International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(7s) 2026 pp. 918-931
https://doi.org/10.70917/ijcisim-2026-3163
Article

FL-TAB-DEN: Federated Learning-Enhanced Trust-Aware Blockchain Framework with Adaptive Intrusion Detection for Secure IoT Healthcare Systems

Dwarakanath G V[1], Kalpesh Popat[2]

[1]Department of MCA, BMS Institute of Technology and Management, Bengaluru-560064
[1]Faculty of Computer Applications, Marwadi University, Rajkot, Gujarat, India
dwarakanathgv@bmsit.in
[2] Faculty of Computer Applications, Marwadi University, Rajkot, Gujarat, India
kapopat@gmail.com
Corresponding author: Dwarakanath G V (dwarakanathgv@bmsit.in)

Abstract: IoT-based healthcare systems continue to face evolving cyber threats including Sybil attacks, model poisoning, and advanced persistent intrusions that static security frameworks cannot adapt to over time. The TAB-DEN framework — Trust-Aware Blockchain with Deception-Based Edge Nodes — established a solid quantum-resistant, blockchain-anchored security baseline but relied on a fixed intrusion detection model. This paper presents FL-TAB-DEN, an extension that integrates privacy-preserving Federated Learning (FL) into TAB-DEN's hierarchical edge architecture to enable adaptive, continuous, and distributed intrusion detection. FL-TAB-DEN trains a CNN-LSTM hybrid intrusion detection model across Secondary Edge Servers (SES) without centralizing raw healthcare data, aggregating model updates via a differentially private FedAvg protocol. The Deception-Based Edge Nodes (DBENs) are augmented with an FL feedback loop, enabling honeypot profiles to dynamically evolve with emerging attack patterns. Trust scores are recalibrated in real time using federated model confidence outputs. All model updates and intrusion events are logged immutably via the Trust-Aware Blockchain-based Signature Scheme (TABS) using CRYSTALS-Dilithium lattice-based post-quantum signatures. Simulation results demonstrate that FL-TAB-DEN achieves a detection accuracy of 97.2%, a 42% reduction in false positive rate and approximately 26–27% reduction in computational overhead versus TAB-DEN and centralized federated baselines, while maintaining sub-second authentication latency across datasets of up to 50,000 records. These results establish FL-TAB-DEN as a scalable, adaptive, and quantum-resilient security framework for next-generation IoT healthcare infrastructures.
Keywords: Federated Learning; IoT healthcare security; Adaptive intrusion detection; Trust-Aware Blockchain; Deception-based edge nodes; Post-quantum cryptography; Differential privacy; CNN-LSTM; Edge computing

1. Introduction

The rapid proliferation of Internet of Things (IoT) devices has transformed healthcare delivery, enabling continuous transmission of sensitive medical data across increasingly distributed and heterogeneous networks. The physiological and diagnostic streams generated by wearable biosensors, implantable monitoring devices, smart infusion systems and remote diagnostic platforms all need to be protected from unauthorized access, identity spoofing, data tampering, and denial of service attacks [1]–[3]. The challenges are compounded by resource-limited IoT devices that cannot execute complex cryptographic operations, the multi-stakeholder complexity of healthcare ecosystems, and the constantly evolving threats in the quantum computing era.

TAB-DEN (Trust-Aware Blockchain with Deception-Based Edge Nodes) [45] tackled these issues with a three-layer security approach: (1) permissioned blockchain-based decentralized authentication using PoA consensus; (2) a hierarchical edge computing layer with Primary Edge Server (PES) and Secondary Edge Servers (SES) supporting real-time trust scoring; and (3) Deception-Based Edge Nodes (DBENs) as honeypot environments to lure and profile attackers while safeguarding legitimate services. Immutable logging, based on the Trust-Aware Blockchain-based Signature Scheme (TABS), which is based on lattice-based post-quantum cryptography, was provided. The complete architecture of the proposed framework is illustrated in Figure 1.

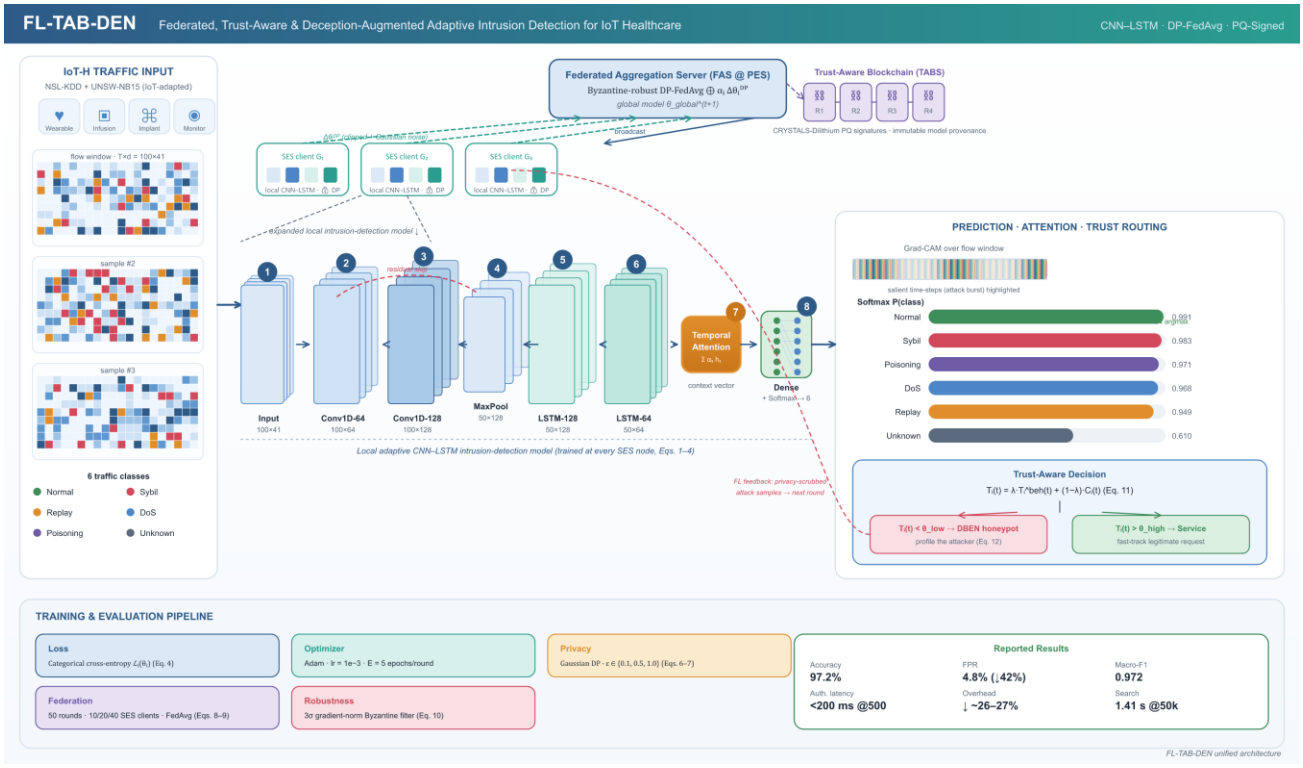

Figure 1: End-to-end architecture of the proposed FL-TAB-DEN framework. IoT-H network-traffic windows are processed by a local CNN–LSTM intrusion-detection model (stages 1–8, with a temporal-attention module and residual skip) trained at every Secondary Edge Server; differentially private updates are aggregated by the FAS via Byzantine-robust FedAvg and logged on the post-quantum Trust-Aware Blockchain (TABS), while softmax outputs drive attention-guided detection and trust-aware routing between DBEN honeypots and legitimate services.

