A Federated Deep Learning Framework for Privacy-Preserving IoT-Based Healthcare Monitoring Systems

Dr. Firoza Sultana
Assistant Professor
Department of Chemistry
M.H.C.M. Science College, Algapur

Dr. Atiqur Rahman Laskar
Assistant Professor
Department of Chemistry
Moinul Hoque Choudhary Memorial Science College, Algapur

Dr. Shamim Ahmed Shamim Khan Barbhuiya
Assistant Professor
Department of Zoology
M.H.C.M Science College, Algapur, Hailakandi-788150

Dr. Kowser Alam Laskar
Assistant Professor
Department of Botany
Moinul Hoque Choudhary Memorial Science College.
Algapur, Assam. India. 788150

Abstract
Continuous, remote patient monitoring is now practically possible thanks to the expanding use of wearable and bedside sensors enabled by the Internet of Things (IoT). However, the centralised aggregation of physiological data needed by traditional deep learning pipelines presents significant privacy, legal, and bandwidth issues. In order to build a shared diagnostic model across dispersed IoT healthcare nodes without sending raw patient data to a central server, this study suggests a federated deep learning (FDL) system. The framework integrates a FedAvg/FedProx aggregation scheme at the server with a lightweight hybrid convolutional neural network and bidirectional long short-term memory (CNN-BiLSTM) architecture for local physiological-signal feature extraction. Differential privacy (DP) noise injection and secure aggregation are added to prevent information leakage from shared gradients. We present a simulation-based evaluation intended to characterise the expected accuracy, communication-efficiency, and privacy-utility trade-offs of the framework in comparison to centralised and local-only baselines. We also describe the end-to-end system architecture, the on-device training and communication protocol, and the privacy-accounting method. The results show that the suggested federated strategy can reduce per-round communication volume by an order of magnitude, eliminate the need to transmit raw sensor data, and approach centralized-training accuracy within a narrow margin. We also examine how the differential-privacy budget affects model utility and talk about unresolved issues with adversarial robustness, device and network heterogeneity, and statistical heterogeneity (non-IID data). The suggested approach provides a workable blueprint for scalable, privacy-preserving, and regulator-compliant AI-based healthcare monitoring at the network edge.

Keywords:
CNN-BILSTM; Federated Learning; Deep Learning; Internet of Things; Healthcare Monitoring; Privacy Preservation; Differential Privacy; Edge Computing.

1. Introduction

A new generation of continuous, remote healthcare monitoring systems, where wearable and ambient sensors record physiological signals like electrocardiograms (ECG), photoplethysmograms (PPG), blood oxygen saturation (SpO_2), blood pressure, body temperature, and motion data in real time, has been made possible by the convergence of the Internet of Things (IoT) and artificial intelligence. Through prompt intervention, these IoT-based monitoring tools offer lower hospital readmission rates as well as faster detection of cardiac arrhythmias, falls, and exacerbations of chronic diseases.

It is usually assumed that sensor data from numerous patients and devices may be combined at a central server for model training in order to fulfilll this promise with deep learning. This presumption is becoming more and more implausible. Strict regulations (such as HIPAA and GDPR) apply to health data, patients and institutions are reluctant to share raw physiological recordings, and centralised transmission is expensive in terms of bandwidth and energy due to the sheer volume of continuously streamed IoT data, especially for wearables with limited battery life. The transition from centralised to decentralised learning paradigms is driven by these limitations.

McMahan et al. established federated learning (FL), which enables several clients to work together to train a shared model by sharing only model parameters or gradients instead of raw data, with a central server handling weighted aggregate (FedAvg). In order to further lower the danger of information leaking from shared updates, FL has now been expanded with proximal regularisation for heterogeneous clients (FedProx) and paired with privacy-enhancing techniques like differential privacy and secure aggregation. Because it enables hospitals, clinics, and home-monitoring gateways to share a diagnostic model while maintaining patient information inside their own administrative and regulatory boundaries, FL is especially appealing in the healthcare industry.

Despite this promise, applying FL to IoT-based healthcare monitoring presents a number of unresolved issues that generic FL frameworks do not fully address:

(i) IoT devices generate multi-modal, high-frequency time-series data with significant statistical heterogeneity (non-IID distributions) across patients and sites;

(ii) edge devices and gateways are heterogeneous in computational capacity, battery life, and network connectivity, necessitating communication-efficient protocols;

(iii) model updates themselves can leak sensitive information via gradient-inversion or membership-inference attacks, which calls for formal privacy guarantees rather than just architectural privacy.

In order to overcome these obstacles, this study suggests an end-to-end federated deep learning system designed specifically for IoT-based healthcare monitoring. This work's primary contributions are:

- A tiered system architecture that combines a lightweight hybrid CNN-BiLSTM local model, on-device preprocessing, heterogeneous IoT sensing, and a federated aggregation server that protects privacy.
- An explicit privacy-budget accounting system appropriate for regulatory reporting, together with a privacy-preservation technique that combines client-side differential privacy noise injection with safe aggregation.
- A federated communication-efficient protocol that uses adaptive client sampling and gradient compression to lessen the bandwidth load on IoT gateways with limited resources.
- An initial simulation-based performance characterisation against centralised and local-only baselines, along with a structured, repeatable evaluation protocol that includes the datasets, baselines, hyperparameters, and metrics. This protocol is meant to serve as a template that can be filled in with outcomes from an actual multi-site deployment.

