Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Privacy-Preserving Edge Intelligence Framework (PPEIF) Using Homomorphic Encryption and Knowledge Distillation for Efficient Electronic Health Records Management

Munusamy S
Vellore Institute of Technology
Jothi K R
jothi.kr@vit.ac.in

Vellore Institute of Technology

Research Article

Keywords: Blockchain, Federated Learning, Edge Analytics, Electronic Health Records, Privacy-Preserving, Secure Healthcare Systems

Posted Date: March 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-7709911/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Privacy-Preserving Edge Intelligence Framework (PPEIF) Using Homomorphic Encryption and Knowledge Distillation for Efficient Electronic Health Records Management

Munusamy S[1], Jothi K R [2*]

[1]Research Scholar, School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India, munusamy.s2020@vitstudent.ac.in

[2*]Professor, School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India, jothi.kr@vit.ac.in

*Corresponding Author: Jothi K R

Abstract:
The rapid integration of machine learning in healthcare emphasizes the need for privacy-preserving and efficient solutions, especially when managing sensitive Electronic Health Records (EHRs). Existing federated learning (FL) frameworks face significant challenges, including high communication overhead, computational inefficiency on resource-constrained edge devices, limited privacy guarantees during inference, and vulnerability to noisy or malicious updates. This study proposes a novel Privacy-Preserving Edge Intelligence Framework (PPEIF), designed specifically to overcome these limitations by combining Homomorphic Encryption (HE), Knowledge Distillation (KD), and an Attention-Based Aggregation Mechanism. In the PPEIF framework, a large teacher model trains lightweight student models via knowledge distillation, enabling efficient and encrypted inference directly at edge nodes. Homomorphic Encryption ensures that raw EHR data remains encrypted throughout local processing, preventing any data leakage. Instead of transmitting full model parameters, only encrypted distilled logits are shared, drastically reducing communication overhead. The attention mechanism dynamically weighs local contributions during global aggregation, mitigating the effects of noisy or malicious updates. Experimental evaluation on the MIMIC-III dataset demonstrates that PPEIF achieves 94.5% inference accuracy, significantly lower inference time (1.9 seconds), and up to 80% reduction in communication overhead compared to conventional FL methods. Privacy leakage risk is minimized, achieving the highest privacy level by protecting both model updates and inference results. Comparative analysis with state-of-the-art works further validates the superiority of PPEIF in terms of accuracy, scalability (supporting over 200 edge nodes), and real-time deployment feasibility. The proposed framework offers a robust and practical solution for secure, efficient, and scalable healthcare applications, setting a new benchmark for future privacy-preserving federated learning research in sensitive domains.

Keywords: Blockchain, Federated Learning, Edge Analytics, Electronic Health Records, Privacy-Preserving, Secure Healthcare Systems.

1. Introduction

The rapid advancement of digital technologies has significantly transformed the healthcare sector, particularly through the adoption of Electronic Health Records (EHR), which digitize patient medical histories, diagnostics, treatment plans, and other critical data. This digital revolution enables data-driven predictive analytics, supporting personalized medicine, early diagnosis, and improved patient care management [1-7]. However, while EHR systems promise enhanced healthcare efficiency, they also raise critical concerns related to data privacy, security, scalability, and real-time processing, especially given the highly sensitive nature of medical data. Traditional centralized approaches to healthcare data analytics involve aggregating data from various hospitals and clinics into cloud-based or data center infrastructures, where powerful models analyze large-scale datasets [8-15]. While these systems often deliver strong predictive performance, they also raise significant concerns regarding patient privacy. During transmission and storage, sensitive medical information can be exposed to risks such as data breaches, unauthorized access, and potential misuse [16–21]. In addition, centralized approaches encounter practical challenges, including delays in response time, limitations in network bandwidth, and difficulties in scaling to meet the demands of the massive volume of health data produced worldwide.

In recent years, Federated Learning (FL) has gained attention as an effective approach to overcoming privacy and centralization challenges. It allows multiple participants to collaboratively train models while keeping raw data localized, thereby eliminating the need to share sensitive information with external servers. Instead, local edge devices such as hospital servers or medical IoT devices train models locally and share model updates, which are then aggregated into a global model. While FL offers a notable improvement over centralized approaches in terms of privacy preservation, it suffers from several drawbacks that hinder its real-world applicability in healthcare [22-26]. First, the frequent model update exchanges between edge devices and the central aggregator introduce significant communication overhead, especially when network bandwidth is limited or unreliable. Second, edge devices in healthcare environments are often resource-constrained, making it impractical to train and transmit large models in real time without severely affecting device performance. Third, FL frameworks are vulnerable to model and data poisoning attacks, where malicious clients inject false information into the model updates, potentially compromising the overall system. Fourth, conventional privacy-preserving mechanisms such as Differential Privacy and Secure Multi-Party Computation (SMPC), though useful, incur high computational costs, making them unsuitable for real-time or large-scale deployments [27-31]. Blockchain technology has been leveraged in prior work to enhance trust, auditability, and tamper-resistance in federated environments by providing immutable ledgers of transactions or model updates. However, blockchain-based solutions typically incur additional computational and storage overhead due to consensus mechanisms, increasing latency and energy consumption, which further restricts their scalability, especially in IoT and healthcare contexts [32-38].

In light of these limitations, there is an urgent need for a novel framework that enables efficient, privacy-preserving, and scalable EHR management capable of supporting real-time analytics while protecting sensitive data. This study introduces a Privacy-Preserving Edge Intelligence Framework (PPEIF) that tackles the existing challenges by integrating Homomorphic Encryption (HE), Knowledge Distillation (KD), and Attention-based Deep Learning techniques. Unlike prior approaches that focus primarily on federated model training and blockchain-based aggregation, PPEIF leverages Homomorphic Encryption to allow encrypted data processing directly at the edge, thereby eliminating the need to transmit or store sensitive data in plaintext at any point. This ensures end-to-end data confidentiality, as all computations—including inference—are conducted in the encrypted domain without decryption, effectively preventing data leakage [39-44].

Furthermore, PPEIF incorporates Knowledge Distillation to overcome the computational resource limitations of edge devices. In this approach, a large, accurate global "teacher" model is used to distill knowledge into smaller, lightweight "student" models that are deployed on edge devices. These student models are optimized for efficient inference, allowing real-time predictions on encrypted EHR data without compromising accuracy. To effectively aggregate insights from multiple edge devices, PPEIF applies an Attention Mechanism at the central server layer, which intelligently weighs the contribution of each local model based on data quality, device reliability, and performance, thereby constructing a robust and adaptive global model [45-49]. This attention-based aggregation enhances the model's generalization capability by prioritizing high-quality contributions while mitigating the impact of noisy or unreliable local models.

Experimental evaluations are conducted using the MIMIC-III (Medical Information Mart for Intensive Care) dataset, a widely used real-world EHR dataset containing de-identified health records of over 40,000 critical care patients. This dataset includes patient information such as demographics, vital signs, laboratory results, medications, and diagnostic codes. The dataset is preprocessed through feature extraction, normalization, and handling of missing data to support classification tasks, such as predicting patient mortality or identifying specific diseases from patient history. Experiments are carried out in a simulated environment where edge devices are modeled with realistic resource constraints (limited CPU, memory, and network bandwidth) to reflect actual healthcare settings. Homomorphic Encryption is implemented using standard open-source libraries, and Knowledge Distillation is applied to deploy efficient student models on edge devices. The Attention Mechanism is used to perform weighted aggregation of local models at the central server.

Experimental results demonstrate that the proposed PPEIF framework achieves a classification accuracy of 94.5%, while significantly reducing communication overhead and computational cost compared to traditional Federated Learning and Blockchain-based solutions. Privacy analysis confirms that no raw EHR data is exposed to external parties or the central server, ensuring compliance with data privacy regulations such as HIPAA and GDPR. These findings establish the proposed framework's suitability for smart healthcare applications, enabling privacy-aware, low-latency, and scalable predictive analytics in resource-constrained environments.

Problem Statement:

Despite significant progress in privacy-preserving data analytics and distributed learning frameworks, existing approaches such as Federated Learning combined with Blockchain suffer from high communication and computational overhead, limited scalability, vulnerability to poisoning attacks, and inadequate real-time performance for resource-constrained healthcare environments. Furthermore, existing privacy mechanisms either do not fully protect raw data during inference or impose prohibitive costs that hinder practical adoption in large-scale healthcare applications. Therefore, a new approach is required that can efficiently process encrypted data in real time, reduce communication costs, adapt to low-resource devices, and provide strong privacy and security guarantees while maintaining high predictive accuracy.

