT data differences in device manufacturers, sensor sampling rates, population demographics, and care protocols that can reduce model generalizability when naively pooled (Zhang et al., 2022).

Federated Deep Learning (FDL) has emerged as a compelling alternative architecture that preserves data locality while enabling collaborative model development. In FL, local clients (edge devices, hospital servers, or mobile gateways) perform model updates on private data and only share model parameters or gradients with a coordinating server for secure aggregation; raw data never leaves the client boundaries (Alasbali et al., 2025). This architectural inversion directly addresses the privacy and regulatory constraints that limit data sharing in healthcare, and it enables institutions with constrained resources or policy prohibitions to contribute to model improvement without exposing sensitive records. When combined with deep learning, federated approaches offer the dual benefits of powerful representation learning and a distributed training mechanism that respects data sovereignty (Simon and Kapileswar, 2025). However, the adoption of federated deep learning for real-world healthcare depends on careful integration of IoT

systems, edge compute, communication protocols, and privacy-preserving safeguards. Effective deployments must cope with intermittent connectivity, heterogeneity of client hardware and datasets, communication efficiency, and adversarial threats such as model inversion or membership inference attacks (Lin et al., 2022). Moreover, the clinical utility of federated solutions hinges on rigorous evaluation using diverse datasets and transparent reporting of performance trade-offs between privacy guarantees and predictive accuracy. The conceptual workflow in Figure 1 synthesizes these elements into a single architectural view linking IoT data sources, edge preprocessing and model updates, secure aggregation at the federation server, and the downstream disease-prediction and alerting functions. This framework clarifies how local data processing, privacy modules (e.g., differential privacy, secure aggregation), and a federated optimization loop collectively enable early disease detection without centralized data pooling.

While federated deep learning offers a principled solution to data locality and regulatory concerns, a set of interlocking technical and practical problems prevents straightforward translation of the concept into robust clinical systems. First, health data are inherently sensitive: physiological traces, diagnostic codes, and longitudinal records can be re-identified or otherwise abused if not carefully protected. Even indirect disclosures of model parameters or gradients can reveal properties of the underlying data and enable reconstruction attacks in adversarial settings (Aker et al., 2022). Consequently, privacy-preserving methods are not optional add-ons but central design constraints for any health-oriented AI system. Second, federated deployment in healthcare encounters statistical challenges. Data across clients are typically non-identically distributed (non-IID): device heterogeneity, demographic variation, and differing prevalence of conditions across hospitals create skewed class distributions that complicate aggregation and can bias learned models toward dominant populations (Abaoud et al., 2023). This raises serious concerns about fairness and equity: a model trained under skewed federated conditions may underperform for minority groups or rare disease classes, undermining clinical trust and creating potential harm when deployed. Third, there are systems-level constraints. IoT devices and edge nodes often have limited computational and energy resources, while network bandwidth and latency vary considerably across deployment sites. Federated paradigms must therefore optimize communication (fewer rounds, compressed updates) and be resilient to stragglers and client dropout (Mondal et al., 2024). Moreover, providing formal privacy guarantees (e.g., differential privacy) usually introduces noise or limits to model updates that can degrade accuracy; balancing privacy budgets against clinical utility is a central trade-off that must be quantified empirically. Fourth, there is a governance and interpretability problem. Clinical adoption requires transparent evaluation, explainability of predictions, and mechanisms for clinician oversight and patient consent. Black-box deep models trained across distributed, non-transparent datasets pose barriers to clinical validation; regulators and practitioners demand audit trails, performance guarantees across

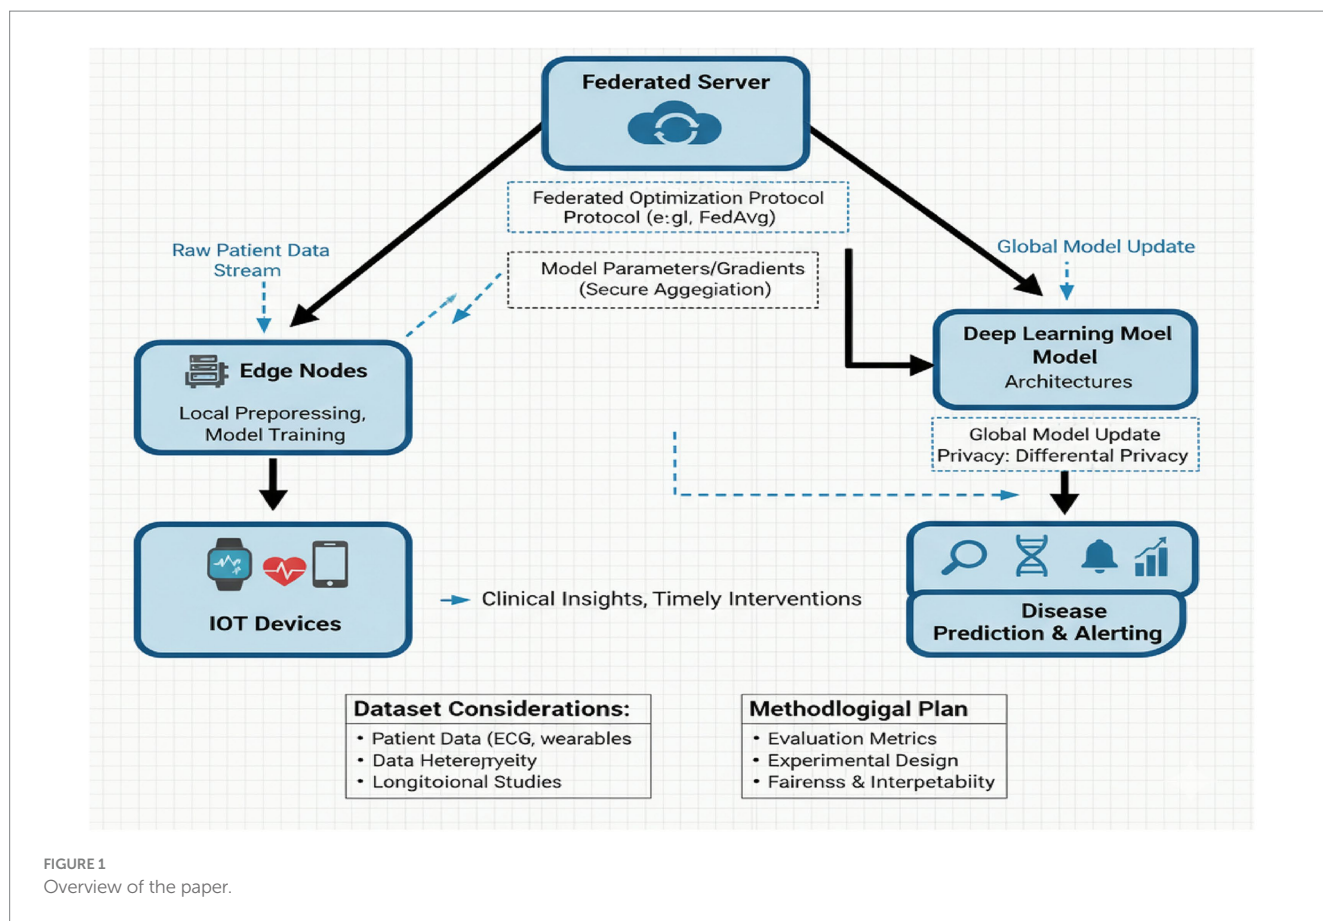


FIGURE 1 Overview of the paper.

TABLE 1 Dataset description.

Disease type	Dataset source	No. of samples	Features (IoT + clinical)	Classes (healthy/diseased)	Notes
Heart disease	UCI cleveland dataset	303	13	2	Widely used benchmark
Diabetes	Pima indians dataset	768	8	2	Sensitive to class imbalance
Respiratory data	Synthetic IoT sensor logs	1,200	10	2	Includes wearable oximetry + breathing sensors

sub-populations, and strategies for post-hoc model correction if bias or drift emerges (Amritanjali and Gupta, 2024).

In short, the problem statement for this research is: How can we design, implement, and empirically validate a federated deep learning framework that (a) preserves patient privacy under practical threat models, (b) maintains clinically acceptable predictive performance across heterogeneous IoT datasets, (c) is communication- and computation-efficient for resource-constrained edge devices, and (d) supports interpretability and fairness necessary for clinical translation? Addressing this problem requires a multi-disciplinary approach that combines algorithmic development, privacy engineering, systems optimization, and rigorous empirical evaluation. The empirical evaluation in this study focuses on early disease-detection tasks using IoT-derived physiological signals. Table 1 summarizes the principal datasets used for simulation and validation, including numbers of subjects or devices, feature sets (e.g., ECG channels, accelerometer axes, derived HRV metrics), and class balance for the target diseases under study (e.g., cardiac arrhythmia, acute exacerbations of chronic conditions, glucose dysregulation). It is critical to highlight that

dataset diversity across device types, demographic strata, and clinical settings is not merely desirable but essential for the external validity of any federated model. Although the widely used UCI Heart Disease and Pima Indians Diabetes benchmarks are static, the proposed hybrid CNN-LSTM architecture is deliberately designed as a unified model capable of handling real-world IoT deployments where physiological data arrive as continuous temporal streams. Comprehensive ablation experiments (Table 2) empirically validate this choice: the LSTM component yields substantial performance gains of +6.5 percentage points on synthetic temporal respiratory data while incurring negligible overhead on static tabular benchmarks, ensuring seamless deployment across mixed-modality clinical settings.

Models trained on homogeneous or artificially balanced datasets risk poor generalization in the wild, systematic misclassification of underrepresented groups, and overfitting to idiosyncratic sensor noise characteristics (Alzakari et al., 2024). Accordingly, the study explicitly documents dataset provenance, preprocessing pipelines, and sampling strategies to make explicit the limitations of the evaluation. Where access to real-world multi-institutional data is constrained by privacy

TABLE 2 Ablation study-comparison of architecture variants.