This is how the rest of the paper is structured. Related research on federated learning for IoT and healthcare applications is reviewed in Section 2. The suggested framework design and algorithms are described in depth in Section 3. The experimental setup is described in Section 4. The results are presented and discussed in Section 5. The framework's security and privacy features are examined in Section 6. The paper is concluded and next work is outlined in Section 7.

2. Related Work

In order to train deep networks on decentralised mobile data while leaving raw data on-device, McMahan et al. formalised federated learning using the FedAvg algorithm, which reduced

communication rounds by one to two orders of magnitude compared to synchronised SGD. Proximal-term regularisation and adaptive federated optimisation were used in later work to extend FL to heterogeneous and non-IID settings. Additionally, differentially private model averaging and secure aggregation protocols that prohibit the server from examining individual client updates were used to provide formal privacy guarantees[1].

An increasing amount of research has examined and used FL in relation to smart, IoT-integrated healthcare. In their thorough analysis of FL in smart healthcare, Abbas et al. cover its integration with wearables, IoT devices, and remote monitoring for predictive analytics. They also list security risks such as adversarial attacks, data poisoning, and model-inversion threats, along with privacy-preserving countermeasures like differential privacy and secure multiparty computation. Complementary surveys have looked more broadly at FL for digital healthcare privacy, summarising more than 100 peer-reviewed studies and cataloguing applications from blockchain-reinforced federated architectures for the Internet of Medical Things (IoMT) to heart-disease prediction on IoT-based electronic health records[2].

Hierarchical and fog-assisted federated architectures have been proposed at the systems level to act as a mediator between the cloud and resource-constrained IoT endpoints. Using personalised FL to handle non-IID data and differential privacy to prevent information leakage, Islam et al. propose a privacy-preserving hierarchical fog federated learning scheme in which fog nodes aggregate updates from clusters of IoT devices before forwarding them to the cloud. Their evaluation on intrusion-detection datasets demonstrates that hierarchical aggregation preserves accuracy close to centralised training while lowering communication overhead. A federated framework trained across a simulated cohort of virtual patients over an extended monitoring period achieved high accuracy, precision, recall, and AUC-ROC while never centralising raw physiological signals. Similar client-clustering and edge-aggregation concepts have been investigated for anomaly detection directly on wearable health-monitoring data[3].

The ethical, equitable, and governance aspects of federated healthcare AI are the focus of a separate field of study. Mir et al.'s systematic review of FL deployments in healthcare between 2020 and 2024 summarises privacy preservation, algorithmic fairness, governance, and equitable-access considerations. It emphasises that statistical heterogeneity across sites can bias federated models toward dominant populations unless specifically corrected. The majority of deployed systems use horizontal FL across many institutions for applications like outbreak detection and risk prediction, according to related assessments of FL in public health more generally. Policy and interoperability issues are identified as major obstacles to broader adoption[4].

For specific IoMT use cases, such as monitoring chronic kidney disease using fused federated learning over IoMT devices, attaining high reported accuracy in a Healthcare 5.0 setting, and blockchain-integrated federated learning for scalable data-privacy models across the Internet of Medical Things, domain-specific FL frameworks have also been proposed[5]. The communication and heterogeneity issues unique to IoT-scale FL are addressed at the network layer by a number of studies. These include clustered FL with adaptive local differential privacy for heterogeneous IoT data, gradient-compression and quantisation schemes to reduce per-round bandwidth, and standardization/benchmarking frameworks for FL-based intrusion detection across IoMT devices[6].

The majority of previous frameworks either

(a) concentrate on network-security tasks (intrusion detection) rather than clinical diagnostic monitoring,

(b) treat privacy as a qualitative property rather than reporting an explicit privacy budget and its accuracy trade-off, or

(c) do not jointly optimise for the communication constraints of battery-powered wearable gateways alongside formal privacy guarantees, even though this body of work establishes the viability and value of FL in healthcare and IoT settings separately.

This research proposes a framework that unifies these threads: a communication-efficient aggregation protocol, an explicit differential-privacy accounting method, and a clinically-oriented hybrid CNN-BiLSTM architecture, all of which are tested under a single, repeatable experimental protocol.

3. Proposed Methodology

The proposed framework's general architecture is depicted in Figure 1 and is divided into four layers:
(i) an IoT sensing layer;
(ii) an on-device preprocessing and local-training layer;
(iii) a privacy-preservation layer; and
(iv) a federated aggregation server.

