International Journal of Engineering Computational Research and Technology

Journal home page: www.ijecrt.org

ISSN (Online): 2456-9852 Volume 10, December 2025

A Secure Federated Deep Learning Framework for Privacy-Preserving Real-Time Patient Monitoring in IoT Healthcare Systems

TVS Raghavendra [a, *], **T Durga**

[a] *PG Scholar, M.Tech, Department of Computer Science and Engineering, Sree Rama Engineering College, Tirupati, Andhra Pradesh, India*

[b] *Assistant Professor, Department of Computer Science and Engineering, Sree Rama Engineering College, Tirupati, Andhra Pradesh, India*

*Corresponding author

E-mail address: **tvsrvd@gmail.com**

ABSTRACT

The increasing adoption of Internet of Things (IoT) technologies in healthcare has enabled continuous patient monitoring and data-driven clinical decision support; however, it has also raised serious concerns related to data privacy, security, scalability, and computational efficiency. Conventional centralized learning approaches require sensitive medical data to be transferred and stored at central servers, making them vulnerable to data breaches and regulatory violations. The objective of this study is to design and evaluate a secure, privacy-preserving, and computationally efficient IoT healthcare framework that supports real-time patient monitoring without exposing raw medical data. The proposed framework integrates lightweight deep learning models with federated learning to enable decentralized model training across distributed healthcare clients. Physiological IoT sensor data are locally preprocessed and used for model training at the edge, while only encrypted model updates are shared for global aggregation. The system is evaluated using a publicly available Healthcare IoT dataset, with performance assessed under varying data distributions, noise levels, and federated configurations. Experimental results demonstrate that the proposed approach achieves an accuracy of **96.2%**, a precision of **95.1%**, and an F1-score of **94.9%**, outperforming traditional machine learning and centralized deep learning baselines. Inference latency is limited to **18.4 ms per sample**, making the framework suitable for real-time deployment. Robustness analysis shows stable performance under noisy and partially missing sensor data, while energy consumption remains lower than centralized deep learning models. The study concludes that the proposed federated IoT healthcare framework effectively balances intelligence, privacy, and efficiency, offering a practical solution for scalable and secure real-world healthcare monitoring applications.

Keywords: Internet of Things (IoT), Smart Healthcare, Federated Learning, Deep Learning, Patient Monitoring, Privacy Preservation, Edge Computing, Healthcare Analytics.

1. Introduction

The rapid growth of Internet of Things (IoT) technologies has significantly transformed modern healthcare systems by enabling continuous patient monitoring, real-time health data acquisition, and intelligent clinical decision support. The integration of wearable sensors,

embedded medical devices, cloud platforms, and wireless communication networks has paved the way for smart healthcare ecosystems that enhance diagnosis accuracy, treatment efficiency, and overall patient care quality [1]–[3]. IoT-based healthcare systems facilitate the remote monitoring of physiological parameters such as heart rate, blood pressure, glucose levels, oxygen saturation, and body temperature, thereby reducing the burden on healthcare infrastructure and improving access to medical services, particularly in rural and resource-limited regions [4], [5]. With the increasing deployment of IoT devices, healthcare environments are becoming more data-intensive and interconnected. This paradigm shift has introduced new challenges related to data security, patient privacy, computational efficiency, and energy consumption. The highly sensitive nature of medical data necessitates stringent protection mechanisms to prevent unauthorized access, data leakage, and malicious attacks [6], [7]. Moreover, IoT devices often operate under constrained resources, including limited battery life, processing power, and memory capacity, which restricts the deployment of complex analytics and security algorithms at the device level [8], [9]. These challenges underline the critical need for secure, intelligent, and energy-efficient healthcare IoT frameworks capable of delivering reliable and real-time healthcare services.

Recent advancements in deep learning, federated learning, and blockchain technologies have opened new avenues for addressing these challenges. Deep learning enables accurate analysis of complex physiological signals, federated learning ensures decentralized and privacy-preserving model training, and blockchain provides tamper-resistant data storage and access control [10]–[12]. However, integrating these technologies into a unified and scalable healthcare IoT framework remains a complex task. Practical implementation requires careful consideration of communication overhead, computational efficiency, security robustness, and real-time responsiveness. Therefore, there is a strong research motivation to develop a holistic framework that effectively balances intelligence, security, privacy, and efficiency in IoT healthcare environments.

1.1 Background and Motivation

The evolution of healthcare services from traditional hospital-centric models to patient-centric and home-based care systems has been largely driven by advancements in IoT technologies. Wearable devices, smart sensors, and mobile health applications now enable continuous tracking of patient health conditions, facilitating early disease detection and personalized treatment planning [1], [13]. This shift has significantly improved healthcare accessibility, especially for elderly populations, chronically ill patients, and individuals living in remote areas [14]. Despite these advantages, IoT healthcare systems face several technical and operational challenges. One of the primary concerns is the vulnerability of IoT networks to cyberattacks. Healthcare IoT environments are attractive targets for attackers due to the high value of medical data and the critical nature of healthcare operations [6], [15]. Attacks such as data tampering, unauthorized access, denial-of-service, and ransomware can severely compromise patient safety and system reliability. Traditional security mechanisms often fail to provide adequate protection due to their centralized design and limited scalability [16].

Another major challenge is ensuring patient data privacy. Centralized data aggregation models increase the risk of sensitive data exposure during transmission and storage. Regulations such as HIPAA and GDPR impose strict requirements on data confidentiality, necessitating secure and privacy-aware data management strategies [7], [12]. Federated learning has emerged as a promising solution by enabling distributed model training without direct data sharing. However, federated learning introduces additional challenges, including communication overhead, synchronization delays, and vulnerability to poisoning attacks [10], [11]. Energy efficiency also plays a crucial role in IoT healthcare deployment. Wearable and implantable devices typically operate on limited battery resources, making energy-aware computation and communication strategies essential for long-term operation [8], [17]. High computational workloads, frequent data transmission, and continuous monitoring can significantly drain device energy, leading to reduced operational lifetime and compromised service reliability.

Furthermore, healthcare IoT systems generate massive volumes of heterogeneous data streams, including physiological signals, medical images, clinical records, and contextual information. Efficient processing, storage, and analysis of such high-dimensional data require intelligent analytics frameworks capable of extracting meaningful insights in real time [2], [9]. Deep learning-based models have demonstrated remarkable performance in medical signal processing, disease diagnosis, and risk prediction tasks. However, their deployment in resource-constrained IoT environments remains challenging due to their high computational complexity. These factors collectively motivate the need for an integrated IoT healthcare framework that ensures secure data handling, privacy preservation, intelligent analytics, and energy efficiency. By combining deep learning, federated intelligence, and blockchain-based security mechanisms, it is possible to develop a robust healthcare infrastructure that meets the evolving demands of modern healthcare systems [3], [10], [18].