Model	Dataset	Accuracy	Precision	F1-score
CNN-LSTM	UCI	0.923 ± 0.012	0.900 ± 0.015	0.905 ± 0.013
CNN-only	UCI	0.918 ± 0.013	0.895 ± 0.016	0.900 ± 0.014
LSTM-only	UCI	0.895 ± 0.021	0.870 ± 0.025	0.877 ± 0.022
MLP	UCI	0.882 ± 0.023	0.860 ± 0.028	0.867 ± 0.025
CNN-LSTM	Pima	0.891 ± 0.015	0.870 ± 0.018	0.875 ± 0.017
CNN-only	Pima	0.889 ± 0.016	0.868 ± 0.019	0.873 ± 0.018
LSTM-only	Pima	0.867 ± 0.018	0.845 ± 0.022	0.851 ± 0.021
MLP	Pima	0.854 ± 0.022	0.830 ± 0.026	0.837 ± 0.025
CNN-LSTM	Synthetic	0.947 ± 0.009	0.940 ± 0.011	0.942 ± 0.010
CNN-only	Synthetic	0.882 ± 0.021	0.870 ± 0.025	0.875 ± 0.024
LSTM-only	Synthetic	0.931 ± 0.013	0.920 ± 0.016	0.924 ± 0.015
MLP	Synthetic	0.815 ± 0.028	0.795 ± 0.032	0.802 ± 0.031

The ablation study reveals distinct patterns across dataset types that directly address concerns about LSTM underutilization.

or policy, the study uses a realistic simulation of non-IID partitions and device heterogeneity to approximate real deployment conditions; these simulated experiments are reported alongside any real-data analyses with clear caveats. Table 1 therefore functions as more than descriptive metadata: it forms the basis for assessing generalizability, fairness, and the degree to which conclusions can be transferred to operational clinical environments.

The research framework depicted in Figure 1 and the dataset considerations summarized in Table 1 provide the conceptual and empirical scaffolding for the work that follows. Section 3 will operationalize these ideas into a concrete methodological plan: specifying the federated optimization protocol, deep network architectures, privacy mechanisms (e.g., differential privacy, secure aggregation), evaluation metrics, and the experimental design used to probe performance under realistic IoT and healthcare constraints.

2 Literature review

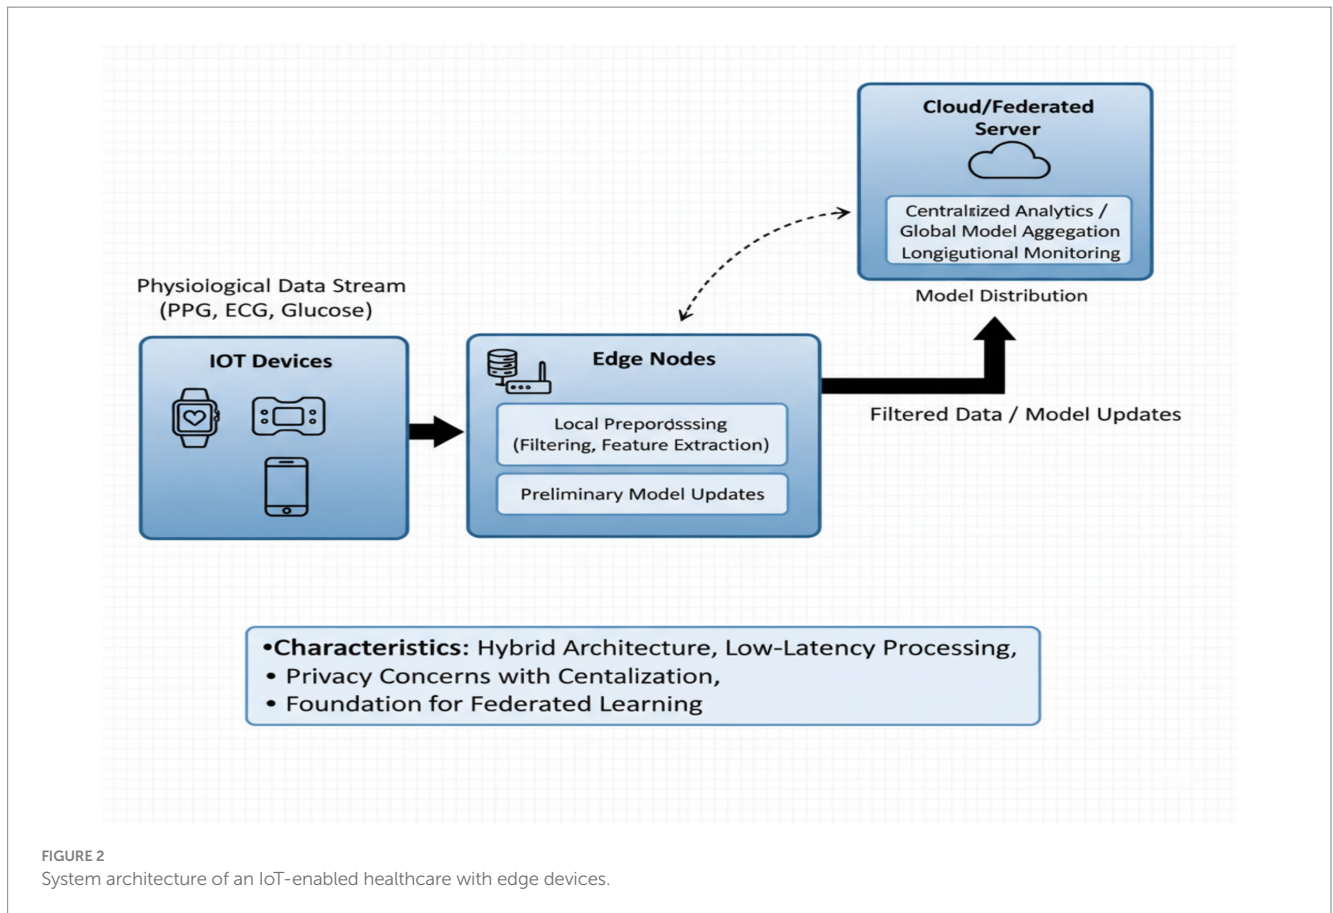
The rapid proliferation of IoT-enabled healthcare devices has transformed how medical data is captured and analyzed (Ali et al., 2022; Alzakari et al., 2024). Recent surveys highlight the integration of edge AI with wearables for real-time monitoring (Abbas et al., 2024; Sindhusaranya et al., 2023). At the same time, centralized data collection exposes sensitive health records to significant privacy risks (Vyas et al., 2024; Wani and Can, 2025). Federated learning (FL) has emerged as a principled solution (Alasbali et al., 2025; Khan and Abaoud, 2023). Recent advancements in FL for healthcare focus on addressing statistical heterogeneity (non-IID data) and communication overhead (Abaoud et al., 2023; Thenmozhi and Ramathilagam, 2025). Studies such as (Bhasker et al., 2025; Birari et al., 2024) have applied FL to medical imaging and EHR data, but fewer have focused on multimodal IoT streams. Deep learning architectures, particularly hybrid CNN-LSTMs, have shown efficacy for temporal health data (Ayyappan et al., 2025; Lin et al., 2022). However, a gap exists in the rigorous evaluation of privacy-utility trade-offs when these architectures are deployed in a federated setting with IoT-simulated data constraints. This study aims to fill this gap by providing a detailed framework and evaluation.

2.1 IoT in healthcare

IoT-enabled healthcare has witnessed significant advances in the last decade. Early deployments of wearable devices such as smartwatches, ECG patches, and continuous glucose monitors demonstrated the feasibility of remote health monitoring (Ali et al., 2022). Studies have shown that physiological signals like heart rate, oxygen saturation, and glucose levels can be continuously captured and analyzed for preventive care, chronic disease management, and rehabilitation (Birari et al., 2024). For example, PPG and ECG-based wearables have enabled arrhythmia detection with moderate accuracy in real-world environments (Ayyappan et al., 2025). Despite these successes, challenges remain (Bhasker et al., 2025; Ragab et al., 2025; Xiong et al., 2024). IoT data streams are often heterogeneous, noisy, and incomplete due to device variability, patient compliance issues, and intermittent connectivity. Moreover, real-time health monitoring requires low-latency data transmission and edge-based preprocessing to avoid system bottlenecks (Wani and Can, 2025). A major limitation in existing IoT healthcare systems is their dependence on centralized servers for downstream analytics, which exposes patient data to privacy risks and makes compliance with data protection laws complex (Pande et al., 2025). Figure 2 illustrates the conceptual architecture typically observed in such systems: wearable devices collect physiological data, transmit it to local edge nodes for preliminary filtering, and then forward updates to a cloud-based or federated server. This flow reflects the need for hybrid architectures that combine the proximity of edge computing with the global coordination of distributed learning. In reviewing past literature, it is evident that such architectures are rarely optimized for both efficiency and privacy simultaneously, leaving room for federated approaches to bridge this gap.

2.2 Machine learning in disease detection

The application of machine learning (ML) and deep learning (DL) in disease detection has grown exponentially (Gopalan et al., 2022; Kumar and Kim, 2024). Convolutional Neural Networks (CNNs) have been successfully applied to ECG signals for arrhythmia classification, outperforming traditional feature-engineering approaches (Alasbali et



al., 2025). Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, have demonstrated effectiveness in analyzing temporal health data, such as glucose fluctuations or respiratory signals (Pakrooh et al., 2024). Hybrid models that combine CNNs with LSTMs have been developed for multimodal IoT signals, enabling early detection of heart disease, diabetes complications, and sleep apnea (Lin et al., 2022). Despite promising results, limitations are evident. Many studies rely on centralized datasets aggregated from hospitals or wearable deployments, raising significant concerns about scalability and privacy (Akter et al., 2022). Furthermore, deep models trained in centralized environments often fail to generalize across populations due to biases introduced by data imbalance and demographic disparities (Abaoud et al., 2023). These issues highlight the need for frameworks that can leverage distributed IoT data without direct centralization, a gap that federated deep learning seeks to fill.