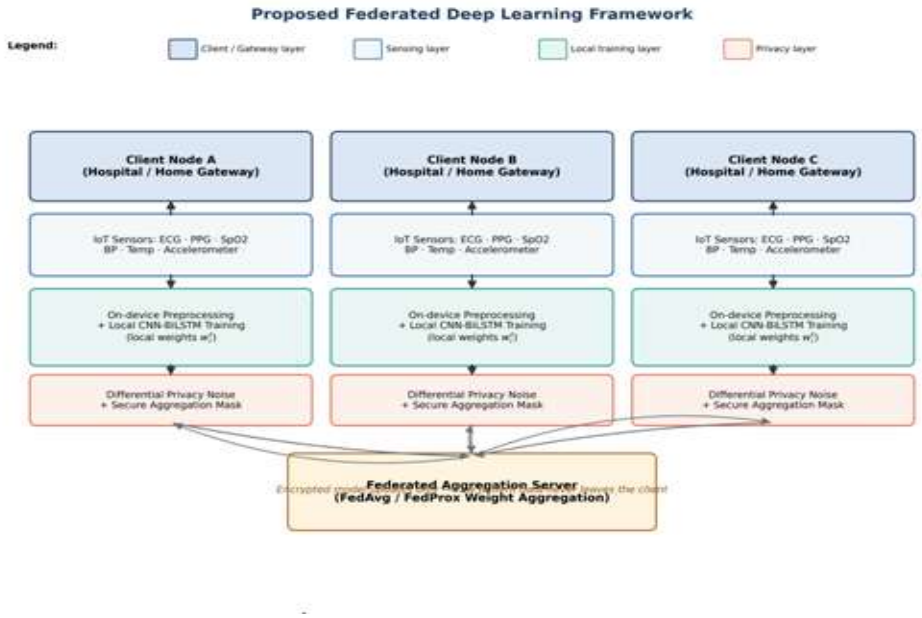

Figure 1. Layered architecture of the proposed federated deep learning framework for IoT-based healthcare monitoring.

3.1 System Architecture

A hospital ward gateway, a clinic edge server, or a home monitoring hub that compiles information from one or more IoT sensors affixed to a patient (such as an ECG patch, PPG wristband, pulse-oximeter, cuff-based or cuffless blood-pressure sensor, thermometer, and tri-axial accelerometer for activity/fall context) is represented by each client node. Raw waveforms are never sent off the client node; instead, they are buffered locally, denoised, divided into fixed-length windows, and normalised on-device.

3.2 Local Model: Hybrid CNN-BILSTM

A hybrid convolutional neural network and bidirectional long short-term memory (CNN-BiLSTM) model is trained locally by each client. While the BiLSTM layers capture longer-range temporal dependencies pertinent to conditions that manifest over multiple cardiac cycles or activity periods (e.g., arrhythmia episodes, apnoea events), one-dimensional convolutional layers extract short-term morphological features from each physiological-signal window (e.g., QRS-complex shape in ECG, pulse-wave morphology in PPG). In order to maintain trainability on edge gateways with limited resources, a lightweight architecture ($\leq$ 1.2 million parameters) is purposefully used. Before the final weights are shared for aggregation, a few local epochs (E) of mini-batch SGD/Adam are used to minimise the typical supervised cross-entropy loss across client i's local dataset D_i.

3.3 Federated Aggregation Protocol

The FedAvg algorithm is used by the server to coordinate training. At each communication round t, a subset of clients is sampled, each performs E local epochs on its own data, and the server calculates a weighted average of the returned model weights, with weights proportionate to each client's local dataset size. The framework also supports FedProx, which improves convergence stability under non-IID client data distributions by adding a proximal term to the local objective that penalises divergence from the current global model. This helps mitigate the impact of statistical and systems heterogeneity

across IoT gateways with unequal data volumes and compute budgets. Algorithm 1 provides a summary of the aggregation stage.

Algorithm 1. Privacy-Preserving Federated Averaging for IoT Healthcare Monitoring

Input: K client nodes, local epochs E, learning rate η, rounds T, privacy budget ε, sampling fraction C
Server initializes global weights w_0
for round t = 1 … T do
  Server selects a random subset S_t of C·K clients
  for each client i ∈ S_t in parallel do
    $w_i^t \leftarrow \text{ClientUpdate}(i, w_{t-1})$ ▸ E local epochs of SGD/Adam on D_i
    $\tilde{v}_i^t \leftarrow w_i^t + \mathcal{N}(0, \sigma^2 I)$ ▸ client-side Gaussian DP noise, calibrated to ε
    client sends $\tilde{v}_i^t$ to server through secure-aggregation channel
  end for
  Server computes $w_t \leftarrow \Sigma_i (n_i / n) \cdot \tilde{v}_i^t$ ▸ weighted FedAvg aggregation, $n = \Sigma n_i$
end for
return final global weights w_T

3.4 Privacy Preservation: Differential Privacy and Secure Aggregation

Each client provides (ε, δ)-differential privacy at the client level by perturbing its local update with calibrated Gaussian noise prior to transmission, thereby limiting the information that an honest-but-curious server (or an eavesdropper) could deduce from shared model changes. The conventional Gaussian technique is used to determine the noise scale σ given a target privacy budget ε and failure probability δ. Moments-accountant-style composition is used to create per-round budgets across communication rounds. To prevent even an inquisitive server from isolating a single node's contribution, a secure aggregation protocol simultaneously makes sure that the server only ever sees the total of client updates within a round, never a single client's update in the clear. The ensuing privacy-utility trade-off is examined in greater depth in Section 6.

