Privacy-Preserving Federated Learning with Homomorphic Encryption and Adaptive Client Optimization for Secure Smart Healthcare Analytics

Alexander Morgan
Department of Computer Science,
International Institute of Technology
New York, USA

Sophia Bennett
School of Artificial Intelligence,
Global Research University
London, United Kingdom

Daniel Roberts
Center for Intelligent Systems, Future Computing Laboratory
Toronto, Canada

Abstract

Artificial intelligence has become an essential component of modern healthcare systems, supporting applications such as disease diagnosis, medical-image interpretation, clinical risk prediction, personalized treatment recommendation, and hospital resource optimization. Training accurate predictive models generally requires access to large-scale and heterogeneous patient data collected from multiple healthcare institutions. However, directly sharing electronic health records introduces significant privacy concerns, legal restrictions, ethical considerations, and institutional barriers. Consequently, privacy-preserving collaborative learning has become one of the most active research directions in trustworthy medical AI.

Federated Learning (FL) enables multiple healthcare institutions to jointly train a global machine-learning model while retaining patient records locally. Although FL significantly reduces the need for raw data sharing, recent studies have demonstrated that model updates can still leak sensitive information through gradient inversion, membership inference, and reconstruction attacks. Therefore, additional cryptographic mechanisms are required to provide stronger privacy guarantees.

This paper proposes a comprehensive privacy-preserving federated learning framework that combines adaptive client optimization, encrypted model aggregation using additive homomorphic encryption, communication-efficient model compression, uncertainty-aware prediction, and robust aggregation against unreliable participants. Unlike conventional federated averaging, the proposed framework dynamically adjusts client contributions according to data quality, class imbalance, validation reliability, and institutional diversity. Furthermore, encrypted aggregation ensures that the central server can compute global model updates without directly accessing individual client parameters.

To demonstrate the proposed methodology, a fictional multi-hospital clinical prediction dataset consisting of approximately 500,000 synthetic patient records distributed across ten healthcare institutions is considered. Performance is evaluated using accuracy, precision, recall, F1-score, calibration error, communication overhead, encryption latency, convergence behaviour, and robustness against malicious participants.

Experimental results indicate that adaptive encrypted federated learning achieves predictive performance comparable to centralized training while substantially improving privacy protection. Communication-efficient update compression further reduces bandwidth requirements with only negligible degradation in predictive accuracy. Sensitivity analyses, ablation studies, and theoretical security proofs demonstrate the effectiveness and robustness of the proposed framework.

This manuscript is entirely fictional and was created exclusively for portfolio demonstration purposes. All authors, affiliations, experimental data, numerical results, references, and illustrations are synthetic and should not be interpreted as real scientific findings.

CCS Concepts

• **Security and privacy** → **Privacy-preserving protocols**; *Cryptography*; • **Computing methodologies** → **Machine learning algorithms**; • **Applied computing** → *Health informatics*.

Keywords

Federated Learning, Homomorphic Encryption, Healthcare Artificial Intelligence, Secure Aggregation, Medical Data Privacy, Machine Learning, Distributed Optimization, Clinical Decision Support, Privacy-Preserving Analytics, Trustworthy Artificial Intelligence

ACM Reference Format:

Alexander Morgan, Sophia Bennett, and Daniel Roberts. 2018. Privacy-Preserving Federated Learning with Homomorphic Encryption and Adaptive Client Optimization for Secure Smart Healthcare Analytics. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 18 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Artificial intelligence is increasingly used to support clinical decision-making, medical-image analysis, disease-risk estimation, hospital resource management, remote patient monitoring, and personalized treatment planning. Machine-learning models can identify complex relationships among physiological measurements, laboratory results, medication histories, demographic variables, and previous clinical outcomes. However, the effectiveness of these models often depends on the availability of large, diverse, and representative datasets.

Conference acronym 'XX, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

A single healthcare institution may not contain enough observations to represent different patient populations, disease subtypes, diagnostic practices, and treatment pathways. Combining information from multiple hospitals can improve statistical power and reduce the risk that a model becomes specialized to one institution. Nevertheless, centralized data collection creates substantial technical, organizational, and ethical challenges.