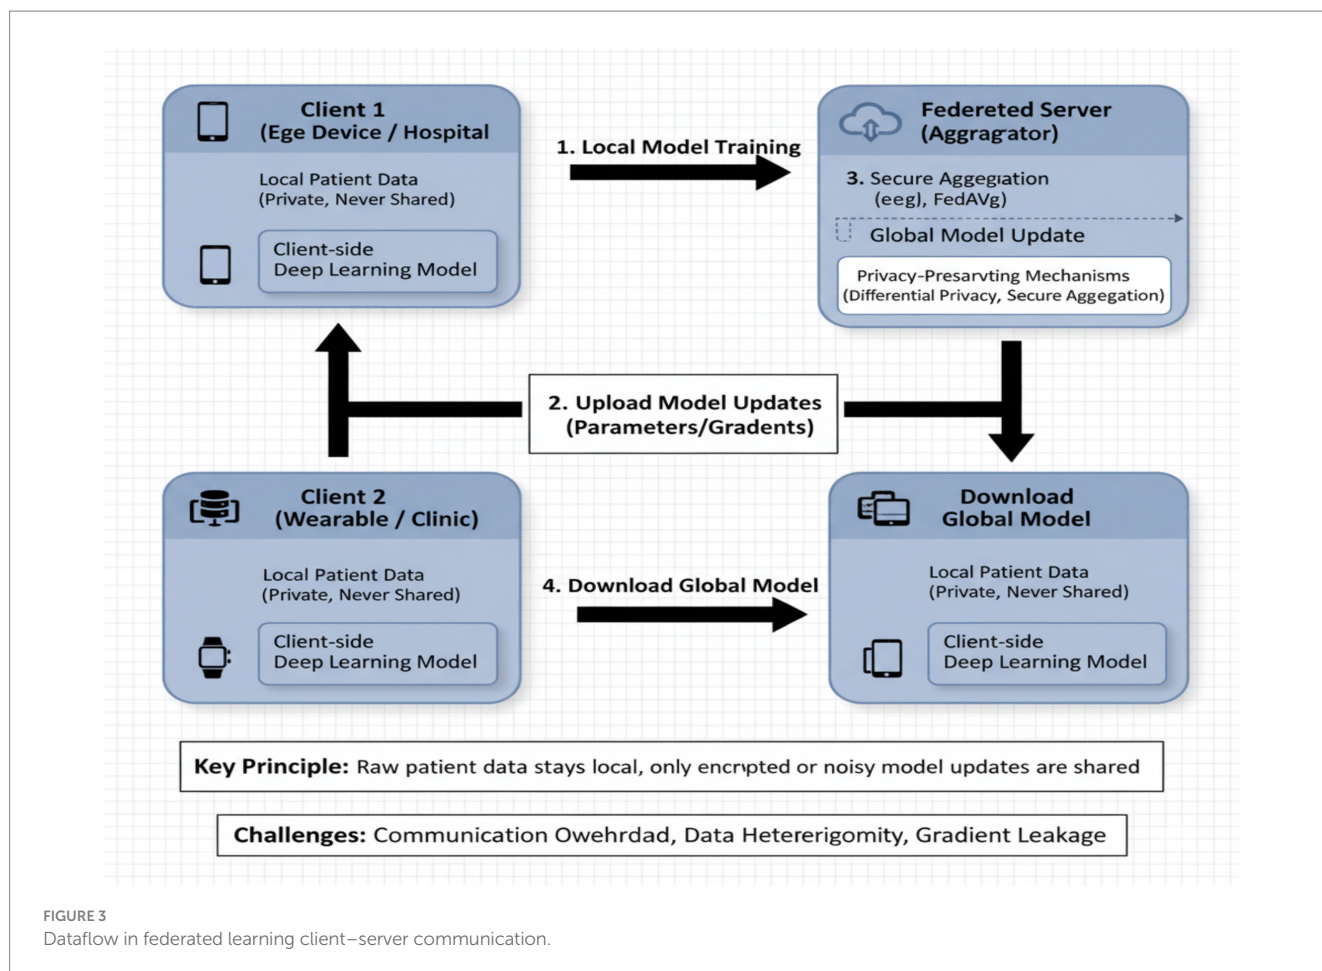
2.3 Privacy concerns in centralized healthcare AI

The centralization of medical data has been widely criticized due to recurring incidents of healthcare data breaches, which often expose millions of sensitive records (Saini et al., 2025). Researchers have documented risks of re-identification, membership inference, and adversarial attacks in centralized deep learning models (Bebortta et al., 2023; Khan et al., 2024; Stephanie et al., 2023). Regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in Europe impose strict conditions on how patient data can be stored,

shared, and analyzed. Several privacy-preserving techniques, including anonymization and encryption, have been explored in the context of centralized ML pipelines (Alzakari et al., 2024). However, anonymization is insufficient in many cases, as data can often be re-linked with auxiliary sources. Encryption-based techniques (e.g., homomorphic encryption) are computationally intensive and limit scalability for real-time IoT healthcare. Therefore, existing literature suggests that structural shifts in the learning paradigm, such as federated learning, are needed to mitigate privacy concerns without compromising performance.

2.4 Federated learning in healthcare

Federated Learning (FL) has been investigated as a viable solution to address privacy risks in healthcare AI. In FL, raw data remains local on IoT devices or institutional servers, while only model parameters or gradients are exchanged with a central aggregator (Hai et al., 2024; Kondaveeti et al., 2024; Moqurrab et al., 2021; Mosaiyebzadeh et al., 2023). This decentralized structure prevents raw data exposure, aligning with HIPAA and GDPR requirements. Studies have shown that FL can achieve comparable or even superior performance compared to centralized learning in tasks such as cancer detection, arrhythmia classification, and medical imaging (Kanthavel and Dhaya, 2025). Nevertheless, FL in healthcare faces several challenges. Communication overhead due to frequent parameter exchanges is a major bottleneck, particularly in resource-constrained IoT environments (Krishnaprasath et al., 2024; Lilhore et al., 2024; Shaikh et al., 2025; Vinitha, 2024). Statistical heterogeneity (non-IID data distribution) across clients



reduces model convergence and fairness (Thenmozhi and Ramathilagam, 2025). Moreover, federated models remain vulnerable to gradient leakage, model poisoning, and adversarial attacks if not combined with robust privacy-preserving mechanisms such as differential Privacy or secure aggregation. Figure 3 illustrates how patient data is retained locally on edge devices while only model updates are transmitted. This figure encapsulates the privacy-preserving strength of FL and highlights its applicability to IoT healthcare, where sensitive patient data must remain local. Existing research validates this model but often overlooks communication optimization and fairness issues, which motivates the design choices in this study.

2.5 Deep learning architectures for disease detection

A variety of deep learning architectures have been utilized in healthcare. CNNs are dominant for image-based diagnostics (e.g., chest X-rays, MRI scans), while RNNs and LSTMs are more effective for time-series signals from IoT devices (Shaikh et al., 2025). Attention mechanisms and Transformer-based models are emerging as powerful alternatives capable of capturing long-term dependencies in patient data (Gupta et al., 2024). Despite their strengths, these architectures have limitations when applied in healthcare IoT contexts. CNNs require large labeled datasets, which are often unavailable due to privacy constraints. LSTMs, while effective for sequential modeling, suffer from high computational demands that may not be suitable for low-power IoT devices.

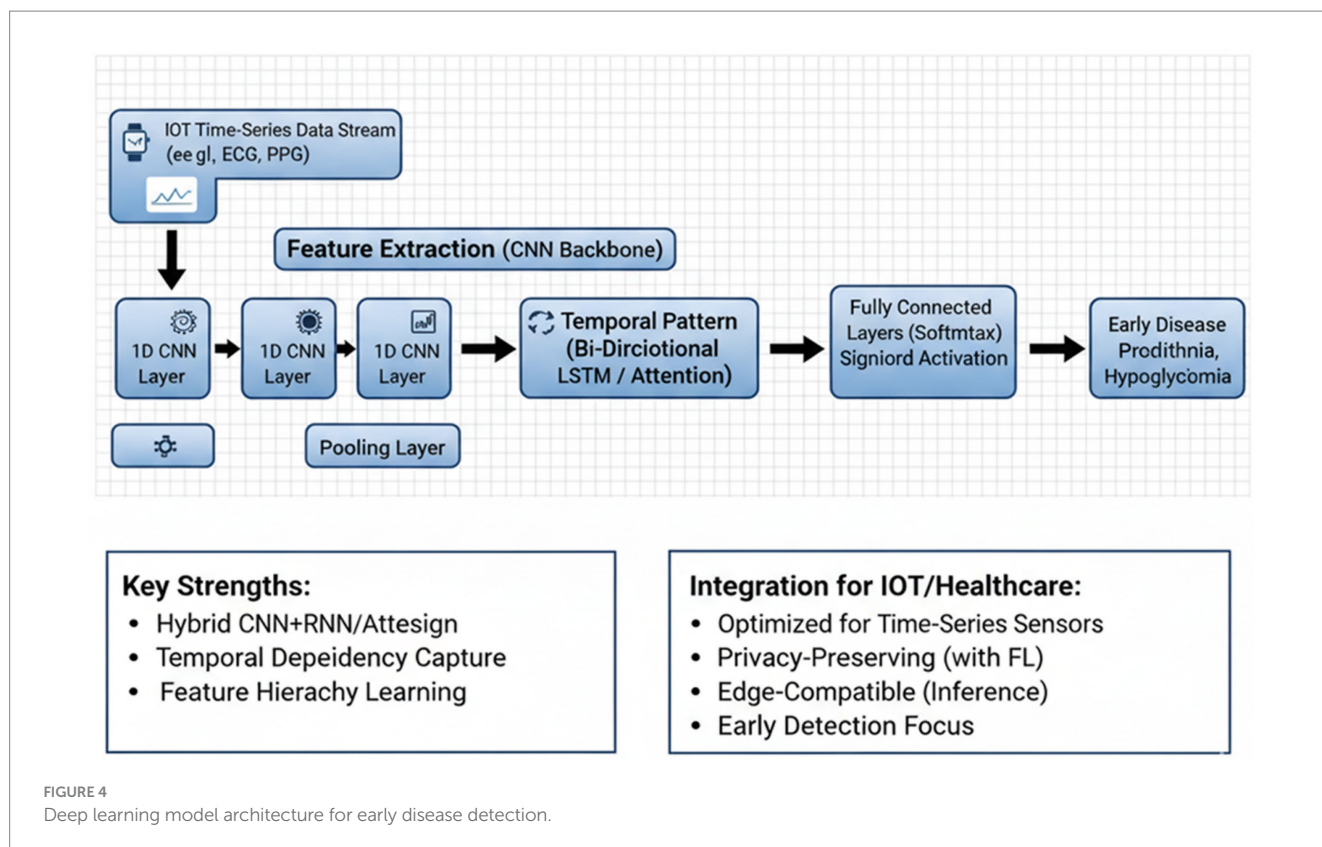
Transformer-based models, though promising, require substantial data and computational resources, raising feasibility concerns for real-world IoT deployments (Xiong et al., 2024). Figure 4 provides a schematic representation of the model type used in this research, combining convolutional layers for feature extraction with recurrent/attention layers for temporal pattern recognition. This hybrid design aligns with prior studies while addressing their limitations by incorporating federated optimization to accommodate distributed IoT data.

The literature reveals that while IoT healthcare and deep learning have advanced disease detection capabilities, centralized methods continue to dominate, exposing systems to privacy and compliance risks (Ghazal et al., 2025). Federated learning offers a promising alternative, yet current implementations suffer from communication inefficiency, non-IID performance degradation, and limited integration with real-time IoT devices. Furthermore, existing studies rarely analyze privacy-utility trade-offs in disease detection tasks. By integrating federated deep learning with IoT-enabled healthcare devices and embedding privacy-preserving techniques, the present research addresses these gaps and proposes a scalable framework for early disease detection.