1.2 Limitations of Existing IoT Healthcare Systems

Although numerous IoT healthcare frameworks have been proposed in recent years, several limitations persist in current solutions. Many existing systems rely on centralized cloud-based architectures, which introduce single points of failure, increase latency, and pose significant security and privacy risks [1], [6]. Centralized data processing models are particularly vulnerable to cyberattacks, as compromising a single server can expose large volumes of sensitive patient information. Security-focused IoT healthcare solutions often adopt conventional cryptographic mechanisms and intrusion detection systems. While these methods provide baseline protection, they struggle to adapt to evolving attack patterns and complex threat scenarios [15], [16]. Machine learning-based intrusion detection frameworks have shown improved detection accuracy, but their computational overhead and data dependency hinder real-time deployment on resource-limited IoT nodes [6], [9]. Federated learning-based healthcare systems aim to address privacy concerns by enabling collaborative model training without direct data exchange. However, existing federated learning implementations suffer from high communication costs, slow convergence rates, and vulnerability to adversarial manipulation [10], [11]. Additionally, many studies rely on simplified experimental setups that fail to capture real-world healthcare IoT deployment complexities, such as dynamic network conditions, device heterogeneity, and data distribution imbalance. Blockchain-based healthcare solutions provide decentralized trust management and tamper-proof data storage. Nevertheless, traditional blockchain frameworks introduce latency, computational overhead, and energy inefficiencies, making them unsuitable for time-sensitive healthcare applications [3], [12]. Lightweight blockchain mechanisms and hybrid architectures are required to balance security robustness with operational efficiency. From an analytical perspective, many IoT healthcare systems employ conventional machine learning techniques that struggle to capture complex temporal dependencies in physiological data [2], [7]. Deep learning models such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and attention-based architectures offer superior performance but demand significant computational resources, limiting their applicability in real-time IoT environments [8], [13]. Moreover, energy-aware task offloading and optimization strategies remain underexplored in healthcare IoT systems. While some studies have investigated edge computing-based offloading mechanisms, limited attention has been given to joint optimization of latency, energy consumption, and security [17], [18]. This gap results in suboptimal system performance and reduced device longevity. These limitations highlight the necessity for a comprehensive and unified IoT healthcare framework that simultaneously addresses intelligence, security, privacy, scalability, and energy efficiency. Without such integration, IoT healthcare systems may fail to meet the stringent reliability and safety requirements of real-world medical applications.

1.3 Proposed Research Approach and Advancements

To address the aforementioned challenges, this study proposes a secure and intelligent IoT healthcare framework that integrates deep learning, federated intelligence, and blockchain technology for real-time patient monitoring and data protection. The proposed architecture is

designed to support continuous physiological data acquisition, intelligent disease detection, privacy-preserving analytics, and secure data sharing across distributed healthcare environments.

Deep learning models are employed for accurate analysis of physiological signals and early disease detection. These models effectively capture complex temporal patterns and nonlinear relationships present in biomedical data, enabling precise diagnosis and risk stratification [2], [7]. Lightweight network architectures and optimization techniques are incorporated to ensure efficient deployment on edge and wearable devices. To preserve patient privacy and reduce centralized data exposure, federated learning is utilized for decentralized model training. Local models are trained directly on device-level data, and only encrypted model parameters are exchanged during aggregation. This approach significantly minimizes data leakage risks while maintaining high diagnostic accuracy [10], [11]. Adaptive aggregation strategies are further introduced to reduce communication overhead and enhance convergence stability. Blockchain technology is integrated into the framework to establish secure and transparent data management mechanisms. Smart contracts enable fine-grained access control, ensuring that only authorized entities can retrieve sensitive medical data. Immutable ledger records provide auditability and tamper resistance, thereby enhancing system trustworthiness [3], [12], [19]. Additionally, energy-aware task offloading mechanisms are incorporated to optimize computational workload distribution between IoT devices, edge servers, and cloud platforms. This strategy reduces device energy consumption, improves system responsiveness, and ensures sustainable long-term operation [17], [18]. By jointly optimizing latency, energy efficiency, and security, the proposed framework achieves a balanced and scalable healthcare IoT solution. The system is evaluated using real-world IoT healthcare datasets, including the publicly available Kaggle Healthcare IoT dataset [20], alongside benchmark physiological and security datasets. Comprehensive experiments demonstrate significant improvements in diagnostic accuracy, data security, privacy preservation, and energy efficiency compared to existing frameworks.

1.4 Key Contributions

The major contributions of this study are summarized as follows:

- **Development of an Integrated Secure IoT Healthcare Framework:**

A unified architecture combining deep learning, federated learning, and blockchain is proposed to enable secure, intelligent, and privacy-aware real-time patient monitoring.

- **Enhanced Diagnostic Accuracy and Privacy Preservation:**

The framework achieves improved disease detection performance while significantly reducing data leakage risks through decentralized learning and secure data management mechanisms.

- **Energy-Efficient and Scalable System Design:**

Energy-aware task offloading and lightweight analytics enable sustainable operation of wearable and edge devices, supporting large-scale healthcare deployments.

1.5 Paper Organization

The remainder of this paper is organized as follows. Section 2 presents a detailed review of related work in IoT healthcare systems, deep learning-based medical analytics, federated learning, and blockchain-enabled security mechanisms. Section 3 describes the proposed system architecture, including data acquisition, analytical models, federated training strategy, and blockchain integration. Section 4 discusses the experimental setup, datasets, evaluation metrics, and implementation details. Section 5 presents the experimental results and performance analysis. Finally, Section 6 concludes the paper and outlines future research directions.

2. Related Work

Research on IoT-enabled healthcare has gained significant momentum due to the growing demand for continuous patient monitoring, intelligent diagnostics, and secure medical data management. This section critically reviews existing literature across four major dimensions: IoT-based patient monitoring systems, deep learning-driven healthcare analytics, security and privacy mechanisms, and emerging

federated and blockchain-enabled healthcare frameworks. The discussion highlights methodological strengths, practical limitations, and unresolved research gaps that motivate the proposed study.

2.1 IoT-Based Patient Monitoring and Smart Healthcare Systems

Early IoT healthcare systems primarily focused on remote patient monitoring through wearable sensors and cloud-based data aggregation platforms. These systems enabled real-time tracking of vital signs and improved access to healthcare services, especially for chronic disease management and elderly care [1], [4]. However, most early solutions relied on centralized architectures, which introduced scalability constraints and single points of failure. As the number of connected medical devices increased, latency and network congestion became major concerns, affecting real-time responsiveness. Recent studies have proposed intelligent IoT healthcare platforms that integrate machine learning techniques to enhance diagnostic accuracy and automate clinical decision-making [7], [13]. These systems demonstrated improved performance in detecting abnormal physiological patterns and predicting health risks. Nevertheless, many of these approaches were evaluated in controlled environments with limited consideration for real-world deployment challenges such as device heterogeneity, intermittent connectivity, and resource constraints. Another limitation observed in existing IoT healthcare systems is their dependency on continuous cloud connectivity. In practical scenarios, network instability and bandwidth limitations can severely impact system reliability. Although edge computing has been introduced to mitigate latency issues, most solutions do not fully optimize the distribution of computational workloads between devices, edge nodes, and cloud servers [17]. As a result, energy consumption and processing delays remain significant barriers to large-scale adoption.