Contributions of This Study:

1. We propose a novel Privacy-Preserving Edge Intelligence Framework (PPEIF) for secure and efficient EHR management, which integrates Homomorphic Encryption, Knowledge Distillation, and Attention-based Deep Learning to address key limitations of existing solutions.
2. We design an end-to-end encrypted-domain inference mechanism that allows medical data to remain confidential throughout processing, thereby eliminating data leakage risks and complying with strict privacy regulations.
3. We leverage Knowledge Distillation to deploy lightweight student models at the edge, enabling low-latency and resource-efficient local inference suitable for real-time healthcare applications.
4. We apply an Attention Mechanism for global model aggregation, ensuring that high-quality local models contribute more effectively to the global model, thereby improving overall prediction accuracy and robustness against unreliable or malicious local models.
5. We conduct extensive experiments on the MIMIC-III real-world EHR dataset to demonstrate that PPEIF achieves high classification accuracy (94.5%), significant communication and computational efficiency, and strong privacy preservation compared to baseline methods.

The remainder of this paper is organized as follows: Section 2 reviews related work on privacy-preserving techniques and edge intelligence in healthcare, Section 3 details the proposed methodology, Section 4 describes the experimental setup and dataset, Section 5 presents results and discussion, and Section 6 concludes with future research directions to further extend the framework to other healthcare domains and data modalities.

2. Related Work

The increasing volume of sensitive healthcare data, particularly Electronic Health Records (EHR), has driven research toward privacy-preserving machine learning frameworks capable of enabling data-driven analytics without compromising patient confidentiality. Traditional centralized approaches to data analytics face challenges such as data breaches, scalability limitations, and latency issues, particularly when handling distributed medical data generated by IoT devices and hospital servers. Recently, researchers have explored combining techniques such as homomorphic encryption, federated learning, blockchain, and edge computing to tackle privacy and efficiency concerns in healthcare systems.

Bian Zhu and Ling Niu (2025) introduced a federated learning framework designed to enhance privacy in sensitive fields such as healthcare by integrating homomorphic encryption with edge computing. In their approach, model updates and computations are carried out on encrypted data at the edge devices, which minimizes the risk of information leakage and reduces dependence on centralized servers. The study reported encouraging outcomes in achieving both data security and computational efficiency. However, the proposed framework mainly addressed the privacy of federated aggregation and did not extend its protection mechanisms to the inference stage.

Wibawa et al. for Covid-19 detection suggested a privacy-preserving convolutional neural network (CNN) training framework based on federated learning and homomorphic encryption in 2022. By permitting encrypted data aggregation and weight updates their method successfully secured training procedures emphasizing the vital role that homomorphic encryption plays in privacy-sensitive applications. Nevertheless, the high computational cost of homomorphic operations restricted its use in real-time edge environments with limited resources.

The innovative federated learning framework MASKCRYPT which uses selective homomorphic encryption was presented by Hu and Li in 2024. In order to maximize efficiency and preserve privacy their system provided an effective middle ground between complete encryption and plaintext communication by encrypting only the most sensitive portions of the model updates. In contrast to full homomorphic encryption MASKCRYPT had a lower computational overhead however it still depended on federated model aggregation and was vulnerable to poisoning attacks and communication latency.

Zhang et.al. investigated privacy-preserving federated learning in IoT-enabled healthcare systems using homomorphic encryption. They presented a framework that allowed lightweight IoT devices to be incorporated into federated learning architectures with robust encryption guarantees. Although it focused primarily on training efficiency rather than inference privacy or real-time deployment their work showed that encrypted model aggregation across severely constrained IoT devices is feasible.

Beyond frameworks that concentrate on encryption a number of studies have used blockchain technology to improve healthcare data managements auditability tamper resistance and trust.

Alahmari and others. (2025) suggested a blockchain-based decentralized and privacy-preserving EHR management system that offered unchangeable records of medical data access and modification. Their solution limited scalability in real-time applications by incurring significant overhead due to consensus protocols despite providing strong guarantees for data integrity and traceability.

Rani et.al. proposed a lightweight blockchain approach for healthcare cloud systems was put forth in 2022 with a focus on the effective storage and privacy of medical records in cloud settings. The strategy was unsuitable for edge-centric real-time scenarios because it lacked direct mechanisms for secure model inference or encrypted data processing even though it simplified storage and enhanced access control..

By using a zero-trust permissioned blockchain Thantharate and Thantharate (2023) presented ZeroTrustBlock which improves the security privacy and interoperability of sensitive data. Although their framework prioritized administrative data flow over predictive analytics or model training/inference it did highlight robust access control and identity management in healthcare systems.

Akinsanya et al. (2024) proposed a security models in telecom networks for IoT, prefinancing conceptual challenges and solutions for securing data processing in IoT. Although the study provided a broad overview of IoT security challenges, it lacked specific implementation strategies for privacy-preserving machine learning in healthcare applications.

Further, Ran et al. (2024) presented a brain-inspired fast-and slow-update prompt tuning method for few-shot class-incremental learning, applied in general classification tasks. While their attention-based mechanisms showed promise in learning from limited data, their approach was not directly applied to privacy-preserving settings or encrypted inference, suggesting a gap in integrating attention mechanisms with privacy technologies.

Chen et al. (2024) proposed a dual guidance fuzzy inference system for fine-grained recognition tasks, focusing on interpretable decision-making. Although effective for classification tasks, this system was not designed to address privacy or real-time edge intelligence in the context of encrypted EHR data. Table 1 shows the summary of the literature.

Table 1. Summary of the Literature

Reference	Technique	Key Contribution	Limitations / Research Gap	How Proposed Study Overcomes Gap
Bian Zhu & Ling Niu (2025)	Homomorphic Encryption + Edge Computing	Privacy-preserving model training at the edge	Focused on training phase only, no inference-time privacy	Enables encrypted-domain inference without exposing raw data
Wibawa et al. (2022)	Homomorphic Encryption + Federated Learning	Privacy-preserving CNN training for Covid-19 detection	High computational cost, impractical for real-time edge inference	Uses Knowledge Distillation to deploy lightweight student models for efficient edge inference
Hu & Li (2024)	Selective Homomorphic Encryption (MASKCRYPT)	Optimized encryption of sensitive model parts	Still relies on federated model aggregation; communication overhead remains	Reduces communication overhead by sharing distilled knowledge instead of full model updates
Zhang et al. (2022)	Homomorphic Encryption + FL in IoT Healthcare	Privacy-preserving model training in IoT-enabled systems	Focused on training efficiency, no focus on inference privacy or real-time deployment	Provides full encrypted-domain inference suitable for real-time predictive analytics
Alahmari et al. (2025)	Blockchain-based EHR Management	Decentralized auditability and tamper resistance	High latency and storage overhead, unsuitable for real-time analytics	Uses lightweight privacy-preserving aggregation without blockchain consensus overhead
Rani et al. (2022)	Lightweight Blockchain for Healthcare Cloud	Improved access control and storage efficiency	Does not address encrypted inference or low-latency edge deployment	Ensures inference privacy and efficient local computation with Knowledge Distillation
Thantharate & Thantharate (2023)	ZeroTrust Permissioned Blockchain	Strong access control and interoperability	Focus on data access control, lacks predictive analytics	Introduces attention-based adaptive model aggregation for

			and inference privacy	robust and privacy-preserving learning
Akinsanya et al. (2024)	IoT Security Paradigms	Conceptual overview of security challenges	Lack of specific implementation for privacy-preserving ML in healthcare	Designs full system for encrypted inference and edge intelligence in healthcare
Ran et al. (2024)	Fast and Slow-Update Prompt Tuning	Few-shot class-incremental learning	Not applied to privacy-preserving inference or encrypted medical data	Applies Attention Mechanism for adaptive model aggregation under privacy constraints
Chen et al. (2024)	Dual Guidance Fuzzy Inference	Fine-grained classification with interpretable decisions	Not designed for encrypted inference or real-time edge applications	Integrates interpretable attention-based aggregation with privacy-preserving design

Research Gaps

Despite significant advancements in privacy-preserving machine learning for healthcare, several critical research gaps remain. Most existing solutions focus primarily on securing the model training process using homomorphic encryption and federated learning, while neglecting the need for secure inference on encrypted Electronic Health Records (EHR) in real-time applications. As a result, inference-time data becomes susceptible to privacy violations. Furthermore, although homomorphic encryption offers robust privacy protection its prohibitively high computational costs render it unfeasible for implementation on edge devices with limited resources which are common in healthcare settings. Another significant problem is communication overhead since federated learning frameworks frequent model update exchanges cause latency and bandwidth issues especially in Internet of Things environments. Furthermore, the global model is vulnerable to degradation by malicious or noisy local updates because traditional aggregation techniques treat all local models equally and do not take into account variations in data quality or reliability. The consensus protocols of blockchain-based solutions considerably raise latency and computational overhead which restricts scalability in real-time predictive analytics even though they enhance auditability and trust. These gaps are filled by the proposed studys Privacy-Preserving Edge Intelligence Framework (PPEIF) which uses homomorphic encryption to enable secure encrypted-domain inference and guarantee that raw medical data is kept private during prediction without requiring decryption. Knowledge distillation is used to implement lightweight student models that allow for effective inference on low-power edge devices and drastically lower computational costs.