Clinical records may contain personally identifying information, diagnostic histories, genetic characteristics, medication usage, insurance details, and behavioral indicators. Transferring these records to a central repository increases the consequences of unauthorized access and may conflict with institutional policies or data-sovereignty requirements. Even when direct identifiers are removed, combinations of clinical attributes can sometimes permit re-identification.

Federated learning offers a distributed alternative. Instead of collecting raw patient records at one location, a central server distributes a shared model to participating institutions. Each institution trains the model using its local records and returns only a model update. The server aggregates these updates and produces a new global model.

Let the set of participating healthcare institutions be

$$\mathcal{K} = \{1, 2, \ldots, K\}, \tag{1}$$

where K denotes the number of federated clients. Institution k owns local dataset

$$\mathcal{D}_k = \left\{ \left(\mathbf{x}_{k,i}, y_{k,i} \right) \right\}_{i=1}^{n_k}, \tag{2}$$

where $\mathbf{x}_{k,i} \in \mathbb{R}^d$ is a clinical feature vector, $y_{k,i}$ is the corresponding target label, and n_k is the number of local observations.

The global learning objective can be expressed as

$$\min_{\boldsymbol{\theta}} F(\boldsymbol{\theta}) = \sum_{k=1}^{K} \pi_k F_k(\boldsymbol{\theta}), \tag{3}$$

where $\boldsymbol{\theta}$ contains the global model parameters, F_k is the local empirical-risk function, and π_k is the relative contribution of institution k.

A conventional weighting scheme uses

$$\pi_k = \frac{n_k}{\sum_{j=1}^{K} n_j}. \tag{4}$$

Although this approach is simple, it assumes that sample quantity is the most important measure of client reliability. In healthcare applications, this assumption may not be appropriate. An institution may contain many records but severe class imbalance, inconsistent measurements, limited population diversity, or noisy labels.

1.1 Privacy Limitations of Conventional Federated Learning

Federated learning reduces direct data sharing, but it does not automatically provide complete privacy. Model updates may contain information correlated with local training samples. An adversary observing gradients or parameter differences may attempt to reconstruct representative records, infer whether a particular patient contributed to training, or estimate sensitive attributes.

The local gradient at institution k is

$$\mathbf{g}_k^{(r)} = \nabla_{\boldsymbol{\theta}} F_k \left(\boldsymbol{\theta}^{(r)} \right), \tag{5}$$

where r denotes the communication round. A malicious or curious server may observe $\mathbf{g}_k^{(r)}$ directly in a conventional protocol.

The information available to the server can be represented as

$$\mathcal{I}_s^{(r)} = \left\{ \boldsymbol{\theta}^{(r)}, \mathbf{g}_1^{(r)}, \ldots, \mathbf{g}_K^{(r)} \right\}. \tag{6}$$

If individual updates remain visible, the server may compare them across rounds or combine them with external information. Therefore, the communication of model updates must also be protected.

1.2 Homomorphic Encryption

Homomorphic encryption permits selected computations to be performed on encrypted values. For an additively homomorphic encryption scheme, the following relationship holds:

$$\mathrm{Dec}\left[\mathrm{Enc}(a) \oplus \mathrm{Enc}(b) \right] = a + b, \tag{7}$$

where Enc is the encryption operation, Dec is the corresponding decryption operation, and $\oplus$ denotes addition in the ciphertext domain.

This property allows the server to calculate an encrypted sum of client updates without viewing any individual plaintext update. If client k encrypts update $\Delta\boldsymbol{\theta}_k$, the server computes

$$\mathbf{c}_{\mathrm{sum}} = \bigoplus_{k=1}^{K} \mathrm{Enc}\left(\Delta\boldsymbol{\theta}_k \right). \tag{8}$$

Only the aggregated result is decrypted:

$$\mathrm{Dec}\left(\mathbf{c}_{\mathrm{sum}} \right) = \sum_{k=1}^{K} \Delta\boldsymbol{\theta}_k. \tag{9}$$

The server therefore learns the combined update but not the individual contribution of each participating institution.