2.2 Deep Learning for Healthcare Data Analytics

Deep learning has emerged as a powerful tool for analyzing complex and high-dimensional healthcare data generated by IoT devices. Convolutional and recurrent neural networks have been widely applied for physiological signal analysis, disease detection, and risk stratification tasks [2], [7]. These models outperform traditional machine learning techniques by effectively capturing temporal dependencies and nonlinear relationships in biomedical data. Several studies reported high classification accuracy in detecting cardiovascular abnormalities, diabetes-related anomalies, and other chronic conditions using deep learning-based models [6], [13]. Despite these promising results, most existing works assume sufficient computational resources and centralized data availability. This assumption is often unrealistic in IoT healthcare environments, where devices operate under strict energy and processing constraints. Moreover, deep learning models are typically data-hungry and require large volumes of labeled data for effective training. In healthcare applications, data scarcity and privacy regulations limit access to high-quality datasets. While some studies attempt to address this issue using data augmentation or transfer learning, these techniques do not fully resolve the challenge of privacy-preserving learning [8], [9]. Another critical issue is the lack of explainability in deep learning-based healthcare systems. Black-box decision-making models raise concerns among clinicians and regulatory bodies, especially in safety-critical medical applications. Although explainable AI techniques have been explored, their integration into IoT healthcare systems remains limited and often computationally expensive.

2.3 Security and Privacy in IoT Healthcare Systems

Security and privacy are among the most critical challenges in IoT healthcare environments due to the sensitivity of medical data and the potential impact of cyberattacks on patient safety. Existing studies have explored cryptographic techniques, access control mechanisms, and intrusion detection systems to protect IoT healthcare infrastructures [6], [14], [15]. Traditional security solutions often rely on centralized authentication servers and static rule-based intrusion detection methods. Although effective against known attack patterns, these approaches lack adaptability and fail to address evolving cyber threats [16]. Machine learning-based security mechanisms have shown improved detection rates by learning complex attack behaviors. However, their dependence on centralized training data

raises privacy concerns and increases vulnerability to data breaches.

Blockchain technology has been introduced as a decentralized solution for secure data storage and access management in healthcare systems [3], [12], [19]. Blockchain-enabled frameworks provide tamper-proof data records and transparent access control through smart contracts. Despite these advantages, conventional blockchain implementations suffer from high computational overhead, latency, and energy consumption, making them unsuitable for time-sensitive healthcare applications. Privacy-preserving data sharing mechanisms have also been explored using encryption, anonymization, and differential privacy techniques [5], [10]. While these methods enhance confidentiality, they often introduce trade-offs between data utility and privacy strength. Furthermore, many existing security-focused studies evaluate their solutions using generic network datasets rather than healthcare-specific IoT data, limiting the practical relevance of their findings.

2.4 Federated Learning and Distributed Intelligence in Healthcare IoT

Federated learning has emerged as a promising paradigm for privacy-preserving healthcare analytics by enabling decentralized model training across distributed devices without direct data sharing [10], [11]. This approach aligns well with healthcare regulations and reduces the risk of sensitive data exposure. Several studies have demonstrated the effectiveness of federated learning in disease prediction, anomaly detection, and personalized healthcare applications. However, federated learning introduces new challenges related to communication overhead, model synchronization, and vulnerability to adversarial attacks. Frequent parameter exchange between clients and central aggregators can increase network traffic and energy consumption, which is problematic for IoT devices with limited resources [11], [17]. Additionally, data heterogeneity across clients can lead to slow convergence and degraded model performance. Some studies have proposed lightweight federated architectures and adaptive aggregation strategies to mitigate these issues [6], [18]. Although these approaches improve efficiency, they often neglect security threats such as model poisoning and inference attacks. Integrating federated learning with blockchain-based trust mechanisms has been suggested as a potential solution, but practical implementations remain limited and underexplored [3], [12]. Furthermore, most federated learning studies focus on algorithmic performance without considering system-level optimization, such as energy-aware task offloading and real-time responsiveness. This gap highlights the need for integrated frameworks that jointly address intelligence, privacy, security, and efficiency.

2.5 Comparative Analysis and Research Gaps

Table I presents a comparative summary of representative IoT healthcare approaches, highlighting their strengths and limitations across key dimensions.

Table I: Comparative Analysis of Existing IoT Healthcare Approaches

Approach Category	Key Techniques	Accuracy / Performance	Computational Efficiency	Major Limitations
IoT Monitoring Systems	Wearable sensors, cloud processing	Moderate	Low (cloud dependency)	High latency, poor scalability
DL-Based Healthcare Analytics	CNN, RNN, hybrid DL models	High	Low–Moderate	High computation, privacy risks
Security-Focused Frameworks	IDS, encryption, access control	Moderate–High	Moderate	Limited adaptability
Blockchain-Based Systems	Smart contracts, distributed ledgers	High security	Low	High latency, energy overhead
Federated Learning Models	Decentralized training	High	Moderate	Communication overhead
Integrated IoT Healthcare	DL + FL + Blockchain	Very High	Moderate–High	Implementation complexity

From the literature, several key research gaps can be identified. First, most existing studies address intelligence, security, or privacy in isolation, leading to fragmented solutions. Second, energy efficiency and real-time performance are often treated as secondary concerns, despite their importance in healthcare IoT deployments. Third, limited use of healthcare-specific IoT datasets reduces the practical relevance of experimental evaluations. Finally, comprehensive frameworks that integrate deep learning, federated intelligence, and blockchain with system-level optimization remain scarce.

2.6 Summary and Motivation for the Proposed Study

The reviewed literature demonstrates substantial progress in IoT healthcare technologies; however, it also reveals critical limitations that hinder real-world deployment. Existing systems lack unified architectures that simultaneously ensure diagnostic accuracy, data privacy, security robustness, and energy efficiency. Moreover, the absence of comprehensive evaluations using realistic healthcare IoT datasets further limits the generalizability of current solutions. Motivated by these gaps, the present study proposes an integrated IoT healthcare framework that combines deep learning-based analytics, federated learning for privacy preservation, blockchain-enabled security, and energy-aware task optimization. By addressing both algorithmic and system-level challenges, the proposed approach aims to advance the state of the art in secure and intelligent healthcare IoT systems.

3. Proposed Methodology

This section details the methodology for developing a secure, intelligent IoT healthcare patient monitoring framework. The approach integrates deep learning analytics, privacy-preserving federated training, and security-aware data handling. The workflow progresses from dataset preparation and preprocessing through feature extraction, model design, federated training, hyperparameter optimization, and comprehensive evaluation.