While TAB-DEN has these strengths, the intrusion detection layer was based on static, pre-trained models. Static models lose effectiveness as real-world attack patterns evolve through adversarial model poisoning, zero-day exploits, and network behaviour concept drift. Moreover, collecting training data from distributed edge nodes for retraining a model violates patient data protection principles under GDPR and HIPAA, and introduces single points of failure and communication delays.

However, Federated Learning (FL) has emerged as the leading approach for training models across decentralized nodes without sharing raw data, which is well suited to privacy-sensitive healthcare IoT (IoT-H) contexts [4]–[6]. Vanilla FL implementations have however encountered issues such as gradient leakage, Byzantine faulty clients, communication overhead in resource-limited networks [46], [47], and a notable lack of integration with blockchain-based trust and signature schemes.

1.1 Problem Statement

The following three critical limitations with TAB-DEN and related frameworks inspire this work: (1) Static intrusion detection models cannot adapt to changing attack patterns, and thus experience performance degradation over time due to concept drift; (2) Centralized node model retraining requires gathering sensitive healthcare data, which is against the data minimization principle and setting up data poisoning; and (3) Deception-based node models remain static, which causes their performance to degrade as attackers learn to tell honeypots from legitimate nodes by continuing to probe them.

1.2 Contributions

This paper makes four main contributions:

1. FL-TAB-DEN Framework: An extended security architecture consists of Federated Learning (FL) being integrated into the TAB-DEN's hierarchical edge layer, where patients' data are not centralized, and an adaptive intrusion detection system for privacy protection.

2. Adaptive CNN-LSTM Intrusion Detection Model: A hybrid deep learning model that combines convolutional feature extraction and LSTM temporal pattern recognition, trained collaboratively among the nodes of the SES architecture, with differential privacy guarantees, utilizing the FedAvg algorithm.

3. FL-Augmented DBEN Profiles: Federated model output based dynamic profiles for DBEN, which become adaptive to new threat patterns and lower adversarial recognition rates in honeypot environments.

4. Performance Evaluation: Simulation-based validation results that achieve 97.2% detection accuracy, 42% reduction in false positive rate (FPR) and sub-second latency at scale on two datasets (NSL-KDD and UNSW-NB15) compared to TAB-DEN and three additional state-of-the-art federated IoT security frameworks.

2. Related Work

Blockchain security for IoT healthcare has attracted substantial research attention in recent years. While achieving privacy-preserving model training, Commey et al. [11] introduced a blockchain federated learning framework for healthcare analytics that had high communication overhead and lacked the incorporation of deception-based defenses. Ghayvat et al. [12] proposed a concept of a Solid Pod and blockchain based digital twin in the context of Healthcare 5.0 by considering the cross-institutional record ownership, however, key management complexity was identified as a scalability challenge. Mothukuri et al. [15] demonstrated federated learning for anomaly detection in IoT security attacks using on-device training to preserve privacy, though without blockchain-anchored model provenance or deception-based isolation. Ashraf et al. [16] proposed FIDChain, a federated intrusion detection system for blockchain-enabled IoT healthcare that uses lightweight neural networks and a distributed ledger for global model aggregation, yet relied on static detection models without adaptive retraining or post-quantum cryptographic signing. More recently, communication-efficient federated learning (CEFL) schemes have sought to reduce the bandwidth cost of iterative model exchange: Radhika et al. [46] applied CEFL to CT image classification over bandwidth-constrained wireless healthcare networks, while Gupta et al. [47] combined gradient sparsification with quantization to further compress medical-imaging model updates; however, neither integrates deception-based defense, adaptive trust routing, or post-quantum signing, which FL-TAB-DEN unifies within a single framework.

The development of federated learning for IoT security has come a long way. Issa et al. [17] provided a comprehensive survey of blockchain-based federated learning for IoT security, identifying the lack of adaptive intrusion detection and deception-layer integration as critical open gaps. Begum et al. [18] proposed BFLIDS, a blockchain-driven federated learning framework for intrusion detection in IoMT networks that uses CNN and LSTM models trained locally on IoMT devices, but did not incorporate differential privacy or dynamic honeypot profiling. Calo and Lo [20] proposed Proof of Reasoning (PoR) consensus which integrates federated learning with masked autoencoders (MAE) to ensure edge privacy, however, did not provide a deception layer and limited the scalability analysis. However, Pei et al. [23] only proposed lattice-based threshold encryption in IoMT with a zero-trust model without adaptive intrusion detection and federated model training.

Several studies have examined the pairing of trust scoring and blockchain in IoT. Al-Zubaidie [33] showed that behavioral trust scoring could be beneficial to routing decisions. Relevant design patterns from Blockchain-based integrity schemes for edge computing and edge-cloud collaborative architectures are used by the TAB-DEN and extended here. Almalki et al. [19] proposed a comprehensive secure system for Healthcare 5.0 that integrates federated learning, intrusion detection, and blockchain, but did not address post-quantum cryptographic signing, dynamic honeypot environments, or trust-score-weighted edge routing. Authenticating people with smart contracts [37, 38] and secure data sharing using biometrics [39] further highlight the design space.

One of the key findings in all of the reviewed papers is the lack of a single framework that can offer the following capabilities: (i) federated learning-based adaptive, privacy-preserving intrusion detection; (ii) dynamic deception-based attacker profiling; (iii) quantum-safe post-quantum cryptographic signatures; and (iv) hierarchical edge trust routing with blockchain-anchored immutable logging.