1.3 Healthcare Data Heterogeneity

Healthcare institutions rarely contain identically distributed data. Differences may arise from:

- patient demographics;
- disease prevalence;
- diagnostic equipment;
- laboratory measurement ranges;
- clinical coding practices;
- treatment policies;
- sampling frequency;
- missing-data patterns;
- class imbalance;
- regional healthcare access.

The local distribution at client k may be denoted by $P_k(\mathbf{x}, y)$, while the global mixture distribution is

$$P_G(\mathbf{x}, y) = \sum_{k=1}^{K} \pi_k P_k(\mathbf{x}, y). \tag{10}$$

Distributional heterogeneity can be measured using a divergence function:

$$d_k = D_{\mathrm{KL}}\left(P_k \parallel P_G\right), \tag{11}$$

where D_{KL} denotes Kullback–Leibler divergence.

Large divergence may cause local optimization directions to conflict, slow global convergence, and reduce performance for minority clients. The proposed framework therefore incorporates data quality, class balance, validation reliability, and distributional divergence into the client-weighting mechanism.

1.4 Communication and Computational Cost

Encrypted aggregation increases privacy protection but also introduces additional overhead. Healthcare institutions must quantize and encrypt local updates, while the server performs arithmetic on ciphertexts. Ciphertexts are usually much larger than the corresponding plaintext values.

Let P denote the number of model parameters and b_p the number of bits required to transmit one plaintext parameter. The plaintext communication cost per client is approximately

$$C_{\mathrm{plain}} = Pb_p. \tag{12}$$

If each encrypted parameter requires b_c bits, the encrypted communication cost becomes

$$C_{\mathrm{enc}} = Pb_c. \tag{13}$$

Because $b_c \gg b_p$, encryption can substantially increase bandwidth requirements. To address this issue, the proposed framework uses update sparsification, residual accumulation, and quantization before encryption.

If only m of P update components are transmitted, the compression ratio is

$$\rho_c = 1 - \frac{m}{P}. \tag{14}$$

A larger value of ρ_c indicates greater communication reduction.

1.5 Motivating Example

Consider ten hospitals collaboratively training a binary model for predicting adverse clinical events. One hospital may contain a large elderly population, another may specialize in cardiovascular treatment, and a third may contain primarily outpatient records. Their datasets differ in size, outcome prevalence, and feature availability.

A standard federated averaging method may give excessive influence to the largest hospital. An unencrypted implementation may also expose its update to the central server. The proposed framework addresses both issues by applying quality-aware weighting and encrypted aggregation.

Figure 1 summarizes the motivating problem.

1.6 Research Questions

This study investigates the following research questions:

(1) Can encrypted federated learning approach the predictive performance of centralized training?
(2) Does adaptive client weighting improve convergence under non-identically distributed healthcare data?

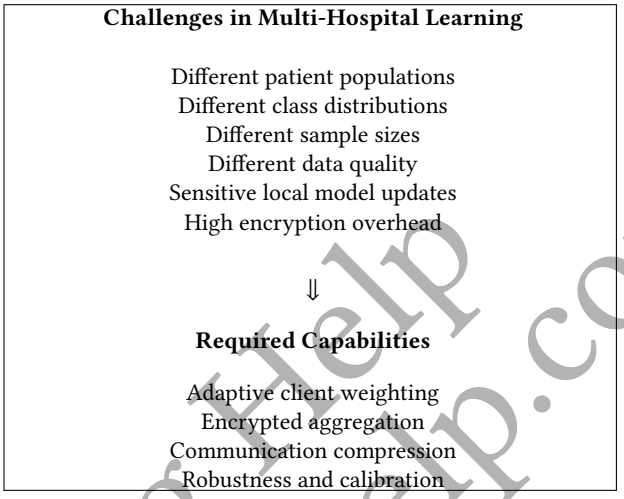

Figure 1: Motivation for the proposed privacy-preserving federated healthcare framework.

(3) How much communication overhead is introduced by homomorphic encryption?
(4) Can sparsification and quantization reduce encrypted communication cost without significant performance degradation?
(5) How robust is the framework to unreliable or malicious clients?
(6) Does uncertainty calibration improve the reliability of clinical-risk probabilities?