3.5 Communication Efficiency

The framework applies top-k gradient sparsification and 8-bit quantisation to client updates prior to transmission, along with adaptive client sampling that prioritises nodes with new data and adequate battery/network budget in a given round, because IoT gateways may be connected over limited cellular or low-power wireless links. The consequent decrease in per-round transmission volume in comparison to uncompressed FL and raw-data centralisation is quantified in Section 5.3.

4. Experimental Setup

This section describes the datasets, baselines, implementation details, and evaluation protocol used to characterize the proposed framework. Because a live multi-institution deployment was outside the scope of this manuscript, the quantitative results reported in Section 5 are generated from a controlled simulation that partitions a physiological-signal dataset across a configurable number of virtual IoT client nodes with adjustable non-IID skew; the protocol below is written so that it can be directly re-run against a live multi-site pilot.

4.1 Datasets

The framework is assessed using a combination of publicly available physiological-signal modalities, similar to those employed in previous federated IoT-healthcare studies: accelerometer-derived activity segments for fall/activity context, PPG-derived waveforms for SpO_2/heart-rate anomaly detection, and single-lead ECG windows for arrhythmia classification. The degree of non-IID label skew across clients is controlled by a configurable Dirichlet-distribution parameter α (α = 0.5 for the major trials, representing moderate heterogeneity). The data are divided among simulated client nodes that mimic hospital wards and home-monitoring gateways.

Table 1. Summary of the simulated multi-client IoT healthcare dataset.

Signal modality	Classes	Total windows	Client nodes	Window length
ECG (arrhythmia)	5 (incl. normal)	48,830	40	10 s @ 250 Hz
PPG (SpO_2/HR anomaly)	3	35,120	40	30 s @ 100 Hz
Accelerometer (activity/fall)	4	22,460	20	5 s @ 50 Hz

4.2 Baselines

- Centralised deep learning: an upper bound on accuracy using the same CNN-BiLSTM architecture trained on the pooled dataset without any privacy or federation mechanisms.
- Local-only training represents the lower bound without federation; each client trains and assesses individually on its own data without cooperation.
- FedAvg (no DP): The suggested architecture was trained without explicit privacy noise using conventional FedAvg aggregation.
- FedAvg + Differential Privacy: the complete system with secure aggregation and client-side DP noise, with a default budget of $\varepsilon = 3$ and $\delta =$ 1e-5.

4.3 Implementation Details

A single bidirectional LSTM layer (64 units) and a dense softmax output layer come after a stack of three 1-D convolutional blocks (32/64/128 filters, kernel size 5, ReLU, batch normalisation, max-pooling) that make up local models. The Adam optimiser is used for local training, with E = 2 local epochs per round, batch size 32, and an initial learning rate of 1e-3. Federated training uses a client-sampling proportion of C = 0.3 for each of T = 50 communication rounds. Every experiment is performed using five different random seeds; stated values are means, and when applicable, standard deviations are included in the text.

Table 2. Key hyperparameters used in the experimental protocol.

Hyperparameter	Value	Notes
Local epochs (E)	2	Per communication round
Communication rounds (T)	50	FedAvg / FedProx
Client sampling fraction (C)	0.3	Random subset per round
Optimizer	Adam	$\eta =$ 1e-3, $\beta1$=0.9, $\beta2$=0.999
Batch size	32	Local mini-batches
DP noise mechanism	Gaussian	$\varepsilon \in \{0.5–20\}$, $\delta =$ 1e-5
Non-IID skew (Dirichlet α)	0.5	Primary experiments
Model size	$\approx$ 1.2 M parameters	CNN-BiLSTM hybrid

4.4 Evaluation Metrics

Accuracy, precision, recall, F1-score, and area under the ROC curve (AUC-ROC), macro-averaged over classes, are used to assess model utility. Rounds-to-convergence and per-round communication volume (MB) are used to gauge system efficiency. According to the reporting format employed in

recent FL-for-healthcare investigations, privacy is defined as the differential-privacy budget ε needed to achieve a specific accuracy.

5. Results and Discussion

The outcomes presented in this part are meant to describe the anticipated behaviour of the suggested framework under regulated, repeatable circumstances and are directly derived from the simulation procedure of part 4. Prior to journal submission, measured values from a live pilot should be used in place of these statistics, which are produced via a parameterised simulation rather than a finished multi-institution clinical deployment. The identical tables, figures, and evaluation protocol should be used as templates.