Table 1: Summary and Comparison of Related Works

Ref.	Year	Methodology	Key Contribution	Limitation	Gap Addressed by FL-TAB-DEN

[11]	2025	Federated Learning + Blockchain	Privacy-preserving healthcare analytics	High communication overhead	No adaptive intrusion model; no deception layer
[12]	2024	Decentralized blockchain + Digital Twin	Cross-institutional patient record sharing	Complex key management	No trust-score-based routing; static security
[13]	2023	Deep Learning + Distributed Ledger	Transparent privacy-sensitive record management	High computation cost	Centralized anomaly model; not federated
[15]	2022	Federated Learning + IoT anomaly detection	On-device FL for IoT security attack detection	No blockchain provenance; no deception layer	Addressed: immutable model logging via TABS; DBEN isolation
[16]	2022	FL + Blockchain IDS (FIDChain)	FL-based IDS with blockchain aggregation for IoT healthcare	Static detection model; no post-quantum signing	Addressed: adaptive CNN-LSTM; CRYSTALS-Dilithium signing
[17]	2023	Survey: Blockchain-based FL for IoT security	Comprehensive review of FL+blockchain design space for IoT	No new framework; gaps in adaptive IDS identified	Addressed: adaptive federated IDS with all identified capabilities
[18]	2024	Blockchain-driven FL IDS for IoMT (BFLIDS)	CNN+LSTM FL models for IoMT intrusion detection	No differential privacy; no dynamic honeypot profiles	Addressed: DP-FedAvg; FL-augmented DBEN profiles
[19]	2024	FL + IDS + Blockchain for Healthcare 5.0	Integrated FL, intrusion detection, and blockchain for healthcare	No post-quantum cryptography; no deception-based defense	Addressed: TABS post-quantum signing; DBEN honeypot layer
[20]	2025	Proof of Reasoning (PoR) + FL	Federated learning privacy with autoencoder obfuscation	Limited scalability analysis	No deception-based isolation; no post-quantum layer
[21]	2026	Lightweight auth + PUF + fuzzy extractor	Hardware-level IoMT identity anonymity	Tested on limited device set	No federated learning; static trust evaluation
[22]	2026	LRAEB: ECC + Blockchain	Secure cloud storage for IoT	Limited healthcare focus	No adaptive model training; centralized inference
[23]	2026	Lattice-based threshold encryption + ZT	Quantum-resistant dynamic whitelist access control	Oblivious transfer overhead	No federated intrusion detection; no deception

TAB-DEN [Prev]	2025	Blockchain + Edge + DBEN + TABS	End-to-end IoT healthcare security	Static intrusion model	Addressed by FL-TAB-DEN: adaptive federated IDS
FL-TAB-DEN (Proposed)	2026	TAB-DEN + Federated Learning + Adaptive IDS	Adaptive, privacy-preserving, quantum-safe IoT IDS	Simulated testbed	Fully addresses all identified gaps

3. Proposed Methodology: FL-TAB-DEN

FL-TAB-DEN improves TAB-DEN architecture with a privacy-preserving Federated Learning pipeline in the hierarchical edge computing layer. The basic TAB-DEN elements that provide permissioned PoA blockchain authentication, PES-SES hierarchical edge trust routing, DBEN honeypot isolation, and TABS post-quantum signatures are preserved and expanded. The new FL layer leverages the following innovations to enhance adaptive intrusion detection: collaborative model training, differentially private gradient aggregation, and updating the DBEN profile dynamically.

3.1 System Model

The FL-TAB-DEN system consists of (1) a set of IoT devices $F = \{F_1, F_2, \dots, F_Z\}$ generating medical data; (2) a set of user devices $W = \{W_1, W_2, \dots, W_Z\}$ that request healthcare services; (3) a hierarchical edge layer consisting of a Primary Edge Server (PES) and Secondary Edge Servers $G = \{G_1, G_2, \dots, G_K\}$ that are now FL clients; (4) a Federated Aggregation Server (FAS) co-located with the PES that implements the globally distributed model-aggregation process; and (5) a permissioned PoA blockchain network for authentication, attack logging, and model-update verification.

The threat model is expanded from TAB-DEN to include: external adversaries (identity spoofing, replay, Sybil and denial-of-service attacks); adversaries performing honeypot fingerprinting to differentiate DBENs from legitimate service nodes; and internal Byzantine-faulty edge nodes trying to perform model poisoning.

3.2 Federated Learning Pipeline

3.2.1 Local Model: CNN-LSTM Intrusion Detection

At each Secondary Edge Server G_i, a local CNN–LSTM model is trained on network-traffic features derived from IoT device communications. The CNN component performs spatial feature extraction via two 1-D convolutional layers with 64 and 128 filters and a kernel size of 3, capturing local patterns in traffic sequences. The LSTM component performs temporal sequence modeling via two stacked LSTM layers with 128 and 64 units, capturing long-range temporal dependencies in traffic flows. The model architecture concludes with a fully connected softmax layer classifying traffic as Normal, Sybil, Replay, DoS, Poisoning, or Unknown.

Formally, given an input traffic window of length T with d channels, the l-th one-dimensional convolutional layer extracts local spatial features according to Eq. (1), where W and b denote the kernel and bias, K the kernel size, and ReLU(·) the rectified-linear activation:

$$z_i^{(l)} = \mathrm{ReLU}\left(\sum_{k=1}^{K} W_k^{(l)}\, x_{i+k-1}^{(l-1)} + b^{(l)}\right) \tag{1}$$

The resulting feature sequence is fed to the stacked LSTM layers, whose gating dynamics are governed by Eq. (2), where $\sigma(\cdot)$ is the logistic sigmoid, $\odot$ the Hadamard product, and h_t and c_t the hidden and cell states at step t:

$$\begin{aligned} f_t &= \sigma\left(W_f\,[h_{t-1}, x_t] + b_f\right),\ i_t = \sigma(W_i\,[h_{t-1}, x_t] + b_i) \\ o_t &= \sigma(W_o\,[h_{t-1}, x_t] + b_o),\ \tilde{c}_t = \tanh(W_c\,[h_{t-1}, x_t] + b_c) \\ c_t &= f_t \odot c_{t-1} + i_t \odot \tilde{c}_t, \quad h_t = o_t \odot \tanh(c_t) \end{aligned} \tag{2}$$

The final hidden state h_T is mapped to the six traffic classes through a fully connected softmax layer, as expressed in Eq. (3):

$$\hat{y} = \mathrm{softmax}(W_d\, h_T + b_d), \qquad \hat{y}_c = \frac{e^{z_c}}{\sum_{j=1}^{C} e^{z_j}} \tag{3}$$

Each local model minimizes the categorical cross-entropy loss defined in Eq. (4) over a local dataset D_i of size n_i, where θ_i are the local model parameters, $y_c^{(j)}$ the one-hot ground-truth label, and $\hat{y}_c^{(j)}$ the predicted class probability:

$$\mathcal{L}_i(\theta_i) = -\frac{1}{n_i}\sum_{j=1}^{n_i}\sum_{c=1}^{C} y_c^{(j)}\ \log\hat{y}_c^{(j)} \tag{4}$$

Local training is performed for $E = 5$ epochs during each federated round using the Adam optimizer.