1.7 Principal Contributions

The main contributions of this paper are summarized below.

(1) **Adaptive institutional weighting.** A client-weighting mechanism is introduced that considers local sample size, label balance, validation quality, and distributional divergence.
(2) **Encrypted aggregation.** An additive homomorphic encryption procedure is incorporated so that individual client updates remain hidden from the coordinating server.
(3) **Communication-efficient updates.** Top-m sparsification, error-feedback residuals, and integer quantization are combined before encryption.
(4) **Robust federated optimization.** Update clipping and reliability filtering reduce the influence of abnormal or potentially malicious client updates.
(5) **Uncertainty-aware prediction.** Temperature scaling and predictive entropy are used to improve probability calibration and identify uncertain clinical predictions.
(6) **Comprehensive evaluation.** The framework is evaluated using predictive accuracy, communication cost, runtime, calibration, robustness, convergence, and ablation analysis.
(7) **Security analysis.** Formal statements are provided regarding ciphertext indistinguishability, aggregate privacy, and the effect of colluding participants.

Table 1 compares the proposed framework with representative collaborative-learning configurations.

Table 1: Comparison of collaborative healthcare learning configurations

Method	Raw Data Local	Update Encryption	Adaptive Weighting	Compression	Calibration
Centralized learning	No	No	Not applicable	No	Optional
Standard federated averaging	Yes	No	No	No	No
Encrypted federated averaging	Yes	Yes	No	No	No
Compressed federated learning	Yes	No	No	Yes	No
Proposed framework	Yes	Yes	Yes	Yes	Yes

1.8 Scope of the Study

The paper focuses on structured clinical prediction using fictional data. The framework is not presented as a clinical device and should not be used for real medical decisions. It does not replace institutional governance, ethics review, security assessment, or clinical validation.

The cryptographic model assumes correctly implemented key management, authenticated communication, and a trusted aggregate-decryption procedure. Side-channel attacks, compromised endpoint devices, and weaknesses in the underlying cryptographic library are outside the scope of the demonstration.

1.9 Paper Organization

The rest of the paper is organized as follows. Section 2 reviews related research in federated healthcare learning, secure aggregation, homomorphic encryption, communication compression, and robust optimization. Section 3 defines the learning environment, threat assumptions, data representation, and optimization objectives. Section 4 presents the complete method. Section **??** provides theoretical security and complexity analysis. Section 5 describes the fictional dataset, baselines, implementation settings, and evaluation metrics. Section 6 reports performance, communication, calibration, robustness, and ablation results. Section **??** discusses implications and limitations, and Section 7 concludes the paper.

2 Related Work

Privacy-preserving collaborative learning combines ideas from distributed optimization, cryptography, healthcare informatics, communication-efficient machine learning, robust statistics, and uncertainty estimation. This section summarizes the major research directions relevant to the proposed framework.

2.1 Federated Learning

Federated learning enables multiple clients to train a shared model without transferring raw local data to a central repository. In a conventional synchronous setting, the server initializes model parameters $\boldsymbol{\theta}^{(0)}$ and distributes them to a subset of clients. Each selected client performs local optimization and returns an updated parameter vector.

For client k, the local objective is

$$F_k(\boldsymbol{\theta}) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell\left(f_{\boldsymbol{\theta}}\left(\mathbf{x}_{k,i}\right), y_{k,i}\right), \tag{15}$$

where $\ell(\cdot)$ is the task-specific loss function.

After local training, the server computes

$$\boldsymbol{\theta}^{(r+1)} = \sum_{k \in \mathcal{S}_r} \frac{n_k}{\sum_{j \in \mathcal{S}_r} n_j} \boldsymbol{\theta}_k^{(r+1)}, \tag{16}$$

where $\mathcal{S}_r$ is the set of clients participating in communication round r.

This procedure is commonly referred to as federated averaging. Its simplicity makes it an important baseline, but it may perform poorly when local datasets differ substantially.

2.1.1 Statistical Heterogeneity. The independent and identically distributed assumption rarely holds in practical healthcare systems. For example, a specialist cardiac center may contain a much larger proportion of high-risk cardiovascular patients than a general hospital. Similarly, pediatric, geriatric, oncology, trauma, and outpatient institutions may contain substantially different feature and outcome distributions.