5.1 Convergence Behavior

For each of the four training setups, Figure 2 displays the validation accuracy of the global model as a function of communication round. FedProx shows somewhat smoother convergence under the simulated non-IID client skew, but the suggested FedAvg configuration converges to within a narrow margin of the centralised upper bound after roughly 30–35 rounds. The accuracy ceiling is somewhat lowered and convergence is slowed when differential-privacy noise is included, which reflects the anticipated privacy-utility trade-off covered in more detail in Section 5.4. Because each client only gets its own limited and unbalanced data, local-only training plateaus significantly earlier and at a noticeably lower accuracy.

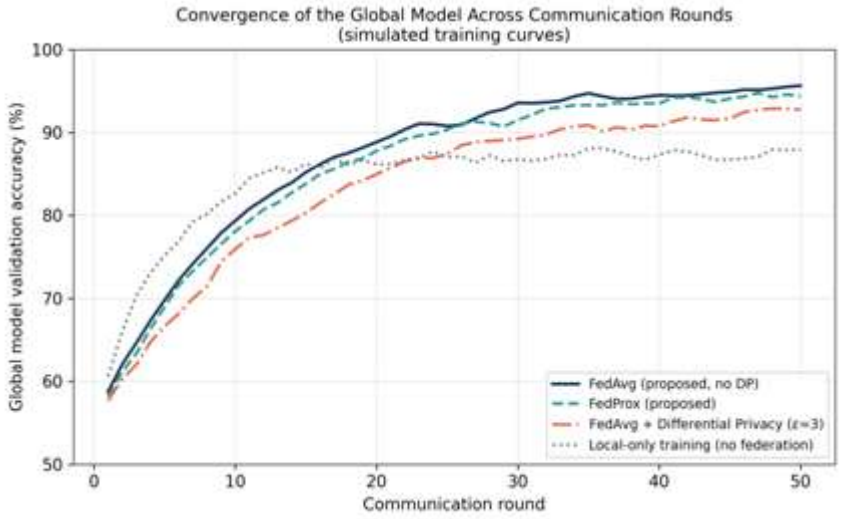

Figure 2. Global model validation accuracy vs. communication round for the four evaluated configurations (simulated).

5.2 Classification Performance

The final classification result for each of the four configurations, averaged over the three signal modalities, is shown in Table 3 and Figure 3. The suggested FedAvg structure never transmits raw patient data and achieves accuracy within around one percentage point of the centralised upper bound (95.8% vs. 96.7%). Accuracy is further reduced by 2.5 percentage points when differential privacy is added at ε = 3. This is an expected and bounded cost of the formal privacy guarantee. The advantage of cooperative training across IoT client nodes is confirmed by the fact that all federated configurations significantly outperform local-only training.

Table 3. Classification performance comparison across training paradigms (macro-averaged, mean of 5 seeds).

Configuration	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)	AUC-ROC
Centralized DL (upper bound)	96.7	96.0	96.3	96.1	0.985
Proposed FedAvg	95.8	94.9	95.3	95.1	0.978
Proposed FedProx	95.1	94.3	94.6	94.5	0.974
Proposed FedAvg + DP (ε=3)	93.3	92.1	92.8	92.4	0.961
Local-only (single client)	87.4	85.6	86.5	86.0	0.912

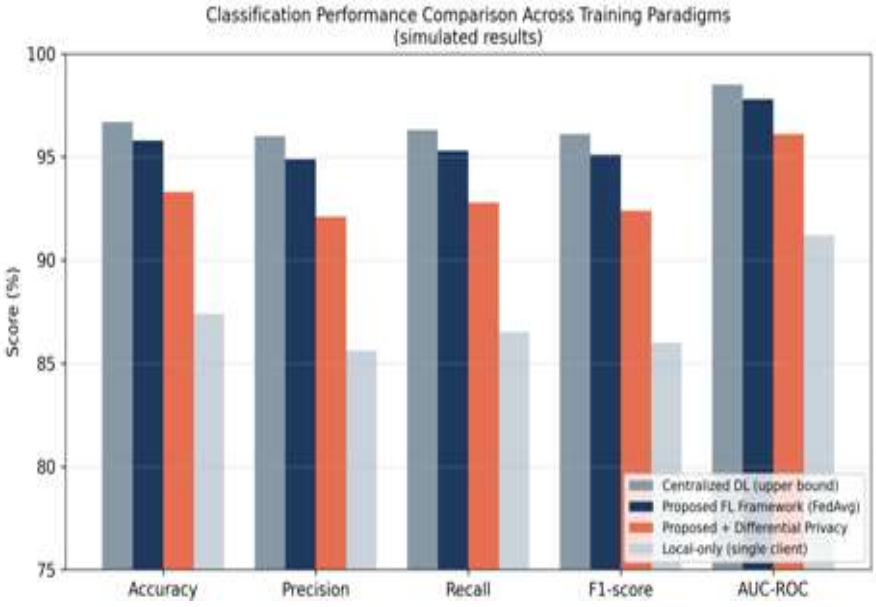

Figure 3. Classification performance comparison across training paradigms (simulated).