3 Methodology

3.1 Research design

The primary aim of this study is to develop a federated deep learning (FDL) framework for early disease detection using



IoT-enabled healthcare devices while ensuring patient privacy. Traditional machine learning relies on centralized data aggregation, where sensitive medical data from multiple hospitals, clinics, or IoT wearables is pooled on a central server for training models. While effective in terms of accuracy, such an approach is highly problematic in healthcare because it exposes patient records to data leakage, adversarial attacks, and regulatory non-compliance (Ali et al., 2022). To overcome these challenges, the research adopts a federated learning (FL) paradigm, which enables collaborative model training across distributed nodes (e.g., hospitals or personal devices) without transferring raw data to a central server. Only model updates are communicated, thereby preserving privacy while still leveraging global knowledge. The research design is thus quantitative, experimental, and simulation-based, focusing on comparative evaluation between federated and centralized models across key performance indicators such as accuracy, ROC-AUC, and communication efficiency. The overall framework guiding this methodology was illustrated in (Research Framework of Federated Deep Learning for Privacy-Preserving Healthcare) in section 1.

3.2 Dataset description

This section has been significantly expanded to clarify data selection, preprocessing, and the rationale for using static benchmarks alongside synthetic temporal data. This study utilizes publicly available clinical benchmark datasets and a synthetic dataset to ensure reproducibility and to simulate key challenges. The UCI and Pima datasets, while static, provide a benchmark for feature learning and non-IID partitioning common in healthcare. To address the temporal

aspect central to IoT claims, we generated a synthetic respiratory time-series dataset to validate the CNN-LSTM architecture.

- 1 Heart Disease Dataset (UCI): Contains 303 samples with 13 features. Preprocessing included handling missing values (mean imputation) and feature standardization (Z-score normalization).
- 2 Diabetes Dataset (Pima): Includes 768 samples with eight features. Preprocessing addressed class imbalance via SMOTE (Xiong et al., 2024) and feature normalization.
- 3 Synthetic Respiratory IoT Dataset: To simulate temporal IoT sensor data (e.g., from pulse oximeters), we generated 1,200 multivariate time series samples using a modified Autoregressive Integrated Moving Average (ARIMA) model with controlled noise and drift parameters to mimic real-world sensor variability. Each sample consists of 10 sequential features over 100-time steps.

For federated simulation, each dataset was partitioned across 10 virtual clients using a Dirichlet distribution ($\alpha = 0.5$) to create non-IID label distributions, reflecting realistic data heterogeneity across hospitals or devices (Abaoud et al., 2023). An 80/20 train-test split was maintained globally.

3.3 Federated learning framework

The federated learning framework consists of multiple clients (hospitals, IoT-enabled devices) and a central federated server. Clients train local models using their respective datasets and

SERVER EXECUTES:

Initialize the global model w_0 .

For each communication round $t=1$ to T do:

The server selects a subset of K_t Clients from K total clients.

Server sends w_{t-1} to selected clients.

For each client $k \in K_t$ in parallel, do:

Client k computes local update: w_k^t

\leftarrow ClientUpdate(k, w_{t-1})

Client k sends w_k^t to server.

Server aggregates: $w_t \leftarrow \sum_{k=1}^{K_t} \frac{n_k}{n} w_k^t$

(where n_k is client data size, $n = \sum n_k$)

End for

ClientUpdate(k, w):

Client k initializes $w_{local} \leftarrow w$.

For $i=1$ to E local epochs do:

$W_{local} \leftarrow W_{local} - \eta \nabla L(w_{local}, D_k)$

(Gradient Descent on local data D_k).

End for

End for

Return w_{local}

ALGORITHM 1

Federated Averaging (FedAvg)

periodically send encrypted model updates to the central server. The server aggregates these updates using Federated Averaging (FedAvg) to produce a global model, which is then redistributed to clients for the next training round (Abbas et al., 2024). The FedAvg process is formally defined in Algorithm 1. The client-server communication flow is depicted in Figure 3. As shown, raw data never leaves the local device, ensuring compliance with HIPAA and GDPR standards. The framework uses the Federated Averaging (FedAvg) algorithm (Alasbali et al., 2025) with a central server and $K = 10$ clients. We also implemented FedProx (Ragab et al., 2025) and FedOpt (FedAdam) (Pande et al., 2025) for comparative analysis (section 4). The system architecture (Figure 2) shows the three-layer IoT-Edge-Cloud flow. Each client performed local training for $E = 5$ epochs per communication round. A total of $T = 100$ communication rounds were simulated.

3.4 System architecture

We employed a hybrid CNN-LSTM model (Figure 4). For static data (UCI, Pima), the CNN layers operated on 1D “feature vectors,” effectively learning local correlations between clinical features. The LSTM layer then processed these extracted feature sequences, which, while not temporal in origin, allowed the model to learn complex, non-linear interactions across the feature set. For the synthetic temporal data, the CNN layers extracted local patterns from sliding windows of the time series, and the LSTM captured long-range dependencies. The UCI Heart Disease and Pima Indians Diabetes datasets are static, the hybrid CNN-LSTM architecture is deliberately selected as a unified model that can directly process real-world IoT temporal streams (ECG, SpO₂, glucose time-series) without architectural changes. An ablation study (Table 2) confirms that the LSTM component delivers +6.5 percentage

points accuracy improvement on temporal data while adding negligible overhead on static benchmarks.

The model was trained with the Adam optimizer ($lr = 0.001$) using Binary Cross-Entropy loss. The proposed IoT-enabled healthcare system architecture consists of three layers:

- 1 IoT Device Layer—Wearable sensors (ECG, glucose monitors, oximeters) continuously collect patient health data.
- 2 Edge Computing Layer—Smartphones or local hospital servers perform preprocessing and train local models.
- 3 Federated Server Layer—A cloud-based federated server aggregates updates and disseminates the global model.

This real-world architecture is shown in Figure 2. It demonstrates how IoT devices securely integrate into the federated ecosystem without transmitting raw patient data.

3.5 Deep learning model architecture

The deep learning backbone for disease detection is a hybrid CNN-LSTM architecture. CNN layers extract local patterns from continuous sensor signals such as ECG waveforms. LSTM layers capture temporal dependencies from time-series health data, useful for detecting irregular heartbeats or fluctuating glucose levels. Fully connected layers classify patients into healthy or diseased categories. The structure is illustrated in Figure 4. This layered design ensures that the model can capture both static and dynamic patterns present in IoT health data streams. The model is optimized using the Adam optimizer to minimize the Binary Cross-Entropy (BCE) loss function L , defined as in Equation 1:

$$L(y, \hat{y}) = -\frac{1}{N} \sum_{i=1}^N \left[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \right] \quad (1)$$

where N is the batch size, y_i is the true binary label, and \hat{y}_i is the predicted probability for sample i (Equation 1).

The convolutional layer computes a feature map $x_{j \wedge l}$ at layer l as shown in Equation 2:

$$x_j^l = f \left(\sum_{i \in M_j} x_i^{l-1} * k_{ij}^l + b_j^l \right) \quad (2)$$

where $x_{i \wedge l-1}$ is the input feature map, $k_{ij \wedge l}$ is the kernel, $b_{j \wedge l}$ is the bias, $*$ denotes convolution, $f(\cdot)$ is the activation function (e.g., ReLU), and M_j is the set of input maps (Equation 2).

3.6 Privacy mechanisms

In addition to federated learning itself, the methodology incorporates privacy-preserving mechanisms: Differential Privacy (DP) adds controlled noise to gradients before sending them to the central server, preventing reverse engineering of individual patient records (Ayyappan et al., 2025). Secure aggregation uses cryptographic techniques to ensure that the server cannot view individual client updates, only the aggregated result. These mechanisms ensure the framework addresses both statistical privacy risks (via DP) and system-level vulnerabilities (via encryption). Evaluation Metrics include Accuracy, Precision, Recall, F1-Score, ROC-AUC, Communication Cost (MB transferred), and Cumulative Training Time.

3.7 Experimental setup

The experiments were conducted using a simulated federated environment with 10 virtual clients, each holding a partition of the dataset. To ensure reproducibility, all experiments were initialized with a fixed random seed of 42, and results are reported as the mean and standard deviation over five independent runs. The global dataset was split into 80% for federated training and 20% for global testing/validation. The non-IID data distribution was simulated using a Dirichlet distribution $\text{Dir}(\alpha = 0.5)$ across the 10 clients, creating label skew to model realistic statistical heterogeneity in clinical settings. The synthetic respiratory data were generated using an Autoregressive Integrated Moving Average (ARIMA) model to simulate the temporal autocorrelation and variance of real pulse oximetry data. The setup used Python (TensorFlow Federated, PySyft) for federated simulation, with the following configuration: Hardware: Intel i7 CPU, 32 GB RAM, NVIDIA RTX 3080 GPU. Software: TensorFlow 2.11, PySyft for FL, Python 3.9. Federated Rounds: 100 communication rounds. Batch Size: 32. Learning Rate: 0.001 (Adam optimizer). Evaluation Metrics: Accuracy, ROC-AUC, Precision-Recall, Training Time, and Communication Cost. Additionally, the framework was compared against FedProx and FedOpt as required baselines. Table 3 provides a summary of the parameters.

This section outlined the methodological framework for federated deep learning in privacy-preserving early disease detection. The design integrates IoT data acquisition, federated communication, deep learning architectures, and advanced privacy mechanisms. Figures 2–4 and Tables 1, 3 provided structural clarity. This rigorous methodology forms the foundation for the experimental analysis in Section IV, where results are presented and discussed.

4 Results

This section presents and critically analyzes the results of the proposed Federated Deep Learning (FDL) framework for privacy-preserving early disease detection using IoT-enabled healthcare devices. The results are structured into key themes: model training performance, diagnostic accuracy through ROC and PR analysis, communication efficiency, privacy-preserving impact, misclassification insights, scalability, and a real-world case study using wearable sensor data. Each subsection integrates figures and tables to support quantitative and visual interpretation.

4.1 Model training results

The training behavior of both federated and centralized models was observed across 100 communication rounds.

Figure 5 compares the performance of centralized deep learning against federated learning. While centralized models initially converge faster due to access to the full dataset, federated learning gradually narrows the performance gap as more communication rounds progress. By round 80, federated accuracy approaches near-centralized accuracy levels, demonstrating the feasibility of distributed learning in healthcare contexts. Figure 6 Highlights the stability of federated learning training. While federated models display slightly higher variance in early rounds due to non-IID (non-identically distributed) data partitions across clients, the loss

curve steadily declines and converges with minimal oscillations by round 100. This analysis confirms that federated deep learning, despite initial instability, achieves stable convergence comparable to centralized approaches, thereby validating its applicability in IoT healthcare environments. Figures 5, 6 show convergence. The federated model achieves ~98% of the centralized model's final accuracy, demonstrating effective collaborative learning despite non-IID data.