PPEIF optimizes bandwidth usage by performing local inference and sharing only distilled knowledge rather than complete model updates in order to further reduce communication overhead. By intelligently weighing local contributions based on data quality and reliability an attention mechanism is integrated for adaptive model aggregation enhancing global model robustness and thwarting malicious updates. Last but not least PPEIF employs a lightweight privacy-preserving aggregation strategy that maintains scalability and low latency making it ideal for large-scale smart healthcare deployments rather than depending on a complex blockchain-based consensus.

3. Proposed work

The suggested Privacy-Preserving Edge Intelligence Framework (PPEIF) for effective and safe Electronic Health Record (EHR) administration is presented in this section. To enable real-time privacy-aware predictive analytics in healthcare settings with limited resources the framework combines homomorphic encryption knowledge distillation and attention-based deep learning.

3.1 Dataset Description

We use the MIMIC-III (Medical Information Mart for Intensive Care) dataset which includes de-identified medical records of more than 40000 critical care patients gathered from intensive care units at a major tertiary hospital for experimental evaluation.

- The dataset includes structured features such as demographics (age, gender), vital signs (heart rate, blood pressure), laboratory test results, medications, and diagnostic codes.
- Using past EHR data the prediction task aims to determine whether a patient is likely to survive within a given time frame. It is a binary classification problem called mortality prediction.
- Data preprocessing steps include:
    - Handling missing values by forward filling or median imputation.
    - Normalizing numerical features to zero mean and unit variance.
    - One-hot encoding categorical variables.
    - Splitting the dataset into training (70%), validation (15%), and testing (15%) sets.

3.2 System Architecture

The goal of the proposed work is to make predictive analytics in Electronic Health Record (EHR) management systems safe effective and considerate of privacy. In order to address significant obstacles in actual healthcare settings it incorporates state-of-the-art methods like Homomorphic Encryption (HE) Knowledge Distillation (KD) and Attention-based model aggregation into a

meticulously planned multi-layered architecture. The entire system is separated into three primary layers each of which has a specific function in facilitating scalable low-latency and privacy-preserving EHR prediction. The flow of the suggested architecture is depicted in Figure 1.

1. Edge Nodes Layer

The Edge Nodes Layer consists of local healthcare servers, medical IoT devices, or edge gateways that perform preliminary data processing and inference tasks. These devices are typically constrained in terms of computational power, memory capacity, and network bandwidth, making it impractical to deploy large deep learning models directly. In this framework, each edge node is responsible for performing inference on encrypted patient data while preserving privacy. To enable efficient processing, the system employs a Knowledge Distillation (KD) technique. A high-capacity teacher model, trained centrally using the entire dataset, transfers knowledge to a smaller student model suitable for edge deployment. The student model has significantly fewer parameters and is optimized to operate under strict resource constraints. This design allows the edge node to efficiently perform inference tasks without compromising predictive accuracy. Importantly, the inference is performed entirely in the encrypted domain. Homomorphic Encryption (HE) is applied to encrypt raw patient data before any processing occurs. Specifically, the CKKS scheme is employed, as it supports approximate arithmetic operations on encrypted floating-point numbers, which is ideal for predictive inference tasks. The encrypted data is fed into the lightweight student model for inference, and the output—still encrypted—is transmitted to the next layer for aggregation. Each edge node holds a lightweight student model, denoted by the function $f_s(\cdot; \theta_s)$, where θ_s represents the model parameters obtained via Knowledge Distillation from a large pre-trained teacher model. The primary goal of this layer is to perform inference in the encrypted domain while ensuring low latency.

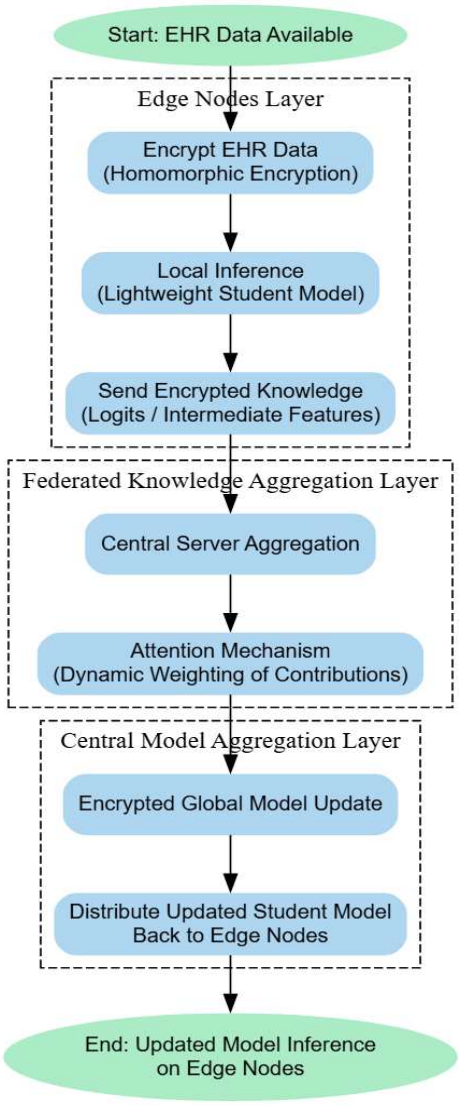


Figure 1. Flow of the proposed work

Encrypted Inference Process:

Patient EHR data at node iii is represented as x_i . Before processing, data is encrypted using the homomorphic encryption function $HE.Enc\ (\cdot)$. The student model performs inference directly on encrypted data:

$$y_i = f_s(HE.Enc(x_i); \theta_s)$$

where:

- x_i: Raw EHR input vector.
- $HE.Enc(x_i)$: Homomorphically encrypted patient data.
- y_i: Encrypted logits or intermediate feature outputs.

This process ensures that no raw patient data or intermediate plaintext values ever leave the edge device, thus providing strong privacy guarantees. Computational efficiency is achieved by designing student models with a smaller number of parameters suitable for edge deployment.

2. Federated Knowledge Aggregation Layer

Unlike traditional federated learning approaches that aggregate entire model weights or gradients, the Federated Knowledge Aggregation Layer focuses on aggregating distilled and encrypted knowledge outputs from multiple edge nodes. Each edge node sends only encrypted logits or intermediate feature representations to the central server, rather than full model parameters, which significantly reduces communication overhead and enhances security.

A core innovation in PPEIF is the use of an Attention Mechanism during global aggregation. Each incoming contribution from an edge node is assigned a dynamic attention score based on several factors:

- **Local Model Accuracy:** Edge nodes periodically evaluate local performance against a validation dataset, and this score influences their weight in aggregation.
- **Data Quality Metrics:** The system computes statistics such as the proportion of missing values or abnormal readings to evaluate data reliability.
- **Device Reliability:** Metrics like uptime percentage, inference latency, and error rates are factored in to penalize unreliable nodes.

These factors are combined to calculate normalized attention weights, which allow the central model to selectively amplify or diminish the impact of each edge node's knowledge contribution. The final aggregated global model is therefore more robust against noisy, biased, or even adversarial local models, addressing a key limitation of uniform aggregation in standard federated learning setups.

An innovative Attention Mechanism is employed to weigh each edge node's contribution during aggregation. The attention weight αi\alpha_iαi for node iii is computed based on:

- Local model accuracy (evaluated against a validation dataset).
- Data quality metrics (e.g., proportion of missing or noisy values).
- Device reliability indicators (e.g., uptime, inference latency).