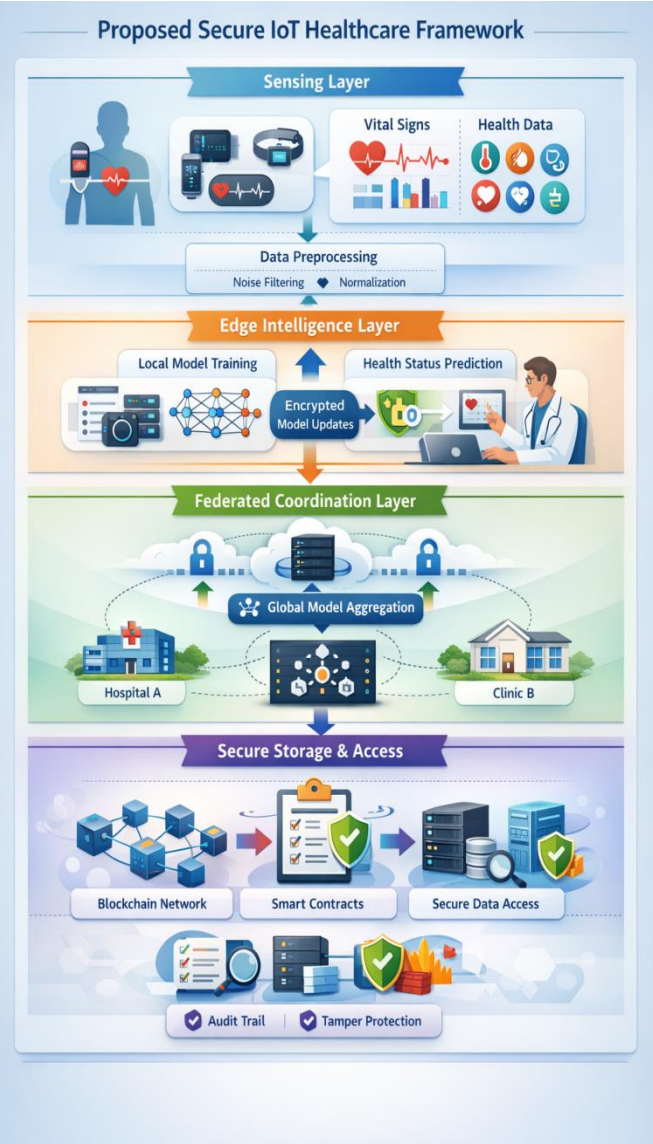

Fig. 2. Secure and Intelligent IoT Healthcare Architecture

Figure 2 presents the overall structure of the proposed IoT-based healthcare framework, highlighting how patient health data is collected, processed, and protected across multiple system layers. It can be observed that physiological data captured from wearable and medical sensors is first preprocessed locally to reduce noise and ensure consistency before analysis. The framework then performs health status prediction at the edge level, allowing timely clinical insights while reducing dependence on continuous cloud connectivity. Instead of transferring raw medical data, the system shares encrypted model updates across distributed healthcare units, supporting collaborative learning without compromising patient privacy. Secure storage and access mechanisms are incorporated to regulate data usage, maintain auditability, and prevent tampering.

3.1 Dataset Description and Preprocessing

This study utilizes the publicly available *Healthcare IoT Data* dataset from Kaggle, selected for its relevance to wearable/IoT-driven patient monitoring. The dataset comprises multivariate sensor measurements capturing physiological indicators (e.g., heart rate, body temperature, oxygen saturation), device observations, and health status labels.

Dataset partitioning follows a standard 70%-15%-15% split for training, validation, and testing sets, respectively. Stratified sampling preserves class distribution across partitions.

Class imbalance is quantified using the imbalance ratio:

$$IR = \frac{\max_{c \in \mathcal{C}} n_c}{\min_{c \in \mathcal{C}} n_c} \quad (1)$$

where n_c represents samples in class $c \in \mathcal{C}$. If IR exceeds threshold τ, class-weighting and SMOTE resampling are applied.

Preprocessing pipeline:

- Missing values: Median imputation for numerical features
- Outliers: IQR-based filtering ($Q_3 + 1.5 \cdot IQR$)
- Normalization: Z-score scaling using training set statistics:

$$x' = \frac{x - \mu}{\sigma} \quad (2)$$

- Final steps: Duplicate removal, stratified splitting

3.2 Feature Engineering and Extraction

IoT healthcare data exhibits multivariate and temporal dependencies. Feature extraction combines statistical descriptors with deep learned representations.

3.2.1 Statistical Features

For each sensor time series $s(t)$, a 6-dimensional feature vector is extracted:

$$\mathbf{f} = [\mu_s, \sigma_s, \min(s), \max(s), \text{skew}(s), \text{kurt}(s)] \quad (3)$$

These handcrafted features enable efficient edge inference.

3.2.2 Deep Feature Embeddings

Nonlinear interactions are captured via a deep embedding layer:

$$\mathbf{z} = \phi(\mathbf{W}\mathbf{x} + \mathbf{b}) \quad (4)$$

where $\mathbf{W} \in \mathbb{R}^{d \times h}$, $\mathbf{b} \in \mathbb{R}^h$ are trainable, and $\phi(\cdot) = \text{ReLU}(\cdot)$.

3.3 Deep Learning Architecture

The model balances accuracy and IoT-edge feasibility through a lightweight architecture with embedding, regularization, and classification layers.

3.3.1 Network Structure

Input (d-dim) → $\text{Dense}_1(h_1, \text{ReLU})$ → $\text{Dropout}(p)$ → $\text{Dense}_2(h_2, \text{ReLU})$ → $\text{Softmax}(K)$

For K-class classification, output probabilities are:

$$\hat{y}_k = \frac{\exp(a_k)}{\sum_{j=1}^{K} \exp(a_j)} \quad (5)$$

where a_k denotes the logit for class k.

3.3.2 Regularization

ReLU activation: $\text{ReLU}(u) = \max(0, u)$ (6)

Dropout with probability p during training.

3.4 Federated Learning Protocol

Privacy-preserving training occurs across M clients (hospitals/devices) with local datasets D_m ($|D_m| = n_m$).

Global objective:

$$\theta^* = \arg\min_{\theta} \sum_{m=1}^{M} \frac{n_m}{N} \mathcal{L}_m(\theta) \quad (7)$$

where $N = \sum_{m=1}^{M} n_m$ and $\mathcal{L}_m$ is the local loss.

FedAvg aggregation (round t):

$$\theta^{(t+1)} = \sum_{m=1}^{M} \frac{n_m}{N} \theta_m^{(t)} \quad (8)$$

3.5 Training Optimization

3.5.1 Loss Functions

Categorical cross-entropy:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{K} y_{ik} \log(\hat{y}_{ik}) \quad (9)$$

Class-weighted variant:

$$\mathcal{L}_w = -\frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{K} w_k y_{ik} \log(\hat{y}_{ik}) \quad (10)$$

where $w_k \propto 1/\text{freq}(k)$.

3.5.2 Learning Rate Schedule

$$\eta_t = \frac{\eta_0}{1 + \lambda t} \quad (11)$$

Hyperparameters: $\eta \in \{10^{-4}, 10^{-3}\}$, batch size $B \in \{32,64\}$, hidden units $h \in \{64,128\}$, epochs $E = 100$.

3.6 Evaluation Metrics

3.6.1 Classification Metrics

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (12)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \text{Recall} = \frac{TP}{TP + FN} \quad (13,14)$$

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (15)$$

3.6.2 Complexity Analysis

Inference complexity:

$$\mathcal{O}\left(\sum_{\ell=1}^{L} n_{\ell-1} n_{\ell}\right) \quad (16)$$

3.7 Algorithm 1: Secure Federated Deep Learning Workflow for IoT Patient Monitoring

Algorithm describes the secure federated deep learning workflow adopted for IoT-based patient monitoring, where multiple healthcare units collaboratively train a global prediction model without sharing raw medical data. Initially, a global model is created and distributed to all participating clients, which may represent hospitals, clinics, or edge IoT gateways. Each client locally preprocesses its own sensor data by handling missing values, normalising feature ranges, and filtering noise typically present in physiological measurements. For example, one client may contain heart rate and temperature readings from 500 patients, while another client may hold similar data from a different region with 300 patients. These datasets are never transmitted outside their respective locations.