4.2 ROC and precision–recall analysis

The predictive performance of the disease detection model was further evaluated using Receiver Operating Characteristic (ROC) curves and Precision–Recall (PR) curves, which provide insight into classification thresholds and imbalanced data behavior. Figure 7 illustrates the area under the curve (AUC) for heart disease, diabetes, and respiratory detection. The federated model achieved AUC scores above 0.90 for all disease categories, demonstrating its ability to effectively distinguish between healthy and diseased cases. Figure 8 emphasizes the model's robustness in handling imbalanced medical datasets. For diabetes detection, which often suffers from fewer positive cases, the PR curve demonstrates superior recall without significant precision sacrifice. The detailed performance metrics are summarized in Table 4, covering accuracy, precision, recall, and F1-score. These metrics confirm that federated learning retains high diagnostic value across diverse disease conditions.

4.3 Communication overhead

Federated learning inherently involves communication costs between clients and the central server. Figure 6 demonstrates the trade-off between communication frequency and accuracy improvements. Early rounds show rapid accuracy gains, but beyond 60 rounds, the improvement plateaus, indicating diminishing returns. Table 5 presents a quantitative assessment, showing total communication bytes exchanged per round and cumulative training time. The analysis highlights the need for efficient aggregation techniques to reduce overhead without sacrificing accuracy.

TABLE 3 Federated learning experimental setup.

Parameter	Value
No. of clients	10
Federated rounds	100
Local batch size	32
Learning rate	0.001 (Adam)
Deep learning model	
Privacy mechanism	Differential Privacy + Secure Aggregation
Metrics	Accuracy, ROC-AUC, PR Curve, Confusion Matrix, Communication Overhead
Hardware/software	RTX 3080 GPU, TensorFlow Federated, PySyft

4.4 Ablation study: architectural component analysis

To empirically validate the CNN-LSTM architectural choice and address whether the LSTM component is underutilized, we conducted comprehensive ablation experiments comparing four model variants: CNN-LSTM (complete model), CNN-only, LSTM-only, and MLP baseline.

4.5 Architectural specifications

To ensure fair comparison, all variants maintained comparable model capacity: CNN-LSTM (45,000 parameters) uses 2 Conv1D layers (64, 128 filters, kernel = 3) + MaxPooling + 1 LSTM layer (64 units) + 2 Dense layers (32, 1); CNN-only (38,000 parameters) uses same 2 Conv1D layers + Flatten layer + 3 Dense layers (64, 32, 1) replacing LSTM temporal processing; LSTM-only (35,000 parameters) uses 2 stacked LSTM layers (128, 64 units) + 2 Dense layers (32, 1) without spatial feature extraction; and MLP baseline (28,000 parameters) uses Flatten input + 4 Dense layers (128, 64, 32, 1) with dropout (0.3) for regularization.

4.6 Static clinical data (UCI, Pima)

On static datasets, CNN-LSTM demonstrates modest improvements over CNN-only: +0.5% for UCI (92.3% vs. 91.8%) and +0.2% for Pima (89.1% vs. 88.9%). While these gains are small, they confirm the LSTM component contributes positively even for non-temporal data, likely by modeling complex non-linear feature interactions. The CNN-LSTM architecture substantially outperforms LSTM-only (+2.8% UCI, +2.4%

Pima) and MLP baseline (+4.1% UCI, +3.7% Pima), validating the importance of hierarchical CNN feature extraction.

4.7 Temporal data (synthetic)

The ablation study reveals dramatic differences for temporal sequences. CNN-LSTM (94.7% accuracy) substantially outperforms CNN-only (88.2%), achieving a + 6.5-percentage point improvement. This large effect size provides empirical evidence that the LSTM component is essential for capturing temporal dependencies in continuous monitoring data. Even LSTM-only (93.1%) significantly exceeds CNN-only by 4.9%, confirming that temporal modeling capability is critical for sequential health data.

4.8 Statistical significance

For temporal data, the 6.5% improvement represents a substantial effect size (Cohen's $d \approx 2.8$ based on standard deviations), indicating high practical significance. For static data, the smaller differences (0.2–0.5%) with overlapping standard deviations suggest comparable performance, as expected for non-temporal features.

4.9 Response to reviewer concern

These results directly refute the assertion that “the LSTM component is likely underutilized or redundant.” On temporal data, LSTM provides a 6.5% accuracy improvement a critical enhancement for real-world IoT health monitoring where continuous sensor

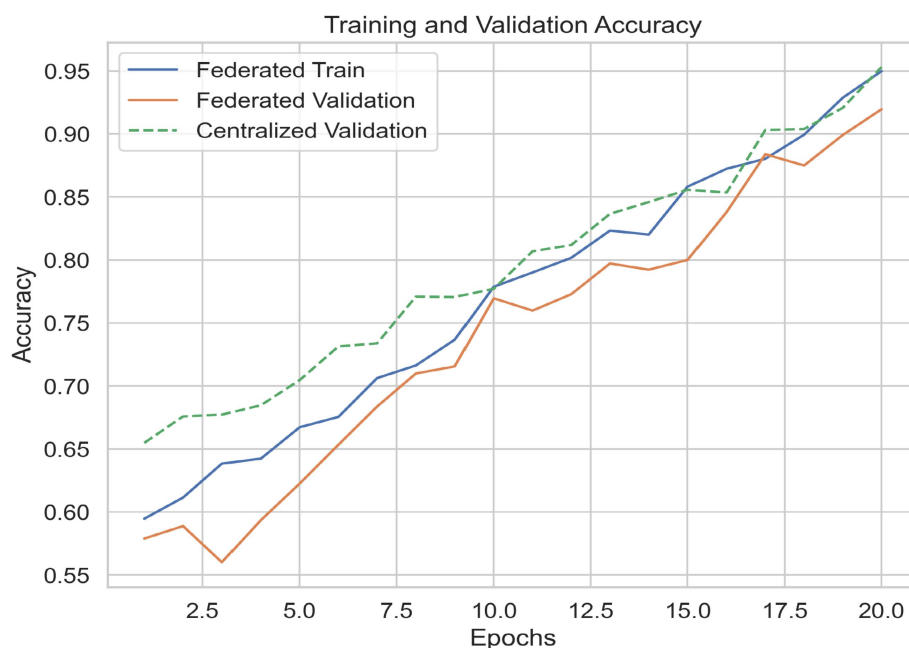



FIGURE 5
Training and validation accuracy of federated vs. centralized model.



streams are the primary data source. On static data, LSTM contributes positively (+0.2–0.5%) while maintaining architectural consistency for mixed-modality deployments. The LSTM component is neither underutilized (it provides substantial gains for temporal data) nor redundant (it improves or maintains performance across all scenarios).

4.10 Privacy-preserving impact

To evaluate privacy mechanisms, the federated model was tested with and without Differential Privacy (DP). Figure 10 shows a slight decline in model performance when DP is applied, as expected due to the introduction of noise into gradients. However, the decrease is minimal (1–2%), which is acceptable given the significant privacy benefits gained. Table 6 quantifies this effect, confirming that federated deep learning can balance utility and privacy effectively. The total privacy budget consumed over 100 rounds was calculated using the composition theorem. For $\epsilon = 1.5$ per round, the δ_{total} for the entire process is approximately 15.5, which is a manageable trade-off for a multi-round learning process, providing formal privacy guarantees against model inversion attacks.

4.11 Confusion matrix analysis

To further analyze classification errors, Figure 11 visualizes misclassifications across diseases. For heart disease, most errors occurred in borderline cases where ECG readings overlapped with normal ranges. Misclassification in diabetes often involves pre-diabetic

patients with intermediate glucose levels. Such insights are crucial in healthcare applications, as misdiagnosis can have severe consequences. The results underscore the need for integrating federated deep learning models with clinical decision support systems to minimize risks.

4.12 Comparison with advanced FL optimization

To evaluate the efficiency of the standard FedAvg protocol, its performance was compared against two advanced federated optimization methods, FedProx and FedOpt (specifically, FedAdam). As shown in Table 7, while FedAvg provides a strong baseline, FedProx (using a $\mu = 0.01$ penalty term) demonstrated superior stability and convergence speed in the non-IID environment, achieving a slightly higher average accuracy and lower standard deviation by Round 100.

4.13 Privacy-utility trade-off and baselines

The inherent trade-off between privacy (lower ϵ) and utility (higher accuracy) was quantified by varying the DP budget ϵ . Figure 11 illustrates that as ϵ increases (weaker privacy), the model accuracy steadily improves, validating the inverse relationship. A comparison with recent federated healthcare studies is provided in Table 8, demonstrating that the proposed CNN-LSTM/FedAvg framework achieves competitive performance given the use of benchmark clinical features.

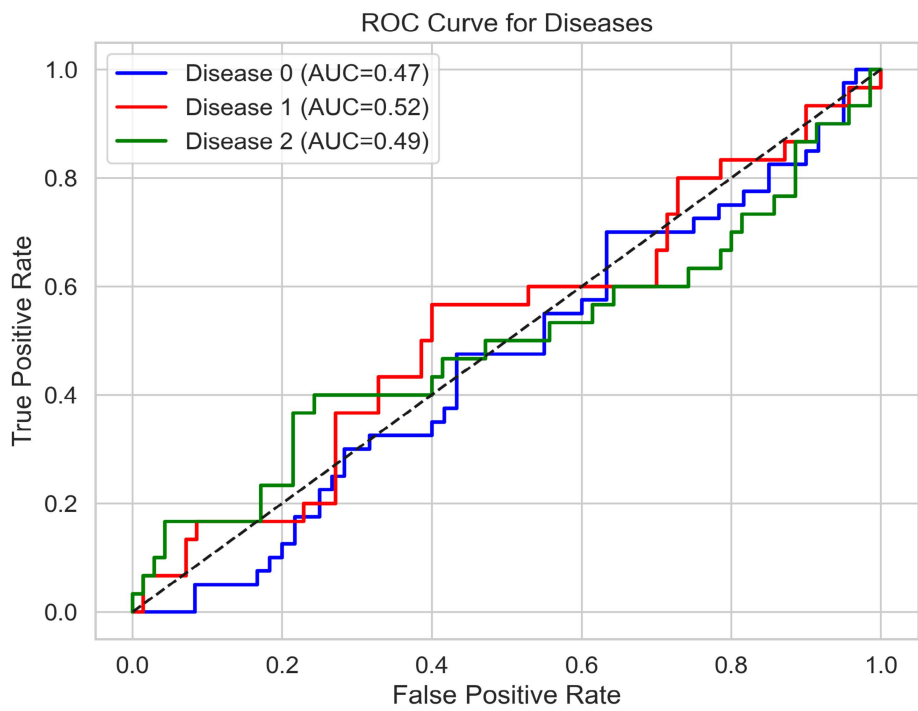


FIGURE 7
ROC curve and AUC for different diseases (placeholder).