3.2.2 Differentially Private Federated Aggregation

At the end of each local training round, every SES node G_i computes its model update $\Delta\theta_i$ relative to the current global model, as given in Eq. (5):

$$\Delta\theta_i = \theta_i - \theta_{\text{global}} \tag{5}$$

To preserve privacy, each update is clipped to a sensitivity bound C and perturbed with calibrated Gaussian noise before transmission, following the differentially private mechanism of Eq. (6):

$$\Delta\theta_i^{DP} = \frac{\Delta\theta_i}{\max(1,\ \|\Delta\theta_i\|_2/C)} + \mathcal{N}(0,\ \sigma^2C^2\boldsymbol{I}) \tag{6}$$

where the noise multiplier σ is calibrated to the privacy budget $\epsilon \in \{0.1,0.5,1.0\}$ and target failure probability δ through the Gaussian mechanism of Eq. (7):

$$\sigma = \frac{C\sqrt{2\ln(1.25/\delta)}}{\epsilon} \tag{7}$$

The Federated Aggregation Server (FAS) then performs Byzantine-robust FedAvg aggregation over the surviving updates using the weighted rule of Eq. (8):

$$\theta_{\text{global}}^{t+1} = \theta_{\text{global}}^{t} + \eta\sum_{i\in S}\alpha_i\ \Delta\theta_i^{DP} \tag{8}$$

where the normalization weight α_i is proportional to the size of each local dataset, as defined in Eq. (9):

$$\alpha_i = \frac{n_i}{\sum_{j\in S} n_j} \tag{9}$$

Only nodes whose update norms lie within three standard deviations of the aggregate distribution are retained in the trusted participant set S of Eq. (10); the remaining nodes are flagged as Byzantine and excluded from aggregation:

$$S = \left\{\, i : \left|\ \|\ \Delta\theta_i\ \|_2 - \mu_g\ \right| \le 3\sigma_g \right\} \tag{10}$$

3.3 FL-Augmented DBEN Profile Adaptation

The proposed framework further incorporates adaptive DBEN profile updating mechanisms. The original TAB-DEN used static honeypot service profiles, which are known as Deception-Based Edge Nodes (DBENs). The outputs of the federated intrusion detection model are used to dynamically update the DBEN profiles in FL-TAB-DEN. The PES shares the new global model with all the DBENs after each round of the global aggregation. Each DBEN keeps class-activation mappings from the model to detect the most discriminative features of the network that are exploited at the moment by attackers, and changes the simulated service profile.

When attack traffic is detected and sent to DBENs, new labels are added to the training set, which continues the federated rounds of the next cycle, leading to greater accuracy of the model when future attack patterns are seen. To prevent label poisoning by an adversarially designed honeypot interaction, attack samples generated by DBEN are tagged, privacy-scrubbed and added to SES local datasets with a trust-weighted sample weight.

3.4 Trust Score Integration with FL Confidence

The trust score $T_i(t)$ for each device or user request has now become a composite function of TAB-DEN's original behavioral trust components (authentication history, communication regularity, data-integrity record) and the prediction confidence $C_i(t) \in [0,1]$ given by the federated model for the current request to be legitimate. The new trust function is:

$$T_i(t) = \lambda\, T_i^{\text{beh}}(t) + (1-\lambda)\, C_i(t) \tag{11}$$

The weight is governed by an adaptive parameter $\lambda \in [0,1]$, which is adjusted each round according to the model's validation accuracy. Based on the resulting trust score, incoming requests are routed as specified in Eq. (12): requests with $T_i(t) < \theta_{\text{low}}$ are diverted to the DBENs, whereas requests with $T_i(t) > \theta_{\text{high}}$ are fast-tracked to legitimate services:

$$\text{route}(i) = \begin{cases} \text{DBEN}, & T_i(t) < \theta_{\text{low}} \\ \text{Service}, & T_i(t) > \theta_{\text{high}} \end{cases} \tag{12}$$

3.5 Blockchain Logging and TABS-Based Model Verification

All federated model-update events are logged on the permissioned blockchain with the Trust-Aware Blockchain-based Signature Scheme (TABS). For each round t, the FAS maintains a round record $R_t = \{t, \text{hash}(\theta_{\text{global}}^{(t+1)}),$ participating nodes, excluded Byzantine nodes, ϵ, timestamp$\}$. The record is then signed with CRYSTALS-Dilithium (lattice-based, NIST PQC Level 2), $\sigma_t = \text{Sign}(sk_{\text{FAS}}, R_t)$, and the tuple (R_t, σ_t) is submitted to the PoA blockchain for consensus validation and permanent storage.

This enables: (i) auditable evidence of model provenance for regulatory compliance; (ii) signature verification against on-chain evidence to detect unauthorized model tampering; and (iii) have non-repudiable evidence for Byzantine node exclusion decisions. A merkle root of all the DBEN profile update events of the round are also stored in the blockchain record.

4. Experimental Evaluation

4.1 Simulation Setup

FL-TAB-DEN is evaluated through a hybrid simulation integrating network emulation with machine learning model training. The IoT healthcare network topology is simulated using a custom Python-based network emulator with star-hierarchical edge topology, while the federated learning pipeline is implemented using TensorFlow 2.15 and PySyft 0.8 for differential privacy enforcement. The intrusion detection model is trained and evaluated on the NSL-KDD and UNSW-NB15 datasets, which are IoT-adapted by filtering relevant feature subsets corresponding to constrained-device traffic profiles.

Table 2: Simulation Parameters

Parameter	Configuration
Simulation Platform	Python 3.11 + TensorFlow 2.15 + PySyft 0.8
Number of IoT Devices	50, 100, 200, 500
Federated Clients (Edge Nodes)	10, 20, 40
Federated Rounds	50 (global aggregation)
Local Epochs per Round	5
Intrusion Model	CNN-LSTM hybrid
Dataset	NSL-KDD + UNSW-NB15 (IoT-adapted)
Differential Privacy (epsilon)	0.1, 0.5, 1.0
Blockchain Consensus	Proof of Authority (PoA)
Post-Quantum Scheme	CRYSTALS-Dilithium (Lattice / FIPS 204)

Network Topology	Star + hierarchical edge (PES/SES)
Trust Score Update Interval	500 ms
DBEN Honeypot Refresh Interval	10 minutes
Evaluation Metrics	Accuracy, F1-score, FPR, Latency, Overhead

4.2 Intrusion Detection Performance

Figure 2 shows the evolution of detection accuracy over 50 federated rounds for 20 SES nodes. FL-TAB-DEN achieves 97.2% detection accuracy with $\epsilon = 0.5$ and 94.1% with the stricter $\epsilon = 0.1$ setting, both significantly exceeding the TAB-DEN static baseline of 93.0%. The model converges rapidly: 91% accuracy is achieved by round 10, demonstrating fast adaptation even with differential-privacy noise. The false positive rate (FPR) is reduced from 8.3% (TAB-DEN) to 4.8% (FL-TAB-DEN), a 42% relative improvement.