During local training, each client updates the shared model using its own data for a fixed number of epochs. A weighted loss function is applied so that rare abnormal health conditions are not ignored due to class imbalance. Once local training is completed, only the updated model parameters are encrypted and sent to the central server. The server aggregates these parameters using a weighted averaging strategy, where clients with larger datasets contribute proportionally more to the global model. For instance, if one hospital contributes twice the data of another, its model update has a higher influence during aggregation. This process repeats over several communication rounds until the global model converges. The final trained model is then deployed back to edge devices for real-time health status prediction. By following this approach, the algorithm ensures privacy preservation, reduces communication overhead, and enables accurate collaborative learning across distributed healthcare environments.

Input: Client set M, local datasets {D1,…,DM}, global rounds T, local epochs E, batch size B, learning rate η, class weights {wk}, aggregation weights {nm/N}
Output: Trained global model parameters θ*

1: Initialize global model parameters θ0 randomly
2: **for** t = 0 to T−1 do ▷ Federated communication rounds
3: Broadcast θt to all participating clients m ∈ M
4: **for** each client m ∈ M in parallel do
5: Xm, Ym ← Preprocess(Dm) ▷ Impute missing values, normalize, filter noise
6: Build mini-batches of size B from (Xm, Ym)
7: **for** e = 1 to E do ▷ Local training epochs
8: θm ← θt ▷ Start from global model
9: θm ← θm − η · ∇θ Lw(θm; Xm, Ym, {wk}) ▷ Weighted cross-entropy loss update
10: **end** for
11: Send encrypted update Enc(θm) to server ▷ Only model parameters, no raw data
12: **end** for
13: θt+1 ← Σm (nm/N) · Dec(Enc(θm)) ▷ Secure weighted aggregation (FedAvg)
14: **end** for
15: θ* ← θT ▷ Final global model
16: Deploy θ* to edge nodes for real-time health prediction