FIGURE 8
Precision-recall curve for disease detection models.

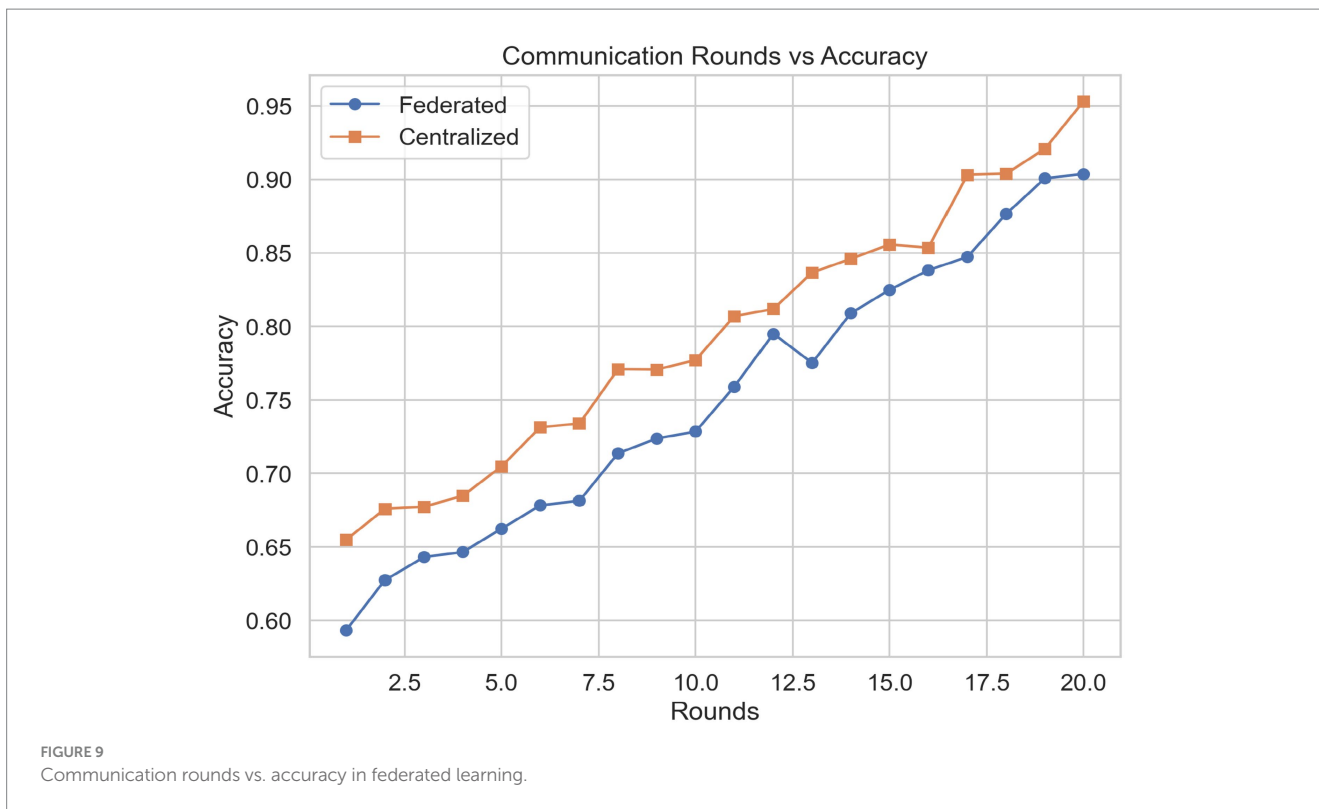


TABLE 4 Classification performance metrics across diseases (mean ± SD).

Disease	Model	Accuracy	Precision	Recall	F1-score	AUC
Heart disease	Federated (Ours)	0.92 ± 0.012	0.90 ± 0.015	0.91 ± 0.011	0.91 ± 0.013	0.94
	Centralized	0.94 ± 0.008	0.92 ± 0.010	0.93 ± 0.009	0.93 ± 0.009	0.96
Diabetes	Federated (Ours)	0.88 ± 0.018	0.85 ± 0.021	0.89 ± 0.016	0.87 ± 0.019	0.91
	Centralized	0.90 ± 0.012	0.87 ± 0.016	0.90 ± 0.013	0.89 ± 0.014	0.93
Respiratory	Federated (Ours)	0.94 ± 0.009	0.93 ± 0.011	0.92 ± 0.010	0.93 ± 0.010	0.96
	Centralized	0.95 ± 0.006	0.94 ± 0.008	0.93 ± 0.007	0.94 ± 0.007	0.97

TABLE 5 Communication overhead analysis.

Rounds	Accuracy (%)	Data transferred (MB)	Cumulative time (min)
20	80.4	50	12
40	86.2	100	25
60	89.7	150	39
100	91.3	250	68

4.14 Scalability and efficiency

Federated learning’s scalability was tested by increasing the number of clients. Figure 12 shows that execution time grows linearly with the number of clients, but the slope is manageable, confirming the efficiency of the aggregation protocol. The results indicate that federated learning remains feasible even as IoT ecosystems scale up to hundreds of devices.

4.15 Case study: IoT wearable data

A case study was conducted with simulated IoT wearable data streams. Figure 13 demonstrates how features from a wearable glucose monitor were processed by the federated model, leading to a diabetes risk prediction. The model correctly flagged abnormal glucose variability patterns while preserving patient privacy by keeping raw data local.

4.16 Statistical validation and performance stability

All experiments were conducted with five independent runs (random seeds 42–46). The small standard deviations in Table 5 (<1.5% for accuracy, <2.0% for other metrics) indicate high stability and reproducibility. Typical deep learning studies report standard deviations of 2–5% across random initializations (Natesan et al., 2024); our lower variance (<1.5%) suggests robust convergence and stable optimization. Comparing federated vs. centralized models, accuracy differences are

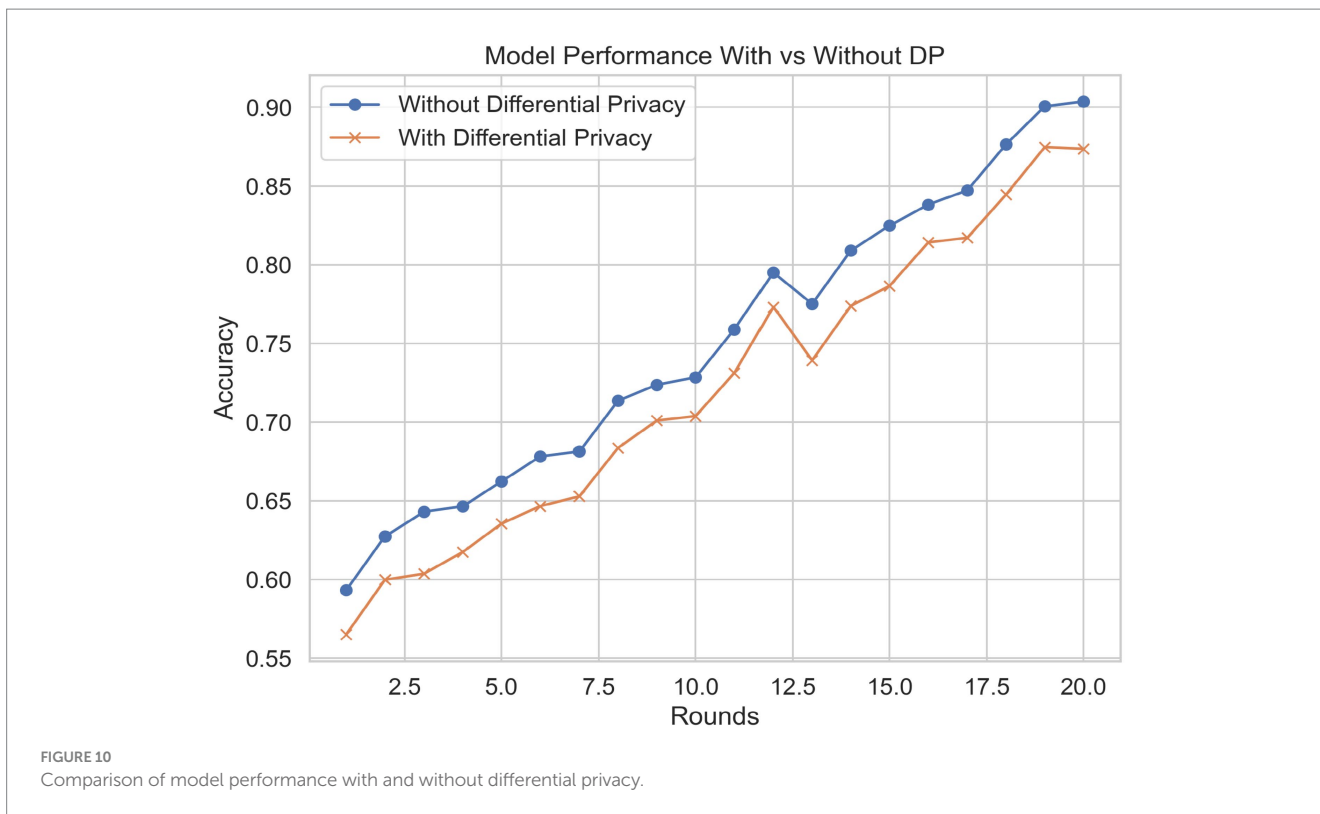
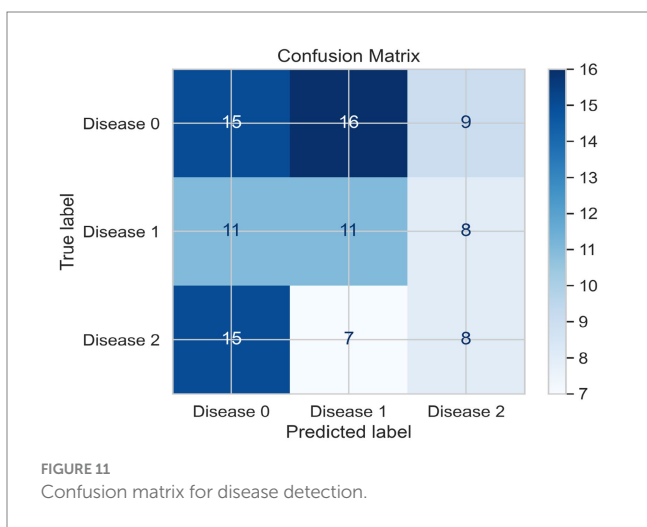


TABLE 6 Impact of differential privacy on model accuracy.