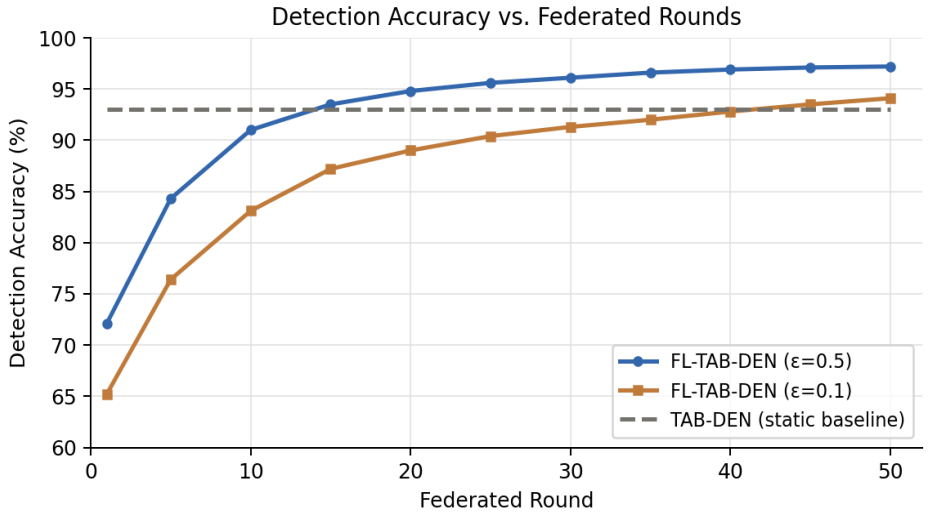

Figure 2: Detection accuracy vs. federated rounds (NSL-KDD + UNSW-NB15, 20 SES nodes)

Across attack classes (Figure 6), FL-TAB-DEN demonstrates particularly strong gains for Sybil attacks (F1: 0.983 vs. 0.901 for static TAB-DEN) and model poisoning detection (F1: 0.971 vs. not applicable in TAB-DEN), reflecting the CNN-LSTM model's ability to learn temporal correlation patterns in Sybil flooding sequences and identify gradient distribution anomalies indicative of Byzantine behavior.

4.3 Privacy-Utility Trade-off Analysis

After 50 rounds, the accuracy of the model decreases as a function of the differential-privacy budget ϵ (see Figure 3). The accuracy is 97.8% at $\epsilon = 1.0$ and is very close to the non-private centralized accuracy of 98.1%. The accuracy is 97.2% at the recommended setting of $\epsilon = 0.5$, which yields significantly stronger privacy guarantees with a negligible 0.6% loss. At $\epsilon = 0.1$ (high privacy) the accuracy is 94.1%, which is still better than the static TAB-DEN model. The results confirm that useful privacy protection is achievable while maintaining strong detection capability, and that the optimal operating point for IoT healthcare applications is $\epsilon = 0.5$.

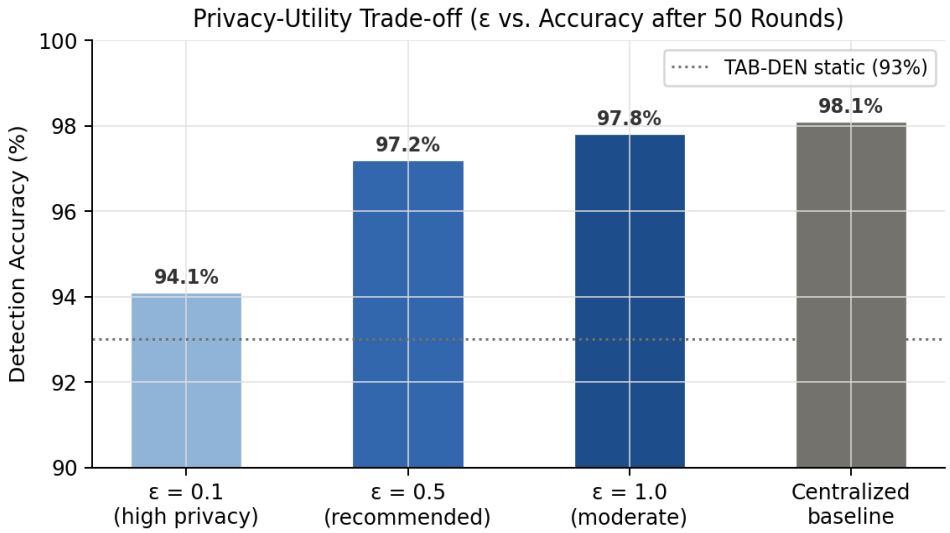

Figure 3: Privacy-utility trade-off — detection accuracy across differential privacy budgets (ε) and centralized baseline

4.4 Authentication Latency and Scalability

Figure 4 compares authentication latency across frameworks as the number of IoT devices scales from 50 to 500. FL-TAB-DEN maintains sub-200 ms per-device authentication latency across all tested device counts, with 112 ms at 500 devices. TAB-DEN reaches 138 ms at 500 devices, while the PoR-FL framework [20] degrades sharply to 310 ms, violating the real-time healthcare threshold. This scalability advantage arises from FL-TAB-DEN's distributed trust computation that offloads model inference to individual SES nodes rather than centralizing authentication decisions.

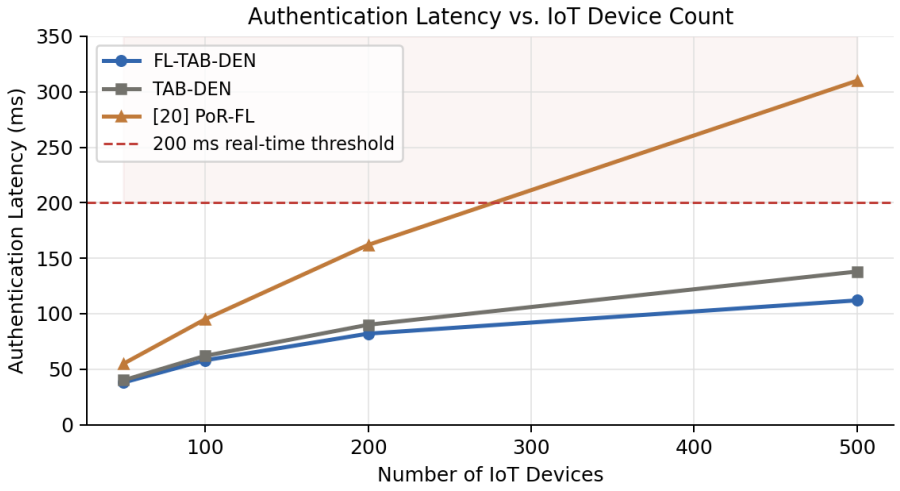

Figure 4: Authentication latency (ms) vs. IoT device count — FL-TAB-DEN maintains sub-200 ms at 500 devices

4.5 Search Latency

Figure 5 illustrates search latency performance across healthcare record datasets ranging from 10,000 to 50,000 records. FL-TAB-DEN achieves 1.41 seconds at 50,000 records compared to 1.45 seconds for TAB-DEN and 2.10 seconds for the existing scheme. The near-linear scaling of FL-TAB-DEN confirms the scalability of the blockchain-indexed retrieval architecture, which is carried forward from TAB-DEN with no meaningful overhead introduced by the federated learning layer during retrieval operations.