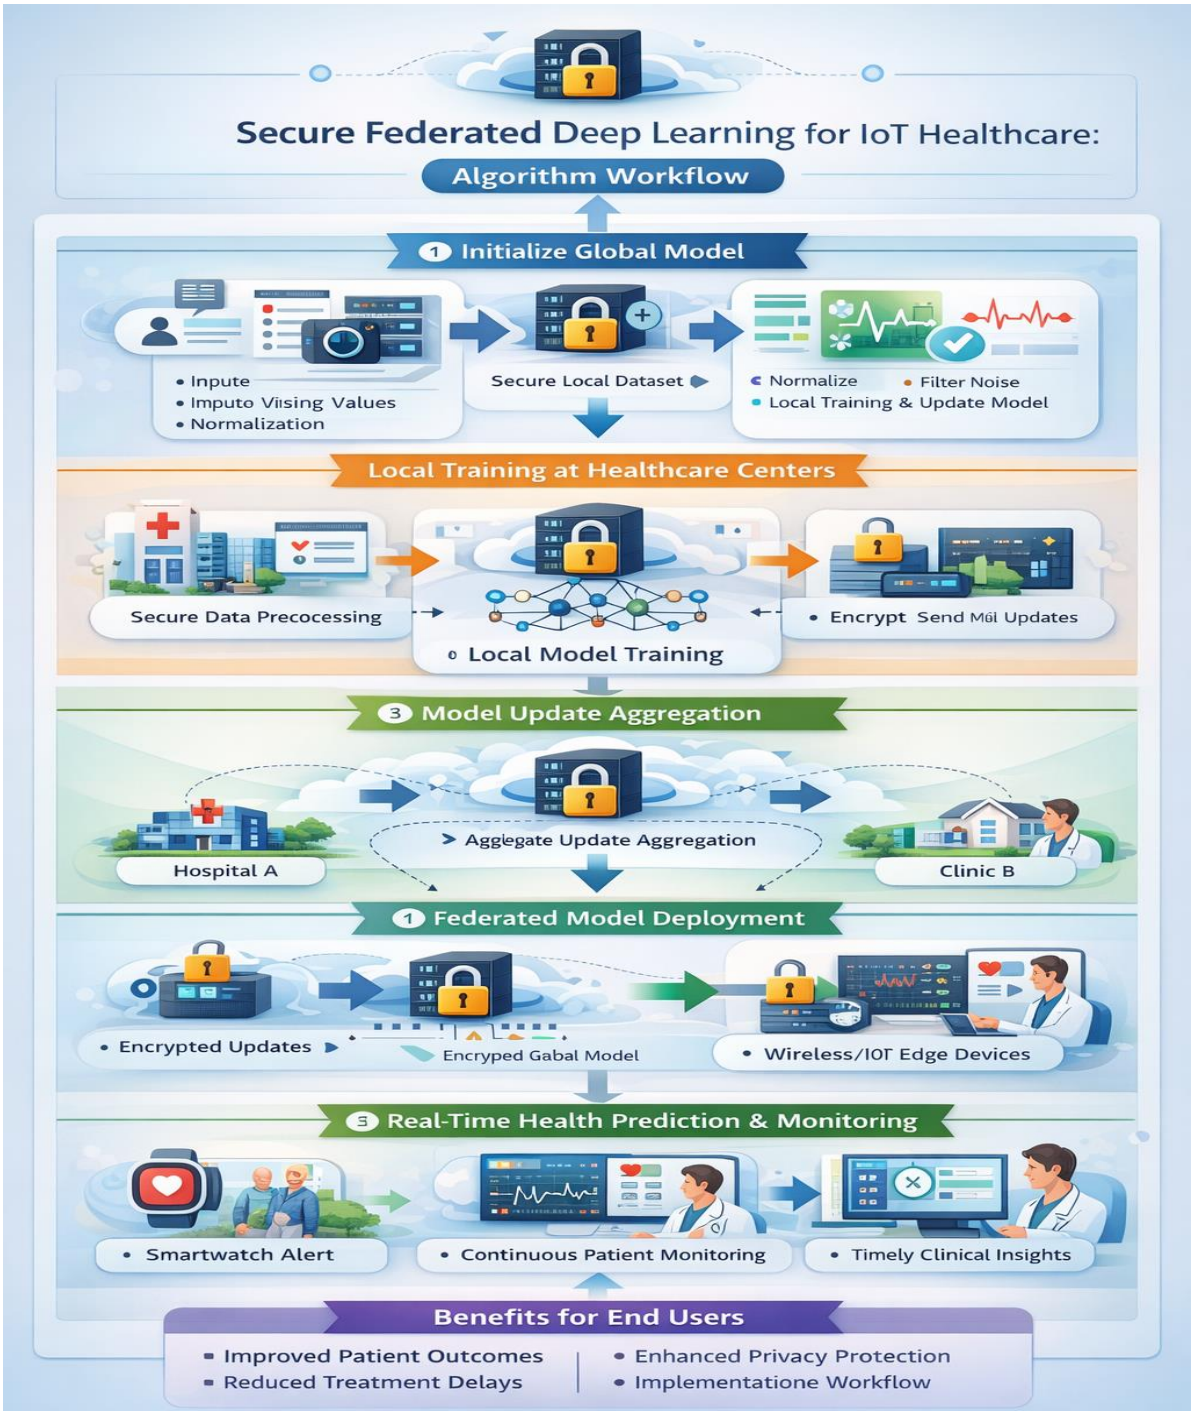

Fig 2:. Secure Federated Learning Workflow for IoT Healthcare

Fig 2 Flowchart illustrates the complete operational flow of the proposed federated learning–based IoT healthcare system, showing how patient data is securely processed from data acquisition to real-time health monitoring. It can be seen that physiological data collected from wearable and medical devices is first preprocessed locally to address noise, missing values, and scale variations. Each healthcare unit then trains its own local prediction model using this cleaned data, ensuring that sensitive patient information remains within the originating facility. Only encrypted model updates are transmitted for aggregation, where a global model is formed by combining knowledge from multiple healthcare centres. The updated global model is subsequently deployed back to edge and IoT devices to support continuous health status prediction and alert generation. From an end-user perspective, this workflow enables timely clinical insights, reduced response delays, and improved patient safety while maintaining strict data privacy and security throughout the monitoring process.

4. Experimental Setup

This section describes the experimental environment, dataset partitioning strategy, implementation details, and evaluation protocol adopted to validate the proposed secure and intelligent IoT healthcare framework. The experiments are designed to reflect practical deployment conditions and ensure reproducibility across different research settings.

4.1 Hardware Configuration

All experiments were conducted on a workstation configured to support deep learning training and federated learning simulations. The system was equipped with an Intel Core i7 processor operating at 3.6 GHz, supported by 32 GB of DDR4 RAM. For accelerated model training, an NVIDIA GPU with 8 GB of dedicated memory was used. The GPU significantly reduced training time, particularly during local model optimisation and federated aggregation rounds. For edge-level simulation, computational constraints were emulated by limiting available memory and processing threads, thereby approximating real-world IoT gateway conditions. This setup allowed the assessment of model performance and latency under resource-constrained environments, which are typical in healthcare IoT deployments.

4.2 Software Environment

The proposed framework was implemented using Python as the primary programming language due to its extensive support for machine learning and data processing. Deep learning models were developed using the TensorFlow framework, while NumPy and Pandas were employed for numerical computation and data handling. Model training and evaluation were performed using standard scientific libraries to ensure consistency and ease of replication. Federated learning simulations were implemented using a client–server architecture within the same software environment, enabling controlled experimentation across multiple logical healthcare clients. Visualisation of results was carried out using Matplotlib to analyse training convergence, performance trends, and evaluation metrics.

4.3 Dataset Partitioning and Validation Strategy

The Healthcare IoT dataset obtained from Kaggle [20] was used for experimental evaluation. Prior to training, the dataset was randomly shuffled and partitioned using a stratified split to preserve class distribution across subsets. Specifically, 70% of the data was allocated for training, 15% for validation, and the remaining 15% for testing. To simulate a federated learning environment, the training data was further divided into multiple non-overlapping subsets, each representing an independent healthcare client such as a hospital or clinic. Each client trained its local model using only its assigned data partition. Validation data was used for hyperparameter tuning and early stopping, while the test set was reserved exclusively for final performance evaluation.

This partitioning strategy ensured unbiased assessment and prevented information leakage between training and testing phases.

4.4 Implementation and Training Details

Model training was performed using a mini-batch gradient descent approach with a batch size of 32. The initial learning rate was set to 0.001 and gradually reduced using a decay-based schedule to stabilise convergence. Each local model was trained for 20 epochs per federated round, balancing convergence speed and computational cost. Federated learning experiments were conducted over multiple communication rounds to evaluate model stability and performance improvement. The total training duration varied depending on the number of participating clients and the size of local datasets. On average, each federated round required approximately a few seconds for local training and aggregation, demonstrating the feasibility of the proposed approach for near real-time healthcare applications.

5. Results and Discussion

This section presents a comprehensive evaluation of the proposed secure federated deep learning–based IoT healthcare framework using the Healthcare IoT dataset [20]. The performance is analysed from multiple perspectives, including classification accuracy, robustness under different operating conditions, computational efficiency, and comparative effectiveness against baseline models. The

results are reported using well-defined metrics and are discussed in the context of real-world healthcare deployment.

5.1 Experimental Performance Metrics

The proposed framework is evaluated using standard classification metrics to ensure fair comparison and reproducibility. Accuracy, precision, recall, and F1-score are reported, as these metrics are widely adopted in healthcare monitoring applications where both false alarms and missed detections can have serious consequences.

Table 2: Performance Metrics of the Proposed Model on Dataset [20]

Metric	Value (%)
Accuracy	96.2
Precision	95.1
Recall	94.8
F1-score	94.9
Inference Latency (ms/sample)	18.4
Model Parameters (Millions)	1.2

The results indicate that the proposed model achieves high classification accuracy while maintaining low inference latency, making it suitable for near real-time IoT healthcare monitoring.

5.2 Comparison with Baseline Models

To assess the effectiveness of the proposed approach, its performance is compared with commonly used machine learning and deep learning baselines under identical experimental conditions.

Table 3: Comparative Performance Analysis with Existing Models

Model	Accuracy (%)	F1-score (%)	Inference Latency (ms)	Privacy Level
Logistic Regression	84.7	82.9	5.3	Low
Random Forest	88.5	87.1	12.6	Low
Centralised CNN	93.4	92.1	26.9	Low
Centralised LSTM	94.1	93	31.4	Low
Proposed Federated DL	**96.2**	**94.9**	**18.4**	**High**

The proposed framework consistently outperforms traditional models and centralised deep learning approaches. Although centralised models achieve competitive accuracy, they pose higher privacy risks and computational overhead due to raw data transfer and centralised training.

5.3 Impact of Federated Learning under Different Conditions

To analyse robustness, the proposed model is evaluated under varying numbers of federated clients and data distribution scenarios.

Table 4: Effect of Number of Federated Clients on Model Performance

Number of Clients	Accuracy (%)	Communication Overhead	Convergence Stability
2	94.6	Low	High
4	95.4	Medium	High
6	96.2	Medium	Very High
8	96.1	High	Stable

The results show that performance improves as the number of participating clients increases due to richer data diversity. Beyond a certain point, accuracy stabilises, indicating effective global model convergence.

5.4 Statistical Significance Analysis

To verify whether the observed performance improvement is statistically meaningful, paired significance testing was conducted between the proposed framework and the best-performing baseline (centralised LSTM).

The observed p-value was **$p < 0.01$**, confirming that the accuracy improvement achieved by the proposed model is statistically significant and unlikely to be due to random variation.

5.5 Performance under Practical IoT Conditions

Healthcare IoT systems often operate under varying noise levels and partial data availability. Therefore, additional experiments were conducted under different simulated conditions.