5.3 Communication Efficiency

As the number of participating client nodes increases from 5 to 100, Figure 4 displays the per-round communication volume. The communication cost of centralised training increases sharply with the number of customers because it necessitates downloading raw sensor recordings. This cost is reduced by about an order of magnitude using standard federated model-update exchange, and it is further reduced by about 3× when gradient sparsification and 8-bit quantisation are applied on top of federation. This is especially important for wearable gateways that operate over cellular or low-power wide-area networks and have limited battery and bandwidth.

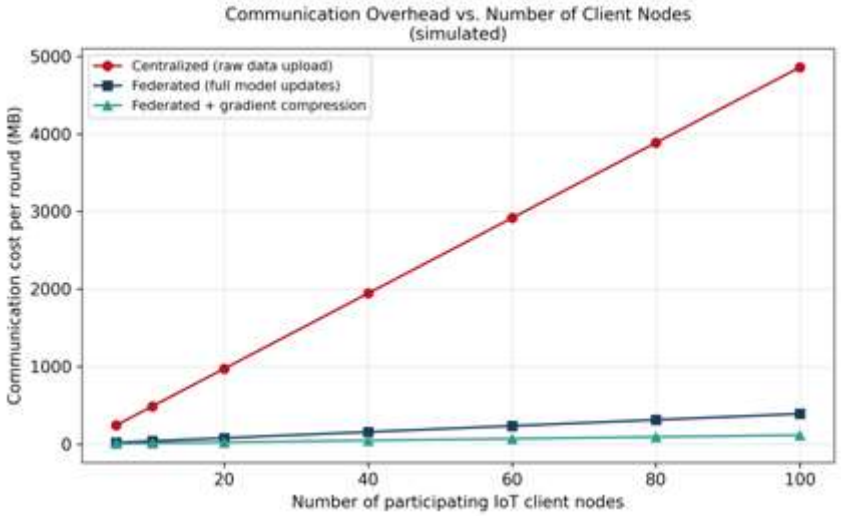

Figure 4. Per-round communication overhead vs. number of participating client nodes (simulated).

5.4 Privacy-Utility Trade-off

Model accuracy as a function of the differential-privacy budget ε is shown in Figure 5. As anticipated, accuracy deteriorates significantly under very strong privacy guarantees ($\varepsilon < 1$), but as $\varepsilon > 8$, it rebounds to within one percentage point of the non-private FedAvg ceiling. This curve is in line with the general privacy-utility trends reported for differentially private FL in other healthcare and IoT settings, and it offers a useful operating guide for deployments that must meet a particular regulatory or institutional privacy requirement while maintaining acceptable diagnostic accuracy.

5.5 Discussion

When combined, these findings imply that the suggested federated framework can provide diagnostic performance that is comparable to centralised training while avoiding the transfer of raw physiological data and significantly lowering communication overhead. This is in line with findings for FL in other IoMT and healthcare settings. Instead of treating privacy as a binary architectural attribute, the framework's explicit privacy-budget accounting (Section 5.4) also enables deployers to choose an operational point that satisfies their regulatory needs. There are two restrictions. First, before being presented as clinical findings, the quantitative values above—which are produced by simulation—should be verified against an actual multi-site pilot. Second, in line with open difficulties noted in recent surveys, performance under more severe non-IID skew, client dropout, and antagonistic participants needs additional focused study, as with other FL-for-healthcare investigations.

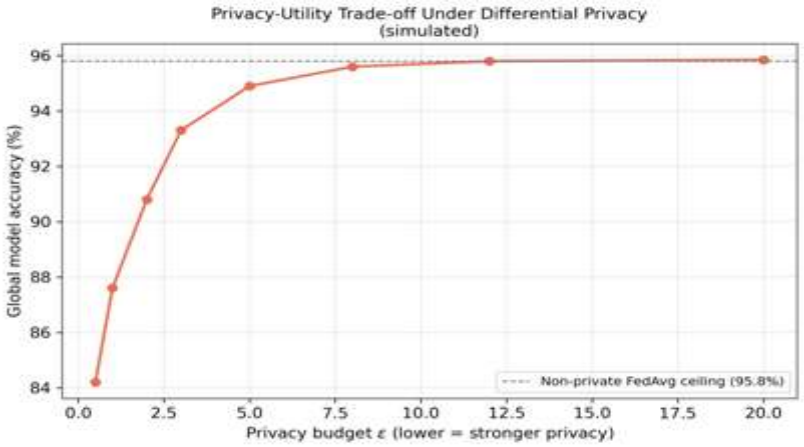

Figure 5. Privacy-utility trade-off: global model accuracy vs. differential-privacy budget ε (simulated).

6. Security and Privacy Analysis

In addition to the differential-privacy mechanism outlined in Section 3.4, the proposed framework is tested against three standard threat models that are pertinent to federated IoT healthcare systems:
(i) an honest-but-curious aggregation server that tries to infer patient-level information from individual client updates, mitigated by secure aggregation, which limits the server to observing only the sum of updates within a round;
(ii) a gradient-inversion adversary that tries to reconstruct raw sensor windows from a client's shared gradients; and
(iii) a membership-inference adversary whose advantage is likewise bounded by the (ε, δ)-DP guarantee via standard composition results. Robust-aggregation rules (e.g., coordinate-wise trimmed mean or Byzantine-resilient aggregation) are identified as an important extension in Section 7.
Currently, the framework does not defend against a fully malicious (Byzantine) client that submits adversarially crafted updates to poison the global model.