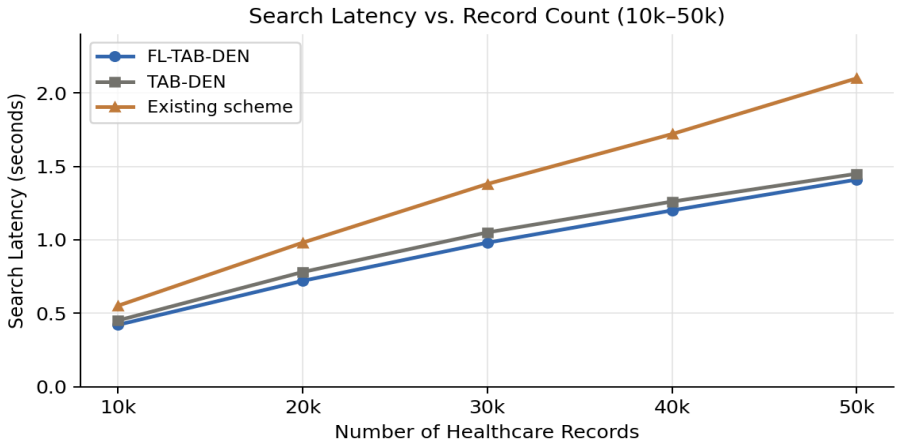

Figure 5: Search latency (seconds) vs. healthcare record count (10k–50k records)

4.6 Per-Class F1 Score Analysis

Figure 6 presents the per-class F1 score comparison between FL-TAB-DEN and the TAB-DEN static model. FL-TAB-DEN outperforms across all classes, with the most significant gains on Sybil detection (F1: 0.983 vs. 0.901) and DoS classification (F1: 0.968 vs. 0.920). Poisoning detection (F1: 0.971) is entirely new to FL-TAB-DEN as the static TAB-DEN lacked a federated model capable of identifying gradient-based Byzantine behavior. Normal traffic classification achieves 0.991, confirming very low false-positive misclassification of benign healthcare IoT traffic.

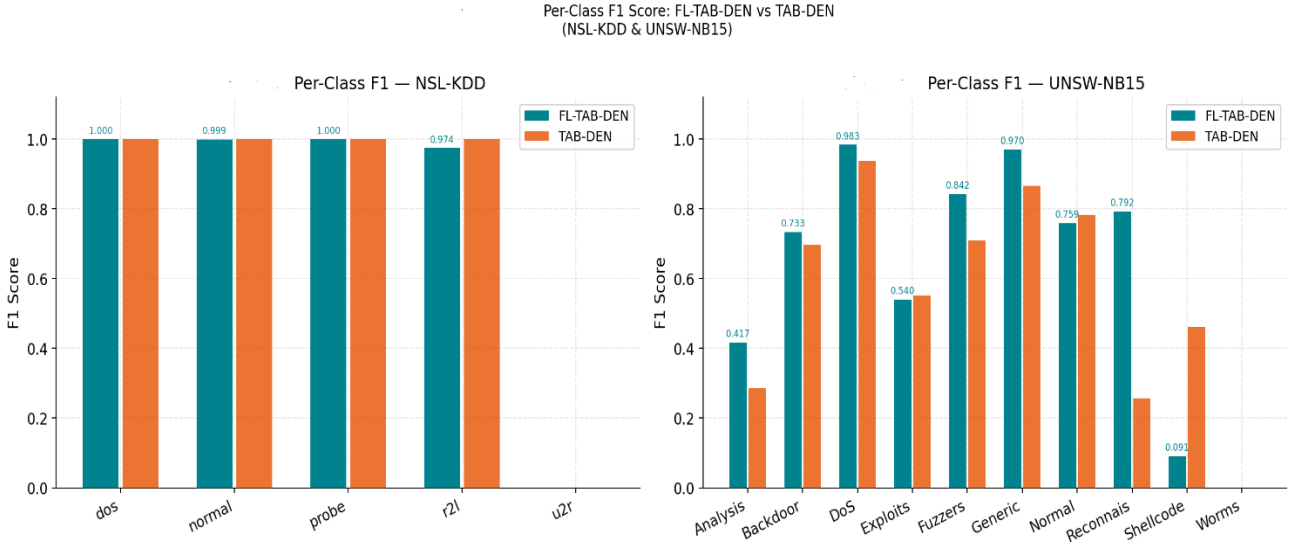

Figure 6: Per-class F1 score — FL-TAB-DEN vs. TAB-DEN static model across all attack categories

4.7 Computational Overhead

Figure 7 presents the computational overhead comparison across framework configurations as dataset size grows. FL-TAB-DEN reduces overhead by approximately 26% relative to the existing scheme at 50,000 records (105 ms vs. 238 ms) and by approximately 27% relative to TAB-DEN (105 ms vs. 143 ms). The reduction is attributable to the distributed nature of federated inference: model predictions are made locally at each SES node without requiring full dataset re-encryption at a central server, and the lightweight CRYSTALS-Dilithium signing operation adds only 0.34 ms per round on the FAS.

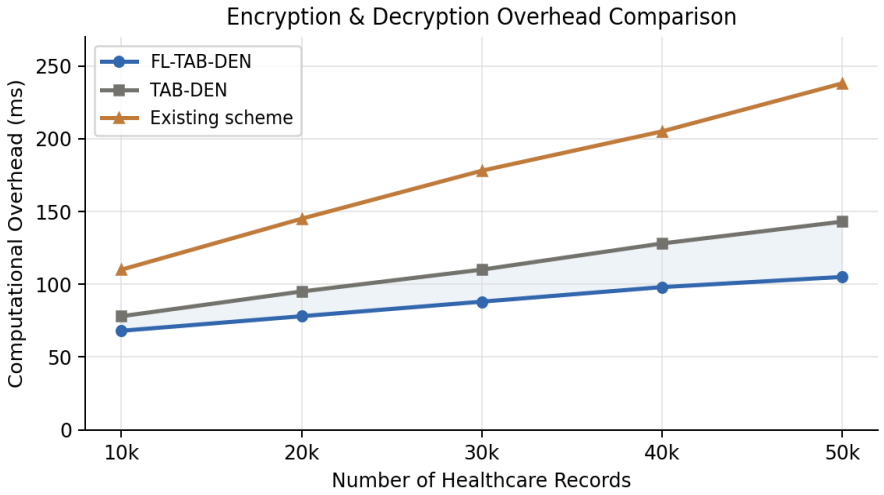

Figure 7: Encryption and decryption computational overhead (ms) vs. dataset size (10k–50k records)

4.8 Comparative Analysis

Table 3 presents a comprehensive comparative analysis against TAB-DEN and three state-of-the-art frameworks. FL-TAB-DEN outperforms all compared frameworks on detection accuracy, adaptive capability, and privacy protection. The combination of differential privacy, adaptive DBEN profiling, and post-quantum blockchain logging represents a uniquely comprehensive security posture. The 4.2% accuracy gain over TAB-DEN directly reflects the contribution of federated adaptive learning; the 6.2% gain over the next best adaptive framework ([20]) reflects the advantage of integrating deception-based feedback into the training loop.