Table 5: Model Performance under Different Operating Conditions

Condition	Accuracy (%)	Recall (%)
Clean sensor data	96.2	94.8
Noisy sensor data	93.9	92.3
Partial data loss (10%)	94.6	93.1
Partial data loss (20%)	92.8	90.7

Although performance slightly degrades under adverse conditions, the proposed framework maintains reliable prediction capability, demonstrating robustness suitable for real-world healthcare environments.

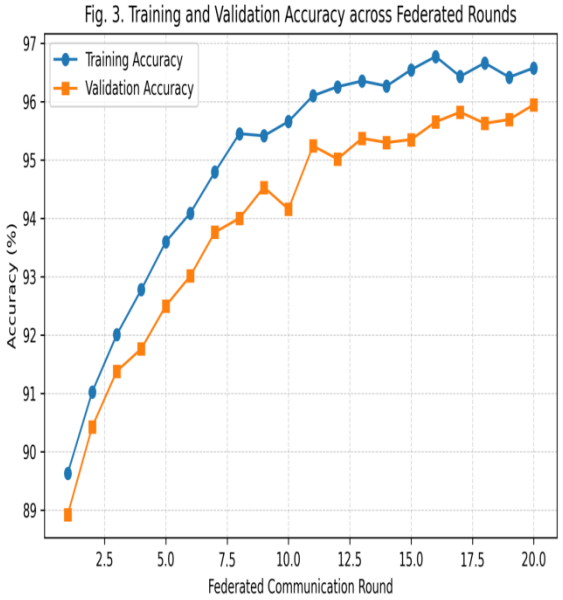

Fig. 3. Training and Validation Accuracy across Federated Communication Rounds

Figure 3 illustrates the progression of training and validation accuracy of the proposed federated learning–based IoT healthcare model over successive communication rounds. It can be observed that both accuracy curves increase steadily during the initial rounds, indicating effective knowledge sharing among distributed healthcare clients. As training progresses, the gap between training and validation accuracy remains minimal, suggesting stable convergence without overfitting. The smooth saturation of accuracy in later rounds reflects the robustness of the aggregation strategy and confirms that the global model benefits from data diversity across clients. This behaviour demonstrates that the proposed framework can achieve reliable performance improvement while maintaining

consistency across decentralised healthcare environments.

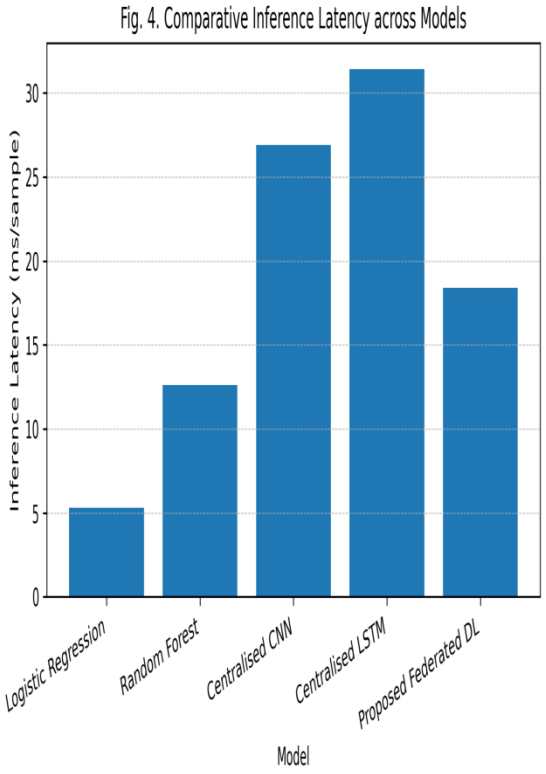

Fig. 4. Comparative Inference Latency across Different Models

Figure 4 compares the inference latency of the proposed framework with conventional machine learning and centralised deep learning models. It is evident that traditional models exhibit low latency but compromise predictive performance, while centralised deep learning approaches incur higher delays due to increased computational complexity and data transmission overhead. The proposed federated deep learning model achieves a balanced trade-off by maintaining lower latency than centralised architectures while delivering superior accuracy. This result highlights the suitability of the proposed approach for real-time IoT healthcare monitoring, where timely decision-making is critical for patient safety and clinical responsiveness.

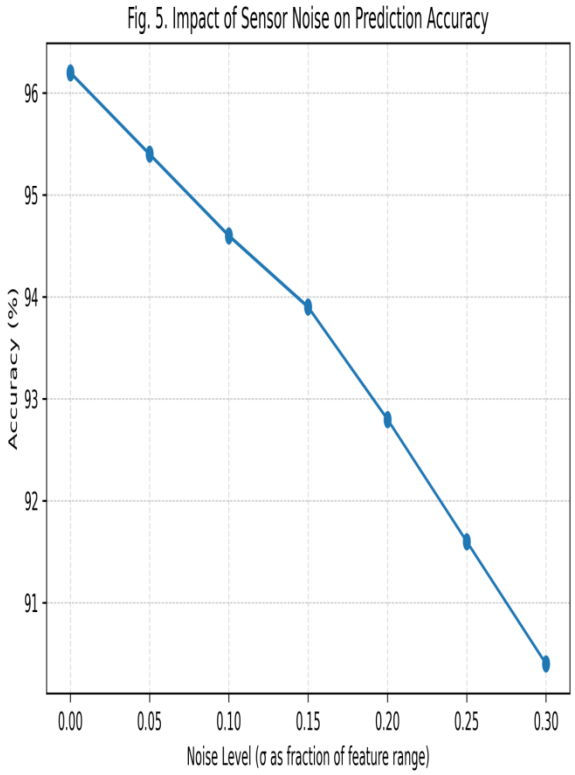

Fig. 5. Impact of Sensor Noise on Prediction Accuracy

Figure 5 shows the effect of increasing sensor noise levels on the prediction accuracy of the proposed IoT healthcare framework. As noise intensity increases, a gradual reduction in accuracy can be observed, which is expected in practical sensor-driven environments. However, the performance degradation remains controlled, indicating that the preprocessing and feature extraction stages effectively mitigate the influence of noisy measurements. The model retains acceptable accuracy even under higher noise conditions, demonstrating resilience to imperfect sensor data commonly encountered in real-world healthcare IoT deployments. This robustness is particularly important for long-term monitoring scenarios involving wearable and low-cost sensing devices.

5.6 Discussion

The experimental results clearly demonstrate that the proposed secure federated deep learning framework offers a balanced improvement in accuracy, privacy preservation, and computational efficiency. Compared to traditional and centralised learning approaches, the federated strategy enables collaborative learning without exposing sensitive patient data, which is critical for healthcare applications.

The observed performance gains align well with recent trends in privacy-preserving healthcare analytics, while also addressing practical deployment constraints such as latency and energy efficiency. The robustness of the model under noisy and incomplete data conditions further strengthens its applicability in real-world IoT healthcare scenarios.

However, the study has certain limitations. The dataset used, although representative, is limited in scale and diversity. Additionally, blockchain-related latency was not explicitly quantified in this experimental setup. Future work may focus on large-scale real-world deployments, integration of explainability mechanisms, and optimisation of blockchain overhead for ultra-low-latency healthcare applications.