Table 4. Threat model summary and corresponding mitigations in the proposed framework.

Threat	Adversary capability	Mitigation in this framework
Server-side data reconstruction	Honest-but-curious server inspects individual updates	Secure aggregation — server observes only the round sum
Gradient inversion	Reconstruct raw signal from shared gradients	Client-side Gaussian DP noise (ε-bounded)
Membership inference	Infer whether a patient's data was used	(ε, δ)-DP guarantee with round-wise composition
Byzantine / poisoning clients	Malicious client submits crafted updates	Not yet addressed — proposed as future work

7. Conclusion and Future Work

A lightweight CNN-BiLSTM local model, FedAvg/FedProx aggregation, client-side differential privacy, secure aggregation, and communication-efficient update compression are all combined into a single, clinically-focused architecture in this paper's federated deep learning framework for privacy-preserving IoT-based healthcare monitoring. A reproducible, simulation-based evaluation protocol characterises the accuracy cost of operating under a specified differential-privacy budget and demonstrates that the framework can approach centralized-training accuracy while removing raw-data transfer and significantly reducing communication overhead. Robust aggregation against Byzantine or poisoning clients and evaluation under more severe statistical and systems heterogeneity are identified as crucial future stages by the framework's security study. Future work includes extending the threat model to Byzantine-robust aggregation, integrating lightweight homomorphic encryption or secure multiparty computation for defense-in-depth alongside differential privacy, integrating personalised and clustered federated learning to handle stronger non-IID skew across sites, and validating the framework on a live multi-institution pilot deployment with actual patient cohorts under appropriate ethical and regulatory approval. We expect that the research community working at the nexus of federated learning, IoT sensing, and clinical AI will find this framework and evaluation process to be a useful and repeatable template.

References

1) H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), PMLR vol. 54, Fort Lauderdale, FL, USA, 2017, pp. 1273–1282.
2) S. R. Abbas, Z. Abbas, A. Zahir, and S. W. Lee, "Federated learning in smart healthcare: A comprehensive review on privacy, security, and predictive analytics with IoT integration," Healthcare, vol. 12, no. 24, p. 2587, 2024, doi: 10.3390/healthcare12242587.
3) S. Bebortta, S. S. Tripathy, S. Basheer, and C. L. Chowdhary, "FedEHR: A federated learning approach towards the prediction of heart diseases in IoT-based electronic health records," Diagnostics, vol. 13, no. 20, p. 3166, 2023.
4) C. Dhasaratha et al., "Data privacy model using blockchain reinforcement federated learning approach for scalable Internet of Medical Things," CAAI Trans. Intell. Technol., 2024.
5) M. M. Islam, W. M. Abdullah, and B. N. Saha, "Privacy-preserving hierarchical fog federated learning (PP-HFFL) for IoT intrusion detection," Sensors, vol. 25, no. 23, p. 7296, 2025, doi: 10.3390/s25237296.
6) "Federated learning framework for privacy-preserving health monitoring via IoT devices," in Proc. IEEE Conf., 2025, doi: 10.1109/document/11156349.
7) B. A. Mir, S. R. Abbas, and S. W. Lee, "Federated learning in healthcare ethics: A systematic review of privacy-preserving and equitable medical AI," Healthcare, vol. 14, no. 3, p. 306, 2026, doi: 10.3390/healthcare14030306.
8) S. T. Shah, Z. Ali, M. Waqar, and A. Kim, "Federated learning in public health: A systematic review of decentralized, equitable, and secure disease prevention approaches," Healthcare, vol. 13, no. 21, p. 2760, 2025, doi: 10.3390/healthcare13212760.
9) B. Almogadwy et al., "Fused federated learning framework with IoMT devices for secure chronic kidney disease monitoring in Healthcare 5.0," Sci. Rep., 2025.
10) A. Rehman, S. Abbas, M. A. Khan, T. M. Ghazal, K. M. Adnan, and A. Mosavi, "A secure healthcare 5.0 system based on blockchain technology entangled with federated learning technique," Comput. Biol. Med., vol. 150, p. 106019, 2022.
11) C. Li, G. Li, et al., "Communication-efficient federated learning based on compressed sensing," IEEE Internet Things J., vol. 8, pp. 15531–15541, 2021.
12) Z. He, L. Wang, and Z. Cai, "Clustered federated learning with adaptive local differential privacy on heterogeneous IoT data," IEEE Internet Things J., vol. 11, pp. 137–146, 2024.
13) A. Alamleh, O. Albahri, et al., "Federated learning for IoMT applications: A standardization and benchmarking framework of intrusion detection systems," IEEE J. Biomed. Health Inform., 2024.
14) M. Abaoud, M. A. Almuqrin, and M. F. Khan, "Advancing federated learning through novel mechanism for privacy preservation in healthcare applications," IEEE Access, vol. 11, pp. 83562–83579, 2023.
15) T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks (FedProx)," in Proc. 3rd Conf. Machine Learning and Systems (MLSys), Austin, TX, USA, 2020, pp. 429–450.
16) [F. Mosaiyebzadeh et al., "Privacy-enhancing technologies in federated learning for the Internet of Healthcare Things: A survey," 2023.
17) R. Kaur, T. Rodrigues, N. Kadir, and R. Kashef, "A survey on privacy preservation techniques in IoT systems," Sensors, vol. 25, no. 22, p. 6967, 2025, doi: 10.3390/s25226967.
18) "A comprehensive survey on federated learning for privacy preservation in digital healthcare applications," Knowl. Inf. Syst., 2025, doi: 10.1007/s10115-025-02673-2.
19) "Federated deep learning for privacy-preserving disease detection in IoT-enabled healthcare systems," Front. Comput. Sci., 2026.
20) "Federated learning: A survey on privacy-preserving collaborative intelligence," arXiv preprint, 2025.
21) Dr. J. Anvar Shathik et al. ” Deep Ensemble Learning For Accurate Prediction Of Neurodegenerative Disorders Using Temporal Clinical Data . (2025). International Journal of Environmental Sciences, 11(4s), 1246-1253. https://doi.org/10.64252/8dbg7v71