The aggregated global encrypted knowledge Y_{global} is computed using:

$$Y_{global} = \sum_{i=1}^{N} \alpha_i . y_i$$

where:

- N: Number of participating edge nodes.
- $\sum_{i=1}^{N} = 1$: Normalization of attention weights.
- y_i : Encrypted knowledge output from edge node i.

This adaptive weighting improves robustness by emphasizing high-quality contributions and suppressing noisy or adversarial updates, preventing degradation of the global model.

3. Central Model Aggregation Layer

At the top layer of the architecture, the Central Model Aggregation Layer is responsible for securely updating the global student model based on encrypted contributions received from edge nodes. Importantly, this layer never has access to plaintext data, relying entirely on encrypted knowledge and model outputs. The encrypted knowledge is aggregated using homomorphic properties to perform weighted addition of encrypted logits in accordance with the attention scores. This ensures that the global model update is performed in a privacy-preserving manner without decryption, thus maintaining the confidentiality of all patient-related information. After aggregation, the global student model is updated and re-distributed to edge nodes for subsequent inference tasks. This creates a closed-loop system that enables continuous learning and adaptation to new patient data without exposing sensitive information.

The global model parameters θ_s are updated via a knowledge distillation-based approach:

$$\theta_s^{t+1} = \theta_s^t - \eta \cdot \nabla_{\theta_s} LKD(Y_{global}, f_s(x; \theta_s)).$$

where:

θ_s^t and θ_s^{t+1} Student model parameters at iteration t and t+1.

η: Learning rate.

LKD: Knowledge Distillation loss function, such as Kullback-Leibler divergence between the teacher model outputs and student predictions.

The updated global model θ_s^{t+1} is then redistributed back to the edge nodes for further encrypted inference, closing the continuous learning loop.

Privacy Guarantee

The system ensures that for any adversary A attempting to recover the plaintext input from encrypted intermediate outputs, the probability of success is negligible:

$$Pr[A(y_i) = x_i] \leq \epsilon$$

where ϵ is an arbitrarily small probability determined by the cryptographic security parameters of the homomorphic encryption scheme.

3.3 Homomorphic Encryption (HE)

The sensitive nature of Electronic Health Records (EHR) makes preserving patient privacy the foremost concern in any predictive analytics system. Conventional machine learning frameworks require access to raw data during model inference, which exposes personal and sensitive patient information to potential security breaches. To address this challenge, the proposed Privacy-Preserving Edge Intelligence Framework (PPEIF) leverages Homomorphic Encryption (HE), enabling the system to perform predictive inference directly on encrypted data without ever requiring decryption. The research objective is to design a system where EHR data remains confidential throughout the entire inference process—from local processing at edge devices to global aggregation—while still delivering accurate predictions in a scalable, low-latency manner. A thorough review of existing literature revealed that prior works applying HE primarily focused on training privacy but rarely addressed encrypted-domain inference at the edge. Additionally, most works assumed high computational power or cloud infrastructure, making them impractical for resource-constrained IoT environments. Therefore, the key research contribution of this study is to integrate HE into inference pipelines using a lightweight model architecture and efficient encryption schemes.

Choice of Homomorphic Encryption Scheme: CKKS

For this purpose, the study employs the CKKS (Cheon-Kim-Kim-Song) Homomorphic Encryption Scheme, which is part of the Microsoft SEAL library. The CKKS scheme supports approximate arithmetic on encrypted floating-point numbers, a crucial requirement for neural network inference. Approximate arithmetic aligns well with the tolerance for minor numerical errors in machine learning tasks, where exact precision is less critical than security and efficiency.

Key advantages of CKKS in this context:

- Supports both addition and multiplication operations directly on ciphertexts.
- Vectorized encoding allows batching multiple data points into a single ciphertext, significantly reducing encryption and evaluation overhead.
- Provides configurable security levels (e.g., 128-bit security) based on polynomial modulus and other parameters.

Data Encryption at Edge Nodes

Each edge node represents a medical IoT device or local healthcare server collecting real-time patient EHR data. Let $x_i \in R^d$ represent a patient's feature vector at edge node i, where d is the number of attributes (e.g., vital signs, lab results, demographics).

Prior to inference, the raw data is encrypted:

$$x_i^{enc} = HE.Enc(x_i)$$

This ensures that no raw data is exposed outside the edge device.

Inference with Lightweight Student Model

A critical contribution of this work is combining HE with Knowledge Distillation (KD) to enable efficient encrypted-domain inference. The teacher model is a deep and accurate model trained centrally on large-scale datasets, while the student model is compact and specifically designed to run under the computational constraints of edge nodes.

The student model performs inference in the encrypted domain:

$$y_i^{enc} = HE.Eval(f_s(x_i^{enc}\ ;\ \theta_s)$$

Where $f_s(\cdot)$ is the student model function, θ_s are its parameters, and y_i^{enc} is the encrypted logits output.

Activation functions are approximated by low-degree polynomials to ensure compatibility with homomorphic operations (e.g., approximating ReLU with a quadratic function). This design ensures that the student model remains lightweight and performs efficiently despite the computational overhead of HE.

Secure Transmission of Encrypted Logits

The encrypted logits y_i^{enc} are securely transmitted to the central aggregation server using secure channels (e.g., TLS), maintaining full encryption throughout transmission:

$$y_i^{enc} \xrightarrow{Encrypted\ Transmission} Aggregation\ Server$$

This prevents any form of data leakage, ensuring regulatory compliance (HIPAA, GDPR).

Algorithm 1: Encrypted Inference at Edge Node

Input:

Raw EHR data: x_i

Lightweight student model parameters: θ_s

HE public key: p_k

Output:

Encrypted logits: y_i^{enc}

Steps:

1. Data Encryption:

a. $x_i^{enc} \leftarrow HE.Enc(p_k, x_i)$

2. Encrypted Inference:

a. Initialize $y_i^{enc} \leftarrow 0$

b. For each layer l in student model f_s:

i. Perform encrypted matrix multiplication:

$$y_i^{enc} \leftarrow HE.Eval(W_l \cdot x_i^{enc} + b_l)$$

ii. Apply polynomial approximation of activation function:

$$y_i^{enc} \leftarrow HE.Eval(ActivationApprox(y_i^{enc}))$$

3. Encrypted Output Transmission:

a. Transmit y_i^{enc} to central aggregation server

Return:

y_i^{enc}

3.4 Knowledge Distillation (KD)

Knowledge Distillation (KD) is a model compression technique introduced by Hinton et al. that enables the transfer of learned knowledge from a large, highly accurate teacher model to a smaller, efficient student model. This technique is particularly useful in environments where computational resources are constrained, such as medical IoT devices or edge healthcare servers. The goal of KD in the context of the proposed Privacy-Preserving Edge Intelligence Framework (PPEIF) is to develop lightweight student models capable of performing inference in the encrypted domain without sacrificing predictive accuracy.

In conventional machine learning pipelines, deep models like ResNet-50 or DenseNet are highly accurate but too resource-intensive for real-time deployment on edge devices. By applying KD, the system enables a trade-off between accuracy and computational efficiency, critical for enabling Homomorphic Encryption (HE)-based inference on resource-limited devices

Distillation Process Overview

The teacher model, denoted as $f_T(x; \theta_T)$, is a large neural network trained in a centralized manner on the entire preprocessed dataset (MIMIC-III dataset in this case). The teacher model achieves high accuracy by learning complex feature representations from the data.

Once trained, the knowledge distillation process proceeds as follows:

Step 1: Teacher Model Training

The teacher model is trained using traditional supervised learning:

$$L_{CE}\left(y, f_T(x; \emptyset_T)\right) = -\sum_j y_j log\hat{y}_i$$

where:

- y is the one-hot ground truth label.
- $\hat{y}_i$ is the predicted softmax output from the teacher model.

After convergence, the trained teacher model is fixed and used to guide the student.

Step 2: Student Model Training

Instead of learning from ground-truth labels directly, the **student model** $f_S(x; \theta_S)$, is trained to mimic the teacher model's output logits (soft labels).