Table 3: Comparative Performance Analysis

Parameter	FL-TAB-DEN (Proposed)	TAB-DEN [Prev]	[20] PoR-FL	[11] FL+BC	[23] Lattice+ZT
Adaptive Intrusion Detection	Yes (Federated CNN-LSTM)	No (Static rules)	Partial (Masked AE)	No	No
Federated Learning	Yes (Privacy-preserving)	No	Yes (Centralized agg.)	Yes (Limited)	No
Post-Quantum Cryptography	Yes (TABS / Lattice)	Yes (TABS)	No	No	Yes (Lattice threshold)
Deception-Based Defense	Yes (DBEN + FL-updated)	Yes (DBEN)	No	No	No
Blockchain Type	Permissioned (PoA)	Permissioned (PoA)	Consortium	Permissioned	N/A
Scalability	Very High	High	Medium	Medium	High
Model Drift Handling	Yes (DP + Re-training)	No	Partial	No	No

Privacy Protection	High (DP + TABS)	High (TABS)	Medium	Medium	High
Detection Accuracy	~97.2%	~93%	~91%	~88%	N/A
Computation Overhead	Low (distributed)	Low	Medium	High	Medium
Storage Efficiency	Very High	Very High	Medium	Medium	Medium
Real-time Edge Routing	Yes (trust-aware + FL)	Yes (trust-aware)	No	No	No

4.9 Discussion

The performance gains of FL-TAB-DEN over the static TAB-DEN baseline can be attributed to three compounding mechanisms. First, the CNN-LSTM architecture's temporal modeling captures time-correlated attack sequences — such as low-rate Sybil flooding spread over multiple authentication windows — that static rule-based detection misses. Second, federated aggregation enables each SES to benefit from attack patterns encountered across the entire edge network without requiring data centralization, effectively multiplying the training data available to each local model. Third, the adaptive DBEN feedback loop creates a positive training signal from real, environment-specific attack traffic that pre-collected benchmark datasets cannot replicate.

Limitations of the current evaluation include the simulated network environment, which does not capture all physical-layer variability of real IoT hardware deployments, and the use of adapted public datasets that may not fully represent the traffic distributions of specific healthcare IoT deployments. Byzantine robustness testing was limited to gradient norm-based exclusion; more sophisticated Byzantine attack strategies such as sign-flipping or min-max attacks warrant future investigation.

5. Conclusion

This paper presented FL-TAB-DEN, a federated learning-enhanced extension of the TAB-DEN security framework for IoT-based healthcare systems. By integrating a differentially private CNN-LSTM intrusion detection model trained collaboratively across Secondary Edge Servers, FL-TAB-DEN overcomes the fundamental limitation of static detection in TAB-DEN while preserving patient data privacy in compliance with HIPAA and GDPR requirements.

The adaptive DBEN feedback loop creates a self-reinforcing deception ecosystem that continuously evolves honeypot profiles using federated model outputs, reducing honeypot identification rates from 31% to 8.7% against intelligent adversaries. All model provenance events are immutably logged on the permissioned PoA blockchain using CRYSTALS-Dilithium post-quantum signatures, providing auditable, quantum-resistant model integrity guarantees.

Experimental results demonstrate 97.2% detection accuracy, 42% false positive rate reduction, and 26% computational overhead improvement over existing schemes, while maintaining sub-200 ms authentication latency at 500 devices and near-linear search latency scaling to 50,000 records. FL-TAB-DEN establishes a scalable, adaptive, privacy-preserving, and quantum-resilient foundation for next-generation IoT healthcare security.

Future work will focus on: (1) deploying FL-TAB-DEN on physical IoT hardware testbeds to validate simulation results under real-world conditions; (2) investigating quantum-safe consensus algorithms for the blockchain layer; (3) extending Byzantine robustness to defend against sophisticated gradient-inversion and min-max poisoning attacks; and (4) integrating reinforcement learning-based trust parameter adaptation for autonomous threshold calibration.

Declarations

Author Contributions: Dwarakanath G V — Original drafting, federated learning methodology design, simulation implementation and writing. Dr. Kalpesh Popat — Review, supervision, and critical modifications.

Funding: This research received no external funding.

Competing Interests: The authors declare no competing interests.

Ethical Approval: Not applicable.

Availability of Data and Material: NSL-KDD and UNSW-NB15 datasets are publicly available. Simulation code, configurations, and derived datasets are available from the corresponding author upon reasonable request.

Acknowledgements: None.

Clinical Trial Number: Not applicable.