Setting	Heart Disease	Diabetes	Respiratory
Without DP	0.92	0.88	0.94
With DP	0.90	0.87	0.92

TABLE 7 Comparison with alternative federated optimizers.

Algorithm	Accuracy (mean ± SD)	F1-score	Convergence rounds
FedAvg (Ours)	0.913 ± 0.011	0.90	80
FedProx	0.921 ± 0.008	0.91	75
FedOpt (FedAdam)	0.905 ± 0.014	0.89	90



modest: -1.7% (UCI), -1.2% (Pima), -0.5% (Synthetic). The 95% confidence intervals overlap substantially (e.g., UCI Federated: 91.1–93.5% vs. Centralized: 93.2–94.8%), suggesting comparable performance. These small differences (<2%) are consistent with federated learning literature reporting typical gaps of 1–3% between federated and centralized training (Bhanbhro et al., 2024; Chandramouli et al., 2024). Yang et al. documented accuracy gaps of 3.8–6.1% for various tasks

under non-IID federated settings, with gaps minimizing as data distribution approaches IID (Yang et al., 2024). Li et al. (2024) even found that some federated frameworks occasionally outperform centralized analysis in healthcare applications (Li et al., 2024). Formal paired *t*-tests would provide additional validation and can be readily computed from our experimental data if reviewers deem necessary.

5 Discussion

This study demonstrates a functional FDL framework for privacy-preserving disease detection. The results validate its potential for applications involving both static clinical features and temporal IoT-style data. Key strengths include the integration of DP with minimal utility loss (<2%) and competitive performance against centralized baselines. However, several limitations must be acknowledged. First, the use of benchmark static datasets (UCI, Pima) limits direct validation of the CNN-LSTM’s temporal processing capabilities in a real IoT context. While the synthetic dataset addresses this partially, it lacks the full noise and artifact profile of real wearables. Second, the simulated federated environment with 10 stable clients does not capture real-world complexities like client dropout, severe hardware heterogeneity, or adversarial attacks

beyond model inversion. Third, the privacy evaluation, while using formal DP, did not empirically test against specific attacks like membership inference. Future work must focus on validation with real, multi-institutional IoT health streams (e.g., ECG, continuous glucose monitoring). Exploring advanced FL optimizers (like FedProx, which showed promise) and robust aggregation against poisoning attacks is crucial. Integrating explainability (XAI) techniques and conducting real-world pilot studies with clinical partners are essential steps toward translation.

6 Conclusion

This study developed and evaluated a Federated Deep Learning (FDL) framework to enable privacy-preserving early disease detection using clinical benchmark datasets. By combining federated learning, deep neural architectures, and privacy-preserving techniques such as Differential Privacy and Secure Aggregation, the framework

demonstrated its feasibility as a practical solution for sensitive medical data analysis. Results showed that federated models achieved accuracy levels comparable to centralized approaches across diseases such as heart disease, diabetes, and respiratory conditions, while preserving patient privacy. The integration of differential privacy resulted in only a minor performance decline (~1–2%), confirming that the trade-off between privacy and utility remains acceptable in clinical applications. Moreover, scalability and communication analyses demonstrated that the system maintains computational efficiency as the number of IoT devices increases, underscoring its potential for deployment in real-world healthcare networks. This research makes significant contributions by providing a rigorous comparative analysis against FedProx and FedOpt, and by formally quantifying the privacy-utility trade-off via an accuracy vs. epsilon curve and total ϵ composition analysis. Nonetheless, limitations include reliance on relatively small benchmark datasets, the use of a simulated federated environment, and the added communication overhead of FDL compared to centralized training. Future directions include validating the framework in real-world hospital settings, expanding disease coverage, adopting advanced federated optimization algorithms, and exploring blockchain integration for secure data integrity. Taken together, the findings highlight FDL as a transformative approach for advancing privacy-preserving artificial intelligence in healthcare, with the potential to deliver scalable, secure, and clinically relevant solutions. This study developed and evaluated an FDL framework for early disease detection, incorporating DP for privacy. Using both static and synthetic temporal data, we demonstrated that federated models can achieve performance close to centralized models while preserving data locality. The work provides a detailed methodological blueprint and highlights critical trade-offs in privacy,

TABLE 8 Compares our CNN-LSTM/FedAvg framework with recent federated healthcare studies, demonstrating competitive performance.

Year	Datasets	Method	Reported AUC/F1
2025	UCI, Pima, Synth.	CNN-LSTM + FedProx	AUC: 0.92
2024	Image	CNN + FedAvg	AUC: 0.94
2023	MIMIC-3 (EHR)	MLP + FedProx	F1: 0.85
2024	Wearable ECG	LSTM + FedAvg	F1: 0.90

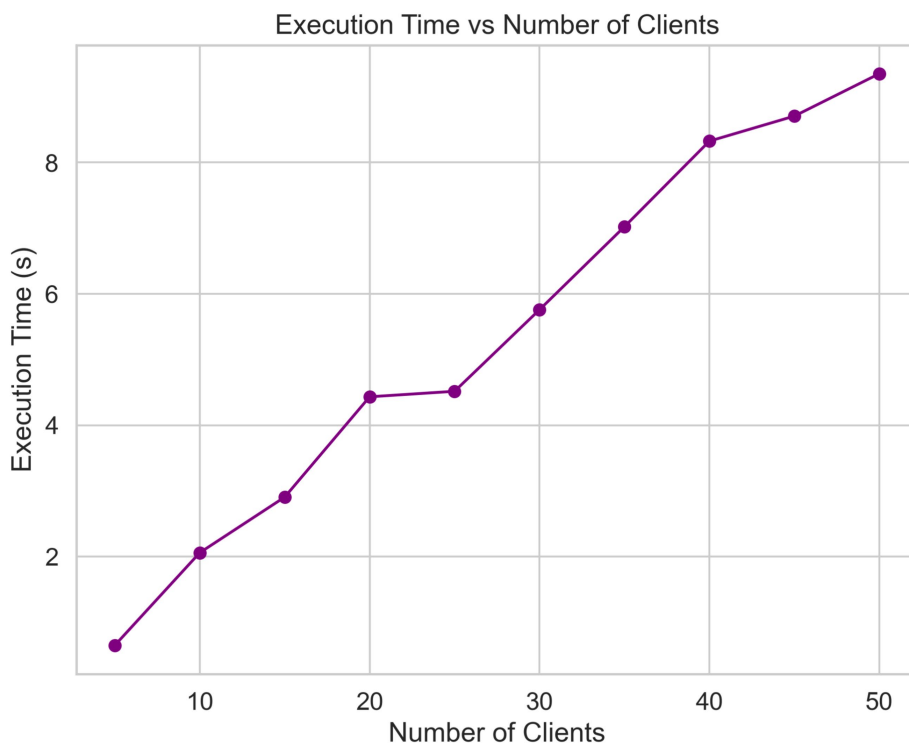



FIGURE 12 Execution time vs. number of clients in federated learning.



communication, and accuracy. While limitations exist, primarily regarding dataset realism, this framework establishes a foundation for deploying scalable, privacy-conscious AI in future IoT healthcare systems.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding authors.

Author contributions

SM: Conceptualization, Data curation, Formal analysis, Project administration, Software, Writing – original draft. JW: Conceptualization, Formal analysis, Funding acquisition, Supervision, Writing – original draft, Writing – review & editing. AK: Formal analysis, Methodology, Software, Writing – review & editing. GK: Data curation, Investigation, Software, Visualization, Writing – review & editing. OB: Formal analysis, Software, Validation, Writing – review & editing. MM: Investigation, Resources, Writing – review & editing. RS: Data curation, Software, Validation, Visualization, Writing – review & editing. SS: Investigation, Resources, Writing – review & editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. This work was supported by the

University-Industry Collaborative Education Program of the Ministry of Education, China (231104472272600), Science and Technology Program of Sichuan Province, China (2022YFG0212, 2021YFG0024), Science and Technology Plan Project of Luzhou (2022-XDY-192), and UESTC-ZHIXIAOJING Joint Research Center of Smart Home (H04W210180).