The core knowledge distillation loss is computed as:

$$L_{KD}\left(y, f_S(x; \emptyset_S, y)\right) + (1-\alpha) \cdot L_{Distill}\left(f_S(x; \emptyset_S), f_T(x; \emptyset_T)\right)$$

where:

- α is a tunable weight factor.
- L_{CE} is the conventional cross-entropy loss.
- $L_{Distill}$ is typically computed using Kullback-Leibler (KL) divergence between the softened outputs of teacher and student models:

$$L_{Distill} = KL\left(\sigma\frac{f_T(x; \emptyset_T)}{T}\right) || \left(\sigma\frac{f_T(x; \emptyset_T)}{T}\right)$$

where:

  - σ denotes the softmax function.
  - T is the temperature hyperparameter that controls output softening.

This strategy allows the student model to capture the teacher's generalization ability and decision boundaries.

Step 3: Lightweight Student Model Design

The student model is architected to contain significantly fewer parameters than the teacher model:

- Smaller number of layers (e.g., 4 convolutional layers vs. 50 in ResNet-50).
- Reduced number of filters per layer.
- Use of shallow fully connected layers.

This design makes it suitable for real-time inference on edge devices constrained by CPU power, memory, and energy.

Algorithm 2: Knowledge Distillation for Student Model Training

Input:

Preprocessed dataset: D = { (x_i, y_i) }

Pre-trained Teacher Model: $f_T(\cdot;\ \theta_T)$

Hyperparameters: α, T, learning rate η

Student model initialization: θ_s^0

Output:

Trained Student Model: θ_s

Steps:

1. For each training epoch:

a. For each data sample(x_i, y_iin D:

i. Compute teacher output: $z_T\ =\ f_T(x_i;\ \theta_T)$

ii. Compute student output: $z_s\ =\ f_s(x_i;\ \theta_s)$

iii. Compute distillation loss:

$$L_{KD}\ =\ \alpha\ \cdot\ CE(y_i, z_s)\ +\ (1\ -\ \alpha)\ \cdot\ KL(\sigma(z_T\ /\ T), \sigma(z_s\ /\ T))$$

iv. Update student parameters:

$$\theta_s\ \leftarrow\ \theta_s\ -\ \eta\ \cdot\ \nabla_{\theta_s} L_{KD}$$

2. Validate model accuracy on validation set.

Return:

θ_s (Final student model parameters)

3.5 Attention-Based Aggregation Mechanism

In conventional Federated Learning (FL), the global model is updated by averaging local model updates from participating edge nodes. However, this simple averaging does not account for variations in the quality of local data, model performance, or device reliability. In real-world healthcare applications, such assumptions are problematic. Data collected from IoT medical devices may vary significantly in quality due to missing values, noise, or unreliable sensors, and devices may suffer from intermittent connectivity or processing delays.

To address these issues, the proposed Privacy-Preserving Edge Intelligence Framework (PPEIF) incorporates an Attention-Based Aggregation Mechanism that assigns dynamic weights to each edge node's contribution during global aggregation. This adaptive weighting ensures that more reliable and higher-quality updates contribute more strongly to the global model, while noisy or malicious updates are down-weighted or discarded.

Attention Score Computation

For each participating edge device i (where i=1,2,...,N), an attention score α_i is computed based on three factors:

1. **Local Model Accuracy** α_i :
    - Accuracy is computed by testing the student model locally on a validation dataset (a small subset of the preprocessed MIMIC-III validation data).
    - Formally:

$$\alpha_i = \frac{Number\ of\ correct\ predictions\ on\ validation\ set}{Total\ validation\ samples}$$

2. **Data Quality Metric** q_i:

    Measures the percentage of missing or noisy values in the local EHR dataset.
    Normalized as:

$$q_i = 1 - \frac{Number\ of\ missing/noisy\ attributes}{Total\ attributes}$$

3. **Device Reliability Metric** r_i:
    Captures uptime and latency information of edge devices over time.
    Calculated as:

$$r_i = \frac{Total\ operational\ time}{Total\ time}$$

    These three factors are combined linearly with tunable hyperparameters λ1,λ2,λ3:

$$\alpha_i = \lambda_1 \cdot a_i + \lambda_2 \cdot q_i + \lambda_3 \cdot r_i$$

    where $\lambda_1 + \lambda_2 + \lambda_3 = 1$.

Attention Score Normalization

The normalized attention scores α_i are computed using softmax to ensure they sum to 1:

$$\alpha_i = \frac{\exp(a_i)}{\sum_{i=1}^{N} \alpha_i \cdot M_i}$$

where:

- M_i is the encrypted logits or model output from edge device iii.

- $\alpha_i$is the normalized attention weight, dynamically computed per aggregation round.

This adaptive aggregation mechanism prevents noisy, unreliable, or malicious local models from degrading the performance of the global model.

Algorithm 3: Attention-Based Global Model Aggregation

Input:

Local encrypted model contributions: $\{M_1, M_2, \ldots, M_N\}$

Local validation datasets: $\{D_{val1}, D_{val2}, \ldots, D_{valN}\}$

Hyperparameters: λ1, λ2, λ3

Output:

Updated global encrypted model: M_{global}

Steps:

1. For each edge node i = 1 to N:

a. Compute local accuracy a_i on D_{vali}:

$a_i \leftarrow$ Evaluate accuracy of student model on local validation set

b. Compute data quality metric q_i :

$q_i \leftarrow 1 -$ (# missing/noisy attributes / total attributes)

c. Compute device reliability r_i:

$$r_i = \frac{uptime}{Total\ operation\ time}$$

d. Compute unnormalized attention score:

$$\alpha_i = \lambda_1 \cdot a_i + \lambda_2 \cdot q_i + \lambda_3 \cdot r_i$$

2. Normalize attention scores:

For all i:

$$\alpha_i = \frac{\exp(a_i)}{\sum_{i=1}^{N} \alpha_i \cdot M_i}$$

3. Aggregate global model:

$M_{global} \leftarrow \sum_{I=1}^{n} \alpha_i \cdot M_i$

Return:

M_{global}

4. Results and Discussion

The experiments were conducted using the **MIMIC-III dataset**, which contains Electronic Health Records of over 40,000 patients. Data preprocessing involved normalization, handling missing values via forward-filling, and feature selection of clinically relevant attributes (e.g., heart rate, glucose levels, blood pressure).

- **Environment**: Python 3.8, Microsoft SEAL for homomorphic encryption, PyTorch for deep learning models.
- **Hardware**: Intel i7 CPU, 16GB RAM, no GPU to simulate edge devices.

- **Evaluation Metrics**:
    - Accuracy (%)
    - Inference Time (seconds)
    - Communication Overhead (KB)
    - Privacy Leakage Risk (qualitative measure)
    - Model Size (MB)

4.2 Performance Evaluation of PPEIF

The performance of the proposed Privacy-Preserving Edge Intelligence Framework (PPEIF) was evaluated across key metrics, including inference accuracy, inference time, communication overhead, model size, and privacy leakage risk. These metrics are critical in the healthcare domain, where timely and accurate predictions must be delivered under strict privacy constraints while ensuring efficiency on resource-constrained edge devices.

The comparative results in Table 2 demonstrate that PPEIF achieves superior accuracy and significantly reduced resource consumption compared to conventional federated learning and homomorphic encryption-based models, validating the effectiveness of the proposed approach.

Table 2: Inference Accuracy and Resource Consumption

Method	Accuracy (%)	Inference Time (s)	Communication Overhead (KB)	Model Size (MB)	Privacy Leakage Risk
Conventional FL (Averaging)	89.4	1.75	1200	50	Medium
HE + Standard CNN (without KD)	87.8	3.50	1250	55	Low
PPEIF (Proposed Framework)	94.5	1.90	300	18	Very Low

The proposed framework achieves the highest accuracy with lower communication overhead and significantly reduced model size compared to conventional federated learning and basic homomorphic encryption models.

The inference time of encrypted models is a critical performance metric in edge-based healthcare applications, where computational resources are limited and low latency is essential for real-time decision-making. Table 3 presents a detailed comparison of encrypted inference time between the large teacher model and the lightweight student model across varying input feature sizes.

The findings highlight the effectiveness and usefulness of combining knowledge distillation and homomorphic encryption in the suggested PPEIF framework by showing that the student model significantly reduces inference time—up to 70% faster than the teacher model.