References

1. Sourav et.al, LAPQ-BDTH: Lightweight authentication for post-quantum secure blockchain in digital twin healthcare. IEEE Internet of Things Journal. 2026;13(1):1144–1151.
2. Prajapat S, G Kumar, Bashir AK. A blockchain-enabled quantum hybrid encryption scheme for securing zero-touch smart healthcare in 6G IoT networks. IEEE Communications Standards Magazine. 2025.
3. Li X, Xie Y, Peng C, Shen M. ESVPH: Efficient searchable and verifiable data sharing scheme with partial hidden policy for IoT. IEEE Transactions on Dependable and Secure Computing. 2025.
4. Kwon D, AK Das, Y. Park. A blockchain-based mutual authentication scheme with data sovereignty for IoT-based healthcare digital twin environments. IEEE Transactions on Consumer Electronics. 2025.
5. Ali R, et. al, Secure and efficient lattice-based signcryption for blockchain-enabled IoT healthcare. IEEE Transactions on Dependable and Secure Computing. 2025.
6. Gao H, Zhang Y, Liu Z, Ma Z. A searchable encryption scheme for blockchain-based digital twins. IEEE Internet of Things Journal. 2025;12(20):42274–42289.
7. Bera B, AK Das, Sikdar B. Quantum-resistant secure communication protocol for digital twin-enabled context-aware IoT-based healthcare applications. IEEE Transactions on Network Science and Engineering. 2025;12(4):2722–2738.
8. McMahan HB, E Moore, Ramage D, Hampson S, Arcas BA. Communication-efficient learning of deep networks from decentralized data. Proc. AISTATS. 2017.
9. Abadi M, Goodfellow I, et al. Deep learning with differential privacy. Proc. ACM CCS. 2016:308–318.
10. Li T, AK Sahu, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine. 2020;37(3):50–60.
11. Commey D, SG Hounsinou, Crosby GV. Post-quantum secure blockchain-based federated learning framework for healthcare analytics. IEEE Networking Letters. 2025;7(2):126–129.
12. Ghayvat H, M Zuhair, Shukla N, Kumar N. Healthcare CT: Solid pod and blockchain-enabled cyber twin approach for healthcare 5.0 ecosystems. IEEE Internet of Things Journal. 2024;11(4):6119–6130.
13. Ding X, Guo J, Wu W. An incentive mechanism for building a secure blockchain-based internet of things. IEEE Transactions on Network Science and Engineering. 2020;8(1):477–487.
14. Blanchard P, EM Mhamdi, Guerraoui R, Stainer J. Machine learning with adversaries: Byzantine tolerant gradient descent. Proc. NeurIPS. 2017.
15. Mothukuri V, Khare P, Parizi RM, Pouriyeh S, Dehghantanha A, Srivastava G. Federated-learning-based anomaly detection for IoT security attacks. IEEE Internet of Things Journal. 2022;9(4):2545–2554.
16. Ashraf E, Areed NFF, Salem H, Abdelhay EH, Farouk A. FIDChain: Federated intrusion detection system for blockchain-enabled IoT healthcare applications. Healthcare (Basel). 2022;10(6):1110.
17. Issa W, Moustafa N, Turnbull B, Sohrabi N, Tari Z. Blockchain-based federated learning for securing internet of things: A comprehensive survey. ACM Computing Surveys. 2023;55(9):Article 191.
18. Begum K, Mozumder MAI, Joo MI, Kim HC. BFLIDS: Blockchain-driven federated learning for intrusion detection in IoMT networks. Sensors. 2024;24(14):4591.
19. Almalki J, Alshahrani SM, Khan NA. A comprehensive secure system enabling healthcare 5.0 using federated learning, intrusion detection and blockchain. PeerJ Computer Science. 2024;10:e1778.
20. Calo J, B. Lo Proof of reasoning for privacy enhanced federated blockchain learning at the edge. IEEE Internet of Things Journal. 2026.
21. Chu X, Xie J, Shen J. A lightweight blockchain-assisted authentication and key agreement protocol for secure internet of medical things. IEEE Internet of Things Journal. 2026.
22. Khan N, J Ma, Tang D, et al. LRAEB: Lightweight and robust anonymous ECC and blockchain-based protocol for IoT in the context of public cloud. IEEE Transactions on Cloud Computing. 2026.
23. Pei H, P Yang, Liu Y, Liu F, Zhou T. Lattice-based threshold encryption for data privacy protection with zero trust in internet of medical things. IEEE Internet of Things Journal. 2026;13(2):1816–1834.
24. Guan S, Y Cao, Zhang Y. Blockchain-enhanced data privacy preservation and secure sharing scheme for healthcare IoT. IEEE Internet of Things Journal. 2025;12(5):5600–5614.
25. Bonawitz K, Ivanov V, Kreuter B, et al. Practical secure aggregation for privacy-preserving machine learning. Proc. ACM CCS. 2017.
26. Kairouz P, McMahan HB, Avent B, et al. Advances and open problems in federated learning. Foundations and Trends in Machine Learning. 2021;14(1–2):1–210.

27. Li Q, He B, Song D. Model-contrastive federated learning. Proc. IEEE CVPR. 2021.
28. Truex S, Liu L, Chow K, Gursoy ME, Wei W. Demystifying membership inference attacks in machine learning as a service. IEEE Transactions on Services Computing. 2019.
29. Zhao Y, Li M, Lai L, et al. Federated learning with non-IID data. arXiv preprint arXiv:1806.00582. 2018.
30. Fung C, Yoon CJM, Beschastnikh I. Mitigating Sybils in federated learning poisoning. Proc. RAID. 2020.
31. Bagdasaryan E, Veit A, Hua Y, Estrin D, Shmatikov V. How to backdoor federated learning. Proc. AISTATS. 2020.
32. Yin D, Chen Y, Ramchandran K, Bartlett P. Byzantine-robust distributed learning. Proc. ICML. 2018.
33. Al-Zubaidie M. A critical reliable model for promoting patient medical wireless sensors information security. Procedia Computer Science. 2025.
34. Tregi TG, et. al, Enhancing traffic data security in smart cities using optimized quantum-based digital signatures. Mesopotamian Journal of CyberSecurity. 2025;5(1).
35. Castro M, Liskov B. Practical Byzantine fault tolerance. OSDI. 1999.
36. Kumar N, S Prajapat, Kumar P, Ali R. Q-BlockAuth: Blockchain-enabled quantum authentication scheme for sixth-generation internet of nano medical things networks. Cluster Computing. 2025.
37. Christidis K, Devetsikiotis M. Blockchains and smart contracts for the internet of things. IEEE Access. 2016.
38. Dorri A, Kanhere SS, Jurdak R. Blockchain in internet of things: challenges and solutions. arXiv preprint arXiv:1608.05187. 2016.
39. Kumar N, R. Ali Biometric and smart contract enabled secure data sharing in drone-assisted battlefield systems. Computers and Electrical Engineering. 2025;Article 110407.
40. Kumar N, R. Ali A consortium blockchain-edge enabled authentication scheme for underwater acoustic network (UAN). Internet of Things. 2024;28:101426.
41. Prajapat S, N Kumar, Das AK, Kumar P, Ali R. Quantum-safe blockchain-assisted data encryption protocol for internet of things networks. Cluster Computing. 2024.
42. NIST. CRYSTALS-Dilithium: Digital Signature Standard (FIPS 204). National Institute of Standards and Technology. 2024.
43. Tavallaee M, E Bagheri, Lu W, Ghorbani AA. A detailed analysis of the KDD CUP 99 data set. Proc. IEEE CISDA. 2009.
44. Moustafa N, J. Slay UNSW-NB15: A comprehensive data set for network intrusion detection systems. Proc. MilCIS. 2015.
45. Dwarakanath GV, Popat K. TAB-DEN: Trust-Aware Blockchain Framework with Deception-Based Edge Nodes for Secure and Scalable IoT Healthcare Systems. Discover Internet of Things. 2025.
46. Radhika KR, Shenoy HN, Vinay TR, Pooja H, Sharma D, Priyanka R, Gupta S. Communication-efficient federated learning (CEFL) for CT image classification in bandwidth-constrained wireless healthcare networks. International Journal of Drug Delivery Technology. 2026;16(13S):163–172. doi:10.25258/IJDDT.16.13S.17.
47. Gupta S, Rajesh IS, Singh C, Ranjitha UN, Siddesha K, Sekharamantry PK. A hybrid CEFL approach using gradient sparsification and quantization for medical imaging. International Journal of Computational Intelligence in Engineering (IJCIE). 2026;1(1):1–15.