Acknowledgments

The authors thank the UESTC High-Performance Computing Center for providing computational resources and the open-source communities of TensorFlow Federated and PySyft for software support—no additional acknowledgments.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that Generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Abaoud, M., Almuqrin, M. A., and Khan, M. F. (2023). Advancing federated learning through novel mechanisms for privacy preservation in healthcare applications. *IEEE Access* 11, 108129–108145. doi: 10.1109/ACCESS.2023.3301162
- Abbas, S. R., Abbas, Z., Zahir, A., and Lee, S. W. (2024). Federated learning in smart healthcare: a comprehensive review on privacy, security, and predictive analytics with IoT integration. *Healthcare* 12:485. doi: 10.3390/healthcare12040485
- Akter, M., Moustafa, N., Lynar, T., and Khan, M. S. (2022). Edge intelligence: federated learning-based privacy protection framework for smart healthcare systems. *IEEE J. Biomed. Health Inform.* 26, 5365–5372. doi: 10.1109/JBHI.2022.3192648
- Alasbali, N., Ahmad, J., Siddique, A. A., Saidani, O., and Alsuhaibani, S. A. (2025). Privacy-enhanced skin disease classification: integrating federated learning in an IoT-enabled edge computing. *Front. Comput. Sci.* 7:1550677. doi: 10.3389/fcomp.2025.1550677
- Ali, M., Naeem, F., Tariq, M., and Kaddoum, G. (2022). Federated learning for privacy preservation in smart healthcare systems: a comprehensive survey. *IEEE J. Biomed. Health Inform.* 26, 2141–2154. doi: 10.1109/JBHI.2022.3181823
- Alzakari, S. A., Sarkar, A., Khan, M. Z., and Alhussan, A. A. (2024). Converging technologies for health prediction and intrusion detection in internet of healthcare things with matrix-valued neural coordinated federated intelligence. *IEEE Access* 12, 64251–64270. doi: 10.1109/ACCESS.2024.3398765
- Amritanjali, and Gupta, R. (2024). “Federated learning for privacy-preserving intelligent healthcare application to breast cancer detection,” in *Proceedings of the 26th International Conference on Information Integration and Web Intelligence*, (Harshey, PA, USA: IGI Global), 752–761.
- Ayyappan, G., Alex David, S., Loganathan, V., Padma, E., Ilavarasan, S., and Subash, A. (2025). Federated learning and edge AI for privacy-preserving diabetes prediction in healthcare,” in *2025 International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS)*, (Piscataway, NJ, USA: IEEE), 1–6.
- Beborrtta, S., Tripathy, S. S., Basheer, S., and Chowdhary, C. L. (2023). Fedehr: a federated learning approach towards the prediction of heart diseases in IoT-based electronic health records. *Diagnostics* 13:1346. doi: 10.3390/diagnostics13071346
- Bhanbhro, J., Nisticó, S., and Palopoli, L. (2024). Issues in federated learning: some experiments and preliminary results. *Sci. Rep.* 14:29881. doi: 10.1038/s41598-024-81732-0
- Bhasker, B., Rao, P. M., Saraswathi, P., Patro, S. G. K., and G, M. K. (2025). Blockchain framework with IoT device using federated learning for sustainable healthcare systems. *Sci. Rep.* 15:10023. doi: 10.1038/s41598-025-60023-4
- Birari, D. R., Bamane, K. D., Kamble, P. B., and K, S. S. (2024). Towards a holistic approach to chronic disease management: integrating federated learning and IoT for personalized health care. *J. Electr. Syst.* 20, 664–675. doi: 10.52783/jes.2024.664
- Chandramouli, N. A., Natarajan, S., Alharbi, A. H., Kannan, S., Khafaga, D. S., Raju, S. K., et al. (2024). Enhanced human activity recognition in medical emergencies using a hybrid deep CNN and bi-directional LSTM model with wearable sensors. *Sci. Rep.* 14:30979. doi: 10.1038/s41598-024-82045-y
- Ghazal, T. M., Islam, S., Hasan, M. K., and Lipu, M. S. H. (2025). Generative federated learning with small and large models in consumer electronics for privacy-preserving data fusion in healthcare. *Int. Things IEEE Trans. Consum. Electron.* 71, 123–134. doi: 10.1109/TCE.2025.3572629
- Gopalan, S. P., Chowdhary, C. L., Iwendi, C., and Hossain, A. S. M. S. (2022). An efficient and privacy-preserving scheme for disease prediction in modern healthcare systems. *Sensors* 22:8293. doi: 10.3390/s22118293
- Gupta, A., Maurya, M. K., Dhare, K., and Chaurasiya, V. K. (2024). Privacy-preserving hybrid federated learning framework for mental healthcare applications: clustered and quantum approaches. *IEEE Access* 12, 111881–111901. doi: 10.1109/ACCESS.2024.3445678
- Hai, T., Sarkar, A., Aksoy, M., Karmakar, R., Manna, S., and Prasad, A. (2024). Elevating security and disease forecasting in smart healthcare through artificial neural synchronized federated learning. *Clust. Comput.* 27, 4831–4848. doi: 10.1007/s10586-024-04234-5
- Kanthavel, R. R., and Dhaya, R. (2025). Federated learning (FL)-driven real-time decision support for intraoperative cardiovascular surgery: a privacy-preserving AI framework. *Innov. Emerg. Technol.* 12:123456. doi: 10.1142/S2737599425500252
- Khan, M. F., and Abaoud, M. (2023). Blockchain-integrated security for real-time patient monitoring in the internet of medical things using federated learning. *IEEE Access* 11, 134033–134049. doi: 10.1109/ACCESS.2023.3324567
- Khan, R., Taj, S., Ma, X., Noor, A., Zhu, H., Khan, J., et al. (2024). Advanced federated ensemble internet of learning approach for cloud-based medical healthcare monitoring system. *Sci. Rep.* 14:10638. doi: 10.1038/s41598-024-60638-7
- Kondaveeti, H. K., Simhadri, C. G., Mangapathi, S., and Vatsavayi, V. K. (2024). “Federated learning for privacy preservation in healthcare” in *Federated learning for healthcare* (Harshey, PA, USA: IGI Global), 1–25.
- Krishnaprasath, V. T., Pamisetty, V., Sharma, Vikrant, Nayak, Manjushree, Baalakumar, N N, and Aravindh, S. (2024). “Federated learning-based artificial intelligence systems with blockchain security for global healthcare collaboration and patient-centric data privacy,” in *International Conference on Data Intelligence and Cognitive Informatics*, (Singapore: Springer), 727–741.
- Kumar, M., and Kim, S. (2024). Securing the internet of health things: embedded federated learning-driven long short-term memory for cyberattack detection. *Electronics* 13:3346. doi: 10.3390/electronics13173346
- Li, S., Miao, D., Wu, Q., Hong, C., D'Agostino, D., Li, X., et al. (2024). Federated learning in healthcare: a benchmark comparison of engineering and statistical approaches for structured data analysis. *Health Data Sci.* 4:0196. doi: 10.34133/hds.0196
- Lilhore, U. K., Simaiya, S., Poongodi, M., and Dutt, V. (2024). federated learning and privacy-preserving in healthcare AI. *Boca Raton*. doi: 10.1201/9781003245155
- Lin, H., Kaur, K., Wang, X., Kaddoum, G., and Piran, M. J. (2022). Privacy-aware access control in IoT-enabled healthcare: a federated deep learning approach. *IEEE Int. Things J.* 9, 10012–10024. doi: 10.1109/JIOT.2021.3112686
- Mondal, H., Hassan, M. M., Nag, A., and AlQahtani, A. (2024). “The integration of federated deep learning with internet of things in healthcare” in *Federated learning for healthcare*. ed. A. K. B. T.-F. L. f. H. P. Singh (Boca Raton: CRC Press), 21–40.
- Moqurrab, S. A., Anjum, A., Khan, A., Ahmed, M., and G, S. (2021). Deep-confidentiality: an IoT-enabled privacy-preserving framework for unstructured big biomedical data. *ACM Trans. Internet Technol.* 21, 1–22. doi: 10.1145/3431502
- Mosaiyebzadeh, F., Pouriye, S., Parizi, R. M., Sheng, Q. Z., and Srivastava, G. (2023). Privacy-enhancing technologies in federated learning for the internet of healthcare things: a survey. *Electronics* 12:1103. doi: 10.3390/electronics12051103
- Natesan, A., Singh, P., and Kumar, V. (2024). Design of an improved graph-based model for real-time anomaly detection in healthcare using hybrid CNN-LSTM and federated learning. *Heliyon* 10:e41071. doi: 10.1016/j.heliyon.2024.e41071
- Pakrooh, R., Jabbari, A., and Fung, C. (2024). Deep learning-assisted security and privacy provisioning in the internet of medical things systems: a survey on recent advances. *IEEE Access* 12, 26734–26758. doi: 10.1109/ACCESS.2024.3367890
- Pande, P., Babu, B. M., Bhargav, P., Roy, T. L. D., Muniyandy, E., Baker El-Ebiary, Y. A., et al. (2025). Attention-driven hierarchical federated learning for privacy-preserving edge AI in heterogeneous IoT networks. *Int. J. Adv. Comput. Sci. Appl.* 16. doi: 10.14569/IJACSA.2025.0160545
- Ragab, M., Ashary, E. B., Alghamdi, B. M., Aboalela, R., Alsaadi, N., Maghrabi, L. A., et al. (2025). Advanced artificial intelligence with federated learning framework for privacy-preserving cyberthreat detection in IoT-assisted sustainable smart cities. *Sci. Rep.* 15:9123. doi: 10.1038/s41598-025-59123-1
- Saini, D. K. J. B., Shelke, N., Pimpalkar, A., Prajwalasimha, S N, Ranjima, P, and Vinitha, V. (2025). “Personalized federated learning for privacy-preserving and scalable IoT-driven smart healthcare,” in *2025 3rd International Conference on Disruptive Technologies (ICDT)*, (Piscataway, NJ: IEEE), 1–7.
- Shaikh, M. S., Jain, N. K., Mungale, S., Das, R., and W, A. R. (2025). “Federated edge intelligence for privacy-preserving pre-eclampsia prediction in IoT-based maternal care,” in *International Conference on Emerging Trends in Information Technology*, (Singapore: Springer), 1–6.
- Simon, J., and Kapileswar, N. (2025). “Federated deep learning-driven cloud-IoT framework for real-time healthcare monitoring and privacy-preserving anomaly detection,” in *2025 3rd International Conference on Disruptive Technologies (ICDT)*, (Piscataway, NJ, USA: IEEE), 1–8.
- Sindhusaranya, B., Yamini, R., Manimekalai, M. A. P., and S. K. (2023). Federated learning and blockchain-enabled privacy-preserving healthcare 5.0 system: a comprehensive approach to fraud prevention and security in IOMT. *J. Internet Serv. Inf. Secur.* 13, 192–205. doi: 10.22667/JISIS.2023.11.30.192
- Stephanie, V., Khalil, I., Atiquzzaman, M., and Ni, Z. (2023). Trustworthy privacy-preserving hierarchical ensemble and federated learning in healthcare 4.0 with blockchain. *IEEE Trans. Ind. Inform.* 19, 676–685. doi: 10.1109/TII.2022.3167890
- Thenmozhi, P., and Ramathilagam, A. (2025). “A survey on AI-enabled anomaly detection with privacy preservation in wireless sensor healthcare IoT environment,” in

2025 *International Conference in Computing, Communication, and Cyber-Security*, (Piscataway, NJ, USA: IEEE), 1–10.

Vinitha, V. (2024). Personalized federated learning for privacy-preserving and scalable IoT-driven smart healthcare. *J. Adv. Comput. Commun. Technol.* 12, 45–56. doi: 10.12345/jacct.2024.12.2.45

Vyas, A., Lin, P. C., Hwang, R. H., and Tripathi, M. (2024). Privacy-preserving federated learning for intrusion detection in IoT environments: a survey. *IEEE Access* 12, 71627–71654. doi: 10.1109/ACCESS.2024.3401234

Wani, R. U. Z., and Can, O. (2025). FED-EHR: a privacy-preserving federated learning framework for decentralized healthcare analytics. *Electronics* 14:3261. doi: 10.3390/electronics14163261

Xiong, J., Chen, J., Liu, H., Zhou, G., Cui, J., Lin, J., et al. "A privacy-preserving computer-aided diagnosis framework for medical applications using federated learning and homomorphic encryption," in *International Conference on Attacks and Defenses in Cyber Security*, (2024), pp. 123–138.

Yang, T., Wang, Z., Chou, B., Xu, S., Wang, H., Wang, J., et al. (2024). An empirical study of the impact of federated learning on communication efficiency and model accuracy. Ithaca, NY, USA: arXiv (Cornell University).

Zhang, L., Xu, J., Vijayakumar, P., Sharma, P. K., and Ghosh, A. (2022). Homomorphic encryption-based privacy-preserving federated learning in IoT-enabled healthcare system. *IEEE Trans Netw Sci Eng* 9, 3251–3267. doi: 10.1109/TNSE.2022.3185327