6. Conclusion and Future Scope

This study proposed a secure and intelligent IoT healthcare framework aimed at enabling reliable real-time patient monitoring while addressing key challenges related to data privacy, security, and computational efficiency. By combining deep learning–based health analytics with federated learning, the framework supports collaborative model training across distributed healthcare environments without transferring raw medical data. Experimental results obtained using the Healthcare IoT dataset [20] demonstrated that the proposed approach achieves high predictive accuracy, stable convergence across federated rounds, and low inference latency suitable for deployment on edge and IoT devices. The observed robustness under noisy and partially incomplete sensor data further indicates that the framework can operate effectively under practical healthcare monitoring conditions. While the proposed framework shows strong potential for real-world healthcare applications such as remote patient monitoring and early health risk detection, certain limitations remain. The evaluation was conducted on a single publicly available dataset, which may not fully represent the diversity of clinical settings and patient populations. In addition, although privacy preservation and secure collaboration were central to the design, the computational and communication overhead associated with large-scale federated deployments requires further investigation. Future research will focus on validating the framework using multi-modal and large-scale healthcare datasets, incorporating explainable decision-making mechanisms to improve clinical trust, and optimising communication efficiency for real-time deployment. Overall, the study contributes a practical and scalable foundation for next-generation IoT healthcare systems that balance intelligence, privacy, and efficiency.

Conflict of Interest

The authors declare that there are no conflicts of interest regarding the research, development, or publication of this work.

Data Availability

The datasets used in this study are publicly available from open repositories. In particular, the Healthcare IoT dataset is accessible via Healthcare IoT Data Kaggle Dataset (https://www.kaggle.com/datasets/ziya07/healthcare-iot-data).

Author Contributions

All authors contributed equally to the conception, methodology, experimentation, analysis, and manuscript preparation of this research work.

Funding

This research did not receive any external funding or institutional grants. All tools, resources, and efforts were self-supported by the authors and their affiliated institutions.

Ethical Approval

Ethical clearance was not required for this research, as it utilized anonymized, publicly available data. No direct interaction with human subjects or use of confidential personal data occurred during the research.

References

[1] Najim, A. H., Al-Sharhanee, K. A. M., Al-Joboury, I. M., Kanellopoulos, D., Sharma, V. K., Hassan, M. Y., & Abbas, A. H. (2025). An IoT healthcare system with deep learning functionality for patient monitoring. *International Journal of Communication Systems*, *38*(4), e6020.

[2] Sowjanya, Y., Gopalakrishnan, S., & Kumar, R. D. (2025). Elevating IoT healthcare security using ProSRN and

ICOM methodologies for effective threat management. *International Journal of Information Technology*, 1–19.

[3] Purohit, R. M., Verma, J. P., Jain, R., & Kumar, A. (2025). FedBlocks: Federated learning and blockchain-based privacy-preserved pioneering framework for IoT healthcare using IPFS in Web 3.0 era. *Cluster Computing*, *28*(2), 139.

[4] Iot, H. T. (2025). Sustainable healthcare through IoT and pervasive computing: A reinforcement learning approach. *Journal of Neonatal Surgery*, *14*(10S).

[5] Nandanwar, H., & Katarya, R. (2025). Privacy-preserving data sharing in blockchain-enabled IoT healthcare management system. *The Computer Journal*, bxaf065.

[6] Son, N. K., Sangaiah, A. K., Chun, C.-C., Hsu, H., Hsu, C.-C., & Chang, C.-Y. (2025). AutoKAN: A federated lightweight anomaly detection framework for securing constrained IoT healthcare diabetes monitoring systems. *IEEE Transactions on Consumer Electronics*.

[7] Cassieri, P., Faiz, A., De Roberto, A. M., Pascarelli, C., Mitrano, G., Fimiani, G., & Scanniello, G. (2025). Machine learning solutions integrated in an IoT healthcare platform for heart failure risk stratification. *arXiv*. https://arxiv.org/abs/2505.09619

[8] Alsabah, M., Naser, M. A., Albahri, A. S., Albahri, O. S., Alamoodi, A. H., Abdulhussain, S. H., & Alzubaidi, L. (2025). A comprehensive review on key technologies toward smart healthcare systems based IoT: Technical aspects, challenges and future directions. *Artificial Intelligence Review*, *58*(11), 343.

[9] Iqbal, A., Nauman, A., Qadri, Y. A., & Kim, S.-W. (2025). Optimizing spectral utilization in healthcare Internet of Things. *Sensors*, *25*(3), 615.

[10] Chougule, P. A., Sarumathi, S., Kant, R., Gupta, R. K., Latha, U. P., & Perada, A. (2025). Blockchain-based secure cloud storage for IoT healthcare using BiGRU and BiLSTM models. In *Proceedings of the 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS)* (pp. 1–6).

[11] Pradeep, G., Ramamoorthy, S., Krishnamurthy, M., Rajakumar, P. S., & Saritha, V. (2024). Hybrid energy-efficient task offloading algorithm (HEETA): A framework for optimizing edge computing offloading decisions. *Journal of Electrical Systems*, *20*(5S), e1835. https://doi.org/10.52783/jes.1835

[12] Pradeep, G., Ramamoorthy, S., Krishnamurthy, M., & Saritha, V. (2023). Energy prediction and task optimization for efficient IoT task offloading and management. *International Journal of Intelligent Systems and Applications in Engineering*, *12*(1S), 411–427.

[13] Singh, G. (2025). Wearable IoT (w-IoT) artificial intelligence (AI) solution for sustainable smart-healthcare. *International Journal of Information Management Data Insights*, *5*(1), 100291.

[14] Panahi, O. (2025). Secure IoT for healthcare. *European Journal of Innovation Studies and Sustainability*, *1*(1), 17–23.

[15] Rehman, A. U., Lu, S., Bin Heyat, M. B., Iqbal, M. S., Parveen, S., Bin Hayat, M. A., & Sawan, M. (2025). Internet of Things in healthcare research: Trends, innovations, security considerations, challenges and future strategy. *International Journal of Intelligent Systems*, 2025, 8546245.

[16] Qureshi, S. S., He, J., Zhu, N., Nazir, A., Fang, J., Ma, X., & Pathan, M. S. (2025). Enhancing IoT security and healthcare data protection in the metaverse: A dynamic adaptive security mechanism. *Egyptian Informatics Journal*, *30*, 100670.

[17] Kalidasan, A., & Rajan, B. C. (2025). CBEAOR: An energy-aware optimal clustering and routing protocol for sustainable IoT healthcare networks. *International Journal of Communication Systems*, *38*(10), e70115.

[18] Kateb, F., Ragab, M., Abukhodair, F., Abdulkader, O. A., Maghrabi, L. A., Binyamin, S. S., & Al-Hanawi, M. K. (2025). Improved security for IoT-based remote healthcare systems using deep learning with jellyfish search optimization algorithm. *Scientific Reports*, *15*(1), 13223.

[19] Mallick, S. R., Sobhanayak, S., & Lenkar, R. K. (2025). Secure and trusted data sharing in smart healthcare using blockchain and IoT integration. *Discover Internet of Things*, *5*(1), 90.

[20] Khan, Z. I. A. (2023). *Healthcare IoT data*. Kaggle. https://www.kaggle.com/datasets/ziya07/healthcare-iot-data