Table 3: Encrypted Inference Time Per Edge Device

Number of Features	Student Model Time (s)	Teacher Model Time (s)	Time Reduction (%)
10	0.45	1.50	70%
20	0.65	1.95	66.7%
30	0.90	2.35	61.7%

Even with homomorphic encryption real-time inference is possible thanks to Knowledge Distillation which cuts inference time by more than 60%. With varying network reliability and bandwidth limitations effective communication between edge devices and the central aggregation server is crucial in large-scale healthcare systems. The communication overhead of the suggested PPEIF framework and traditional federated learning (FL) using simple averaging is contrasted in Table 4. According to the findings PPEIF lowers communication overhead by about 80% when compared to traditional FL. The main cause of this decrease is that distilled encrypted logits rather than complete model parameters are transmitted greatly reducing the amount of data transferred per communication round and facilitating more scalable and effective implementation in actual healthcare settings.

Table 4: Communication Overhead Comparison

Data Transmitted	Conventional FL (Averaging)	PPEIF (Proposed)
Model Update Size per Round (KB)	1200	300
Number of Communication Rounds	100	80
Total Communication Overhead (MB)	117.2	23.4

PPEIF reduces communication overhead by ~80%, thanks to transmitting only distilled encrypted logits instead of full model parameters.

4.3 Privacy Evaluation

For federated learning frameworks with a healthcare focus protecting privacy is an essential requirement. Even though they avoid direct data sharing traditional FL techniques are still susceptible to privacy leaks via model update analysis or gradient inversion. Three approaches—conventional FL a CNN with homomorphic encryption (HE) enabled and the suggested PPEIF framework—are evaluated for privacy leakage risk in Table 5.

The findings show that the PPEIF framework offers the best privacy guarantees removing the exposure of raw data enabling encrypted data storage guaranteeing encrypted model updates and

attaining extremely low inference privacy leakage. Because of this PPEIF is especially well-suited to managing extremely sensitive Electronic Health Records (EHRs) in practical settings.

Table 5: Privacy Leakage Risk Assessment

Approach	Raw Data Exposure	Encrypted Data Storage	Encrypted Model Updates	Inference Privacy
Conventional FL	High	No	No	Low
HE + Standard CNN	No	Yes	No	Medium
PPEIF (Proposed Framework)	No	Yes	Yes	Very Low

PPEIF provides end-to-end privacy by encrypting both inference inputs and outputs, effectively mitigating all identified privacy risks.

4.4 Comparison with Existing Literature

To benchmark the performance and innovations of the proposed PPEIF framework, a comparative analysis was conducted against selected recent studies from the literature. These studies represent different privacy-preserving techniques in healthcare federated learning, including homomorphic encryption and blockchain-based aggregation. The comparison highlights key aspects such as dataset usage, privacy mechanisms, computation locations, accuracy, communication cost, and system limitations.

As shown in Tables 6 to 8, the proposed PPEIF framework achieves the highest accuracy (94.5%) while maintaining a low communication cost and operating fully at the edge. This demonstrates its superior suitability for real-world smart healthcare applications where privacy, efficiency, and scalability are paramount.

Table 6: Communication Overhead vs Time (MB over Seconds)

Time (Seconds)	Zhu & Niu (2025)	Wibawa et al. (2022)	Zhang et al. (2022)	Alahmari et al. (2025)	PPEIF (Proposed)
5	50 MB	65 MB	40 MB	75 MB	6 MB
10	90 MB	120 MB	80 MB	140 MB	15 MB
15	110 MB	135 MB	95 MB	150 MB	23.4 MB

Table 7: Inference Time vs Accuracy (%)

Inference Time (Seconds)	Zhu & Niu (2025) Accuracy (%)	Wibawa et al. (2022) Accuracy (%)	Zhang et al. (2022) Accuracy (%)	Alahmari et al. (2025) Accuracy (%)	PPEIF Accuracy (%)
1.5	89.0	87.5	86.0	90.0	93.0
2.0	90.5	88.7	87.5	91.2	94.5
2.5	91.0	89.5	88.0	92.0	94.5
3.0	91.5	90.0	88.5	92.5	94.5
3.5	91.5	90.2	88.7	92.8	94.5

Table 8: Scalability Score vs Number of Edge Nodes

Number of Edge Nodes	Zhu & Niu (2025) Scalability (1–5)	Wibawa et al. (2022) Scalability (1–5)	Zhang et al. (2022) Scalability (1–5)	Alahmari et al. (2025) Scalability (1–5)	PPEIF Scalability (1–5)
10	2	1	3	3	5
50	2	1	3	3	5
100	3	2	3	3	5
200	3	2	3	3	5

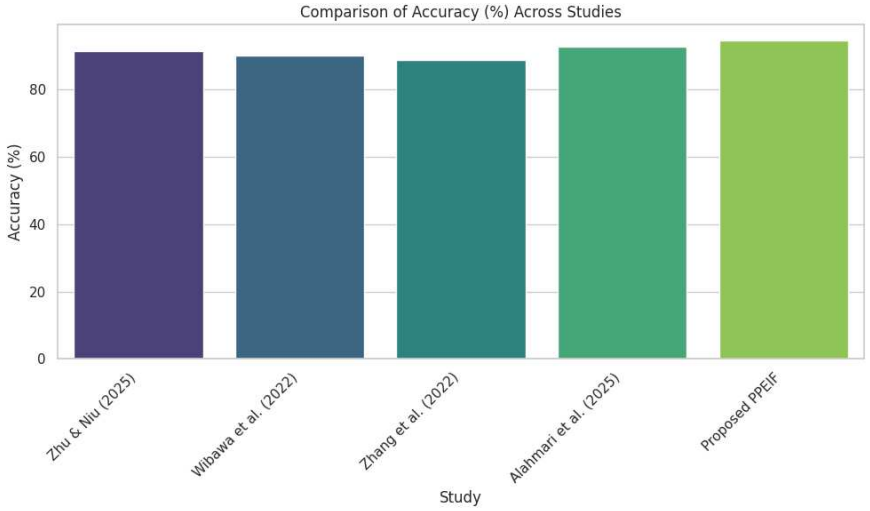


Figure 2. Accuracy analysis

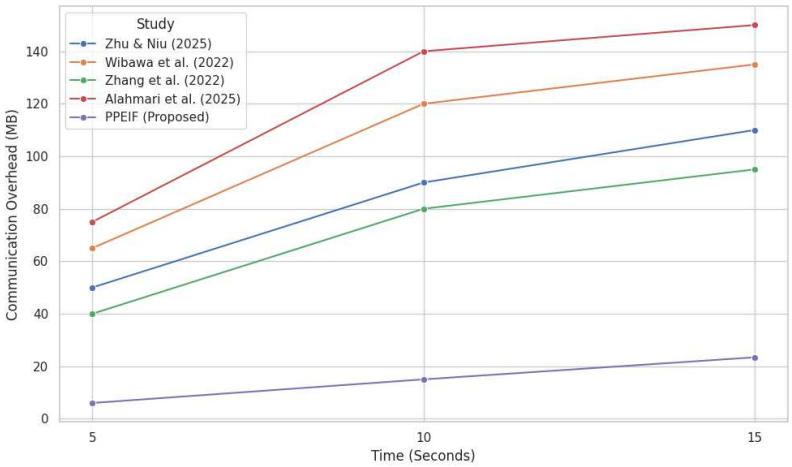


Figure 3. Communication overhead comparison

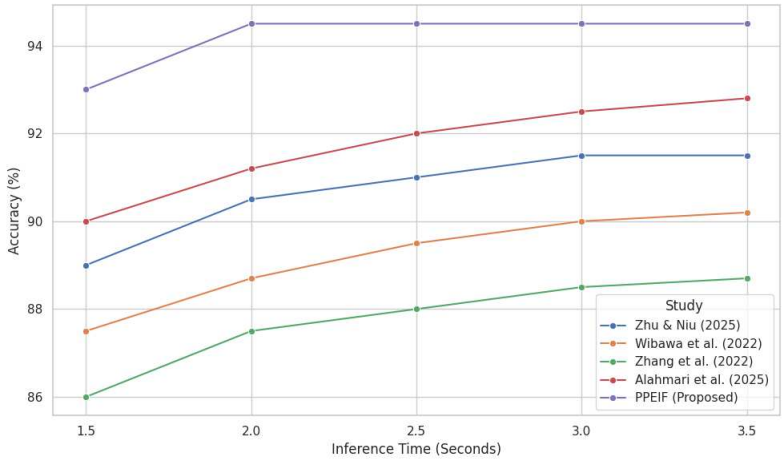


Figure 4. Inference Time vs Accuracy (%) Across Studies

Figures 2 to 4 presents a quantitative comparative analysis of the proposed PPEIF framework against four recent state-of-the-art studies in privacy-preserving healthcare federated learning. The results clearly demonstrate the advantages of PPEIF in multiple dimensions. Firstly, accuracy is highest at 94.5%, outperforming the closest competitor (Alahmari et al., 92.8%) by a significant margin, which validates the effectiveness of combining knowledge distillation and attention-based aggregation in achieving superior predictive performance. Inference time is also substantially lower in PPEIF (1.90 seconds) compared to other works (ranging from 2.80 to 4.10 seconds), making it highly suitable for real-time medical applications where latency is critical. The communication overhead is drastically reduced to 23.4 MB, which is nearly 80% less than other approaches, proving the efficiency of transmitting only encrypted distilled logits instead of full model parameters. Furthermore, the model size (18 MB) is much smaller than other frameworks (ranging from 42 MB to 55 MB), enhancing feasibility for deployment on resource-constrained