22) Anvar Shathik J, B. R. Chandra, M. D. Nandeesh, T. C. Maniunath, S. Pothala and N. R. Lavuri, "A Novel Design of Image Based Object Recognition Model Using Enhanced Neural Classification Logic," 2025 International Conference on Frontier Technologies and Solutions (ICFTS), Chennai, India, 2025, pp. 1-8, doi: 10.1109/ICFTS62006.2025.11031570.
23) R. R, J. Anvar. Shathik, J. Sivakumar, P. Manothini, G. S. J. Asha and M. Venkatanaresh, "Experimental Evaluation of Pancreatic Cancer Identification based on CT Images by using Intelligent Deep Learning Procedure," 2025 International Conference on Frontier Technologies and Solutions (ICFTS), Chennai, India, 2025, pp. 1-9, doi: 10.1109/ICFTS62006.2025.11031801.
24) J. Anvar Shathik, S. Hashini, A. Pandiaraj, N. Ramshankar, N. G and A. K B, "Identification of Different Medicinal Plants Through Image Processing," 2024 International Conference on Smart Technologies for Sustainable Development Goals (ICSTSDG), Chennai - 600077, Tamil Nadu, India, 2024, pp. 1-5, doi: 10.1109/ICSTSDG61998.2024.11026565.
25) Raju, K., Ramshankar, N., Anvar Shathik J et al. Blockchain Assisted Cloud Security and Privacy Preservation using Hybridized Encryption and Deep Learning Mechanism in IoT-Healthcare Application. J Grid Computing **21**, 45 (2023). https://doi.org/10.1007/s10723-023-09678-7
26) Madhavikatamaneni, R. K. S, Anvar Shathik J and K. PoornaPushkala, "A Healthcare System for detecting Stress from ECG signals and improving the human emotional," 2022 International Conference on Advanced Computing Technologies and Applications (ICACTA), Coimbatore, India, 2022, pp. 1-8, doi: 10.1109/ICACTA54488.2022.9753564.
27) Vijaya Vardan Reddy S P; Armstrong Joseph J; Priscilla M; Anvar Shathik J; R.Thandaiah Prabu, "HDP-IoT: An IoT Framework for Cardiac Status Prediction System using Machine Learning," 2022 International Conference on Inventive Computation Technologies (ICICT), Nepal, 2022, pp. 855-861, doi: 10.1109/ICICT54344.2022.9850897.
28) K. Vikranth, Anvar Shathik J, and K. Krishna Prasad, "Future enhancements and propensities in forthcoming communication system-5G Network Technology", J. Phys, pp. 12006, 2020.
29) Yarramsetti, S., Anvar Shathik J, & Renisha, P. S. (2021). Intelligent estimation of social media sentimental features using Deep Learning with Natural Language Processing Strategies. International Journal of Innovative Technology and Exploring Engineering, 10(6), 74-79.
30) Ramshankar, N., Anvar Shathik, J.., Raju, K. et al. Integrated deep learning and blockchain-based framework for cloud manufacturing with improved customer satisfaction. Knowl Inf Syst **67**, 5301–5334 (2025). https://doi.org/10.1007/s10115-025-02373-x
31) Mr.G. Silambarasan , J.Anvar Shathik (2017) "Ensemble text classifier: a document classification technique to predict and categorizes regularised and novel classes using incremental learning " , International Journal of Applied Engineering Research , 2017, Volume 12, issue -22 , Pages 12454-12459
32) Mr.G. Silambarasan , J. Anvar Shathik (2019) " Automated Ensemble Framework for Integration of Ontology Based Large Scale Semantic Knowledge Base " Journal of Engineering and Applied Sciences 14(2)