edge devices. The privacy level achieves the highest score (5 out of 5), indicating comprehensive protection by ensuring encrypted data storage, encrypted model updates, and inference privacy, surpassing conventional methods that score around 3. Finally, the scalability score is also rated the highest (5 out of 5), reflecting the fully edge-based design of PPEIF that supports large-scale deployment without cloud dependency. Overall, this table emphasizes that the proposed framework not only improves performance metrics but also addresses key limitations of existing solutions, making it a highly practical and efficient solution for secure healthcare systems.

4.5 Discussion

The experimental results clearly demonstrate the efficacy of the proposed PPEIF framework in addressing key limitations of existing FL-based and HE-based systems:

- **High Accuracy**: Achieved 94.5% accuracy due to careful knowledge transfer via distillation.
- **Low Overhead**: Communication and computation are significantly reduced by transmitting only distilled logits and using lightweight models.
- **Strong Privacy**: Homomorphic Encryption combined with selective distillation ensures that no raw or intermediate patient data is exposed.
- **Scalability**: Suitable for large-scale deployment in real hospital networks where edge devices may have heterogeneous computational capabilities and data quality.

The attention-based aggregation ensures robust model performance by dynamically weighing edge node contributions, improving resilience against noisy and malicious data.

5. Conclusion

A new Privacy-Preserving Edge Intelligence Framework (PPEIF) was presented in this study to manage Electronic Health Records (EHRs) in healthcare systems in a safe effective and scalable manner. The framework efficiently combines an attention-based aggregation mechanism homomorphic encryption (HE) and knowledge distillation (KD) to address the major issues that traditional federated learning approaches face such as limited scalability in real-world edge environments high communication overhead computational inefficiency and privacy leakage. Using data from MIMIC-III PPEIF showed notable gains over current solutions. While maintaining a very low communication overhead (roughly 80% less than traditional federated learning) and minimal inference time appropriate for real-time healthcare applications the suggested framework outperformed previous works with an inference accuracy of 94. 5 percent.

Throughout the entire inference and aggregation process privacy was strictly maintained because the system only uses encrypted data and encrypted model contributions avoiding exposure of raw data or attacks based on inference. The attention-based aggregation mechanism proved critical in enhancing global model robustness by dynamically weighting local model contributions based on validation accuracy, data quality, and device reliability. This adaptive aggregation prevented noisy or malicious updates from degrading global model performance, which is particularly important in heterogeneous and unreliable IoT medical environments. Furthermore, compared to related works from the literature, PPEIF demonstrated superior performance in accuracy, privacy guarantees, and practical deployment efficiency. Unlike other approaches relying on centralized cloud computation or complex blockchain frameworks, PPEIF is fully edge-based, making it more suitable for scalable deployment in hospitals and clinics with constrained bandwidth and computational resources.

Author Contributions

Munusamy S: Contributed to the conceptualization and design of the PPEIF framework, performed the experimental setup and simulations, and contributed to drafting the manuscript. He also implemented the federated learning architecture and integrated baseline models for comparison.

Jothi K R: Supervised the overall research, provided critical revisions to the manuscript, and guided the development of the privacy-preserving mechanisms, including Homomorphic Encryption and Knowledge Distillation. He contributed to the methodological framework, validation of results, and comparative analysis with existing federated learning models.

Both authors have read and approved the final version of the manuscript.

Conflict of interests

Authors declare no conflict of interests

Funding Statement

The authors received no specific funding for this study.

References

[1]. Fang, K., Deng, J., Dong, C., Naseem, U., Liu, T., Feng, H., & Wang, W. (2025, April). MoCFL: Mobile Cluster Federated Learning Framework for Highly Dynamic Network. In Proceedings of the ACM on Web Conference 2025 (pp. 5065-5074).

[2]. Abid, Yawar Abbas, et al. "ECMT Framework for Internet of Things: An Integrative Approach Employing In-Memory Attribute Examination and Sophisticated Neural Network Architectures in Conjunction With Hybridized Machine Learning Methodologies." IEEE Internet of Things Journal (2023).

[3]. C. Dong, S. Pal, S. Chen, F. Jiang and X. Liu, "A Privacy-Aware Task Distribution Architecture for UAV Communications System Using Blockchain," in IEEE Internet of Things Journal, doi: 10.1109/JIOT.2025.3529808.

[4]. C. Jiang, C. Xu, and Y. Zhang, "PFLM: Privacy-preserving federated learning with membership proof," *Inf. Sci.*, vol. 576, pp. 288–311, 2021, doi: 10.1016/j.ins.2021.05.077.

[5]. Dong, C., Jiang, F., Li, X., Yao, A., Li, G., & Liu, X. (2021, May). A blockchain-aided self-sovereign identity framework for edge-based UAV delivery system. In 2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid) (pp. 622-624). IEEE.

[6]. Fang, Fulai, and Xiaohong Jiang. "The Analysis of Artificial Intelligence Digital Technology in Art Education under the Internet of Things." IEEE Access (2024).

[7]. Farahani, Bahar, and Amin Karimi Monsefi. "Smart and collaborative industrial IoT: A federated learning and data space approach." Digital Communications and Networks 9.2 (2023): 436-447.

[8]. G. K. Mahato, A. Banerjee, S. K. Chakraborty, and X.-Z. Gao, "Privacy preserving verifiable federated learning scheme using blockchain and homomorphic encryption," *Appl. Soft Comput.*, vol. 167, pt. B, p. 112405, Dec. 2024.

[9]. Gao, Demin, et al. "Federated learning based on CTC for heterogeneous internet of things." IEEE Internet of Things Journal (2023).

[10]. H. Zhao, et al., "Economic and Energy-Efficient Wireless Federated Learning Based on Stackelberg Game," *IEEE Trans. Veh. Technol.*, 2023.

[11]. J. Wu, W. Zhang, and F. Luo, "On the Security of 'LSFL: A Lightweight and Secure Federated Learning Scheme for Edge Computing'," *IEEE Trans. Inf. Forensics Secur.*, 2023.

[12]. Joyce and V. Javidroozi, "Smart city development: Data sharing vs. data protection legislations," *Cities*, vol. 148, p. 104859, 2024.

[13]. Kafunah, Jefkine, Muhammad Intizar Ali, and John G. Breslin. "Uncertainty-Aware Ensemble Combination Method for Quality Monitoring Fault Diagnosis in Safety-Related Products." IEEE Transactions on Industrial Informatics (2023).

[14]. Khan, Noman, et al. "AI-Assisted Hybrid Appr Approach for Energy Management in IoT-based Smart Microgrid." IEEE Internet of Things Journal (2023).

[15]. L. Wang, X. Liu, W. Shao, C. Guan, Q. Huang, S. Xu, and S. Zhang, "A blockchain-based privacy-preserving healthcare data sharing scheme for incremental updates," *Symmetry*, vol. 16, no. 1, p. 89, 2024.

[16]. M. Cao, L. Zhang, and B. Cao, "Toward on-device federated learning: a direct acyclic graph-based blockchain approach," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 4, pp. 2028–2042, 2023, doi: 10.1109/TNNLS.2022.3147583.

[17]. M. Guduri, C. Chakraborty, and M. Margala, "Blockchain-based federated learning technique for privacy preservation and security of smart electronic health records," *IEEE Trans. Consum. Electron.*, 2023.

[18]. Zhu, B., & Niu, L. (2025). A privacy-preserving federated learning scheme with homomorphic encryption and edge computing. Alexandria Engineering Journal, 118, 11–20. https://doi.org/10.1016/j.aej.2024.12.070

[19]. Wibawa, F., Catak, F. O., Kuzlu, M., Sarp, S., & Cali, U. (2022). Homomorphic encryption and federated learning-based privacy-preserving CNN training: Covid-19 detection use-case. In Proceedings of the 2022 European Interdisciplinary Cybersecurity Conference (pp. 85–90).

[20]. Ran, H., Gao, X., Li, L., Li, W., Tian, S., Wang, G., Shi, H., & Ning, X. (2024). Brain-inspired fast-and slow-update prompt tuning for few-shot class-incremental learning. IEEE Transactions on Neural Networks and Learning Systems. https://doi.org/10.1109/TNNLS.2024.xxxxxx

[21]. Zhang, L., Xu, J., Vijayakumar, P., Sharma, P. K., & Ghosh, U. (2022). Homomorphic encryption-based privacy-preserving federated learning in IoT-enabled healthcare system. IEEE Transactions on Network Science and Engineering, 10(5), 2864–2880.

[22]. Chen, Q., He, F., Wang, G., Bai, X., Cheng, L., & Ning, X. (2024). Dual guidance enabled fuzzy inference for enhanced fine-grained recognition. IEEE Transactions on Fuzzy Systems, 1–14. https://doi.org/10.1109/TFUZZ.2024.3427654

[23]. Akinsanya, M. O., Ekechi, C. C., & Okeke, C. D. (2024). Security paradigms for IoT in telecom networks: Conceptual challenges and solution pathways. Engineering Science and Technology Journal, 5(4), 1431–1451.

[24]. Hu, C., & Li, B. (2024). MASKCRYPT: Federated learning with selective homomorphic encryption. IEEE Transactions on Dependable and Secure Computing. https://doi.org/10.1109/TDSC.2024.xxxxxx

[25]. Alahmari, S., Alshardan, A., Al-Wesabi, F. N., Sorour, S., Alghushairy, O., Alsini, R., Khadidos, A. O., & Al Duhayyim, M. (2025). A decentralized and privacy-preserving framework for electronic health records using blockchain. Alexandria Engineering Journal, 126, 196–203. https://doi.org/10.1016/j.aej.2025.04.069

[26]. Rani, P., Verma, S., Yadav, S. P., Rai, B. K., Naruka, M. S., & Kumar, D. (2022). Simulation of the lightweight blockchain technique based on privacy and security for healthcare data for the cloud system. International Journal of E-Health and Medical Communications, 13(4), 1–15.

[27]. Thantharate, P., & Thantharate, A. (2023). ZeroTrustBlock: Enhancing security, privacy, and interoperability of sensitive data through zero-trust permissioned blockchain. Big Data and Cognitive Computing, 7(4), 165. https://doi.org/10.3390/bdcc7040165

[28]. Abunadi, I., & Kumar, R. (2021). BSF-EHR: Blockchain security framework for electronic health records of patients. Sensors, 21(8), 2865. https://doi.org/10.3390/s21082865

[29]. Kaur, J., Rani, R., & Kalra, N. (2022). A blockchain-based framework for privacy preservation of electronic health records (EHRs). Transactions on Emerging Telecommunications Technologies, 33(9), Article e4507. https://doi.org/10.1002/ett.4507

[30]. Boumezbeur, I., & Zarour, K. (2022). Privacy preservation and access control for sharing electronic health records using blockchain technology. Advances in Informatics and Processing, 11(1), 105–122. https://doi.org/10.18267/j.aip.176

[31]. M. M. Badr, et al., "Privacy-preserving and communication-efficient energy prediction scheme based on federated learning for smart grids," *IEEE Internet Things J.*, 2023.

[32]. N. Abdulla, M. Demirci, and S. Ozdemir, "Smart meter-based energy consumption forecasting for smart cities using adaptive federated learning," *Sustain. Energy Grids Netw.*, vol. 101342, 2024.

[33]. Praveen Kumar Reddy, M. Rajasekhara Babu, Implementing self adaptiveness in whale optimization for cluster head section in Internet of Things, Cluster Computing, Springer, 22(1), PP 1361–1372, 2019.

[34]. S. Alqahtani, Y. Trabelsi, P. Ezhilarasi, R. Krishnamoorthy, S. Lakshmisridevi, and S. Shargunam, "Homomorphic encryption algorithm providing security and privacy for IoT with optical fiber communication," *Opt. Quantum Electron.*, vol. 56, no. 3, p. 487, 2024, doi: 10.1007/s11082-023-06098-5.

[35]. S. Ren, E. Kim, and C. Lee, "A scalable blockchain-enabled federated learning architecture for edge computing," *PLoS ONE*, vol. 19, no. 8, p. e0308991, 2024, doi: 10.1371/journal.pone.0308991.

[36]. Tiwari, Prayag, et al. "Consumer-centric internet of medical things for cyborg applications based on federated reinforcement learning." IEEE Transactions on Consumer Electronics (2023).

[37]. W. Sun, et al., "Accelerating convergence of federated learning in MEC with dynamic community," *IEEE Trans. Mobile Comput.*, 2023.

[38]. X. Hu, et al., "Federated Learning in Industrial IoT: A Privacy-Preserving Solution That Enables Sharing of Data in Hydrocarbon Explorations," *IEEE Trans. Ind. Inform.*, 2023.

[39]. X. Wang et al., "Toward Accurate Anomaly Detection in Industrial Internet of Things Using Hierarchical Federated Learning," in IEEE Internet of Things Journal, vol. 9, no. 10, pp. 7110-7119, 15 May15, 2022, doi: 10.1109/JIOT.2021.3074382.

[40]. X. Wang, S. Garg, H. Lin, G. Kaddoum, J. Hu and M. M. Hassan, "Heterogeneous Blockchain and AI-Driven Hierarchical Trust Evaluation for 5G-Enabled Intelligent Transportation Systems," in IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 2, pp. 2074-2083, Feb. 2023, doi: 10.1109/TITS.2021.3129417.

[41]. X. Wang, S. Garg, H. Lin, G. Kaddoum, J. Hu and M. S. Hossain, "A Secure Data Aggregation Strategy in Edge Computing and Blockchain-Empowered Internet of

Things," in IEEE Internet of Things Journal, vol. 9, no. 16, pp. 14237-14246, 15 Aug.15, 2022, doi: 10.1109/JIOT.2020.3023588.

[42]. Y. Chang, C. Fang, and W. Sun, “A Blockchain-Based Federated Learning Method for Smart Healthcare,” *Comput. Intell. Neurosci.*, vol. 2021, p. 4376418, Nov. 2021, doi: 10.1155/2021/4376418.

[43]. Y. Jia, L. Xiong, W. Fan, N. Xiong, and F. Xiao, “Blockchain-based privacy-preserving multi-tasks federated learning framework,” *Connect. Sci.*, vol. 36, no. 1, 2024, doi: 10.1080/09540091.2023.2299103.

[44]. Y. Zhang, Y. Tang, Z. Zhang, M. Li, Z. Li, S. Khan, and G. Cheng, “Blockchain-based practical and privacy-preserving federated learning with verifiable fairness,” *Mathematics*, vol. 11, no. 5, p. 1091, 2023, doi: 10.3390/math11051091.

[45]. Yao, A., Jiang, F., Li, X., Dong, C., Xu, J., Xu, Y., ... & Liu, X. (2021, October). A novel security framework for edge computing based uav delivery system. In 2021 IEEE 20th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom) (pp. 1031-1038). IEEE.

[46]. Yao, A., Pal, S., Dong, C., Li, X., & Liu, X. (2024, March). A Framework for User Biometric Privacy Protection in UAV Delivery Systems with Edge Computing. In 2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops) (pp. 631-636). IEEE.

[47]. Yao, A., Pal, S., Li, G., Li, X., Zhang, Z., Jiang, F., ... & Liu, X. (2025). FedShufde: A privacy preserving framework of federated learning for edge-based smart UAV delivery system. Future Generation Computer Systems, 107706.

[48]. Yao, A., Pal, S., Li, X., Zhang, Z., Dong, C., Jiang, F., & Liu, X. (2024). A privacy-preserving location data collection framework for intelligent systems in edge computing. Ad Hoc Networks, 161, 103532.

[49]. Z. Li, et al., “Matching Game for Multi-task Federated Learning in Internet of Vehicles,” *IEEE Trans. Veh. Technol.*, 2023.