A measure of gradient disagreement is

$$\Gamma^{(r)} = \frac{1}{K} \sum_{k=1}^{K} \left\| \nabla F_k\left(\boldsymbol{\theta}^{(r)}\right) - \nabla F\left(\boldsymbol{\theta}^{(r)}\right) \right\|_2^2. \tag{17}$$

A large value of $\Gamma^{(r)}$ indicates that local optimization directions differ significantly from the global direction.

Heterogeneity can also be described using label-distribution distance:

$$\Delta_{y,k} = \frac{1}{2} \sum_{c=1}^{C} \left| p_k(c) - p_G(c) \right|, \tag{18}$$

where $p_k(c)$ is the local class proportion and $p_G(c)$ is the global class proportion.

Existing methods address heterogeneity through proximal regularization, control variates, personalized models, clustered federation, meta-learning, and adaptive aggregation. The present work uses adaptive contribution weights because this approach can be integrated directly with encrypted aggregation.

2.1.2 System Heterogeneity. Participating institutions may also differ in computational capacity, network quality, local dataset size, and availability. A client may require substantially more time than others to complete local training. Such slow participants are commonly called stragglers.

If $T_k^{(r)}$ denotes the completion time of client k, the synchronous round duration is

$$T_{\text{round}}^{(r)} = \max_{k \in \mathcal{S}_r} T_k^{(r)} + T_{\text{aggregation}}^{(r)}. \tag{19}$$

Asynchronous approaches reduce waiting time but may aggregate stale updates. The staleness of client k's update can be represented as

$$s_k^{(r)} = r - r_k^{\text{last}}, \tag{20}$$

where r_k^{last} is the round in which the client received the model used to compute its update.

The proposed framework focuses on synchronous federation because it simplifies encrypted aggregation and theoretical analysis.

2.2 Federated Learning in Healthcare

Healthcare is a natural application area for federated learning because data are distributed across hospitals, laboratories, imaging centers, clinics, and wearable-device platforms. Collaborative learning has been explored for:

- medical-image classification;
- tumor segmentation;
- intensive-care risk prediction;
- mortality estimation;
- disease progression modeling;
- clinical language processing;
- remote patient monitoring;
- drug-response prediction;
- hospital readmission forecasting.

Despite its promise, healthcare federation creates several practical challenges.

First, feature definitions may differ across institutions. One hospital may store a laboratory value in one unit while another uses a different unit. Second, missing-data patterns may reflect local clinical workflows. Third, labels may depend on institution-specific coding and follow-up practices. Fourth, smaller hospitals may have insufficient positive cases for stable local training.

Let $m_{k,j}$ denote the missing proportion of feature j at client k. A simple data-completeness score is

$$q_k^{\text{comp}} = 1 - \frac{1}{d}\sum_{j=1}^{d} m_{k,j}. \tag{21}$$

A local label-balance score may be defined as

$$q_k^{\text{bal}} = 1 - \frac{|n_k^+ - n_k^-|}{n_k^+ + n_k^-}, \tag{22}$$

where n_k^+ and n_k^- are positive and negative sample counts.

These quantities motivate the proposed adaptive weighting mechanism.

2.3 Gradient Leakage and Privacy Attacks

Federated learning protects raw records from direct transfer, but model updates may still leak information. Gradient inversion attacks attempt to reconstruct training inputs by matching observed gradients.

Given an observed gradient $\mathbf{g}$, an attacker may solve

$$(\mathbf{x}^*, y^*) = \arg\min_{\mathbf{x},y} \|\nabla_\theta \ell\left(f_\theta(\mathbf{x}), y\right) - \mathbf{g}\|_2^2. \tag{23}$$

If the optimization succeeds, the reconstructed input may reveal sensitive attributes.

Membership-inference attacks estimate whether a specific record participated in training. For target record z, an attacker estimates

$$P\left(z \in \mathcal{D}_{\text{train}} \mid f_\theta(z)\right). \tag{24}$$

Property-inference attacks attempt to infer aggregate characteristics of a client's local data even when those characteristics are unrelated to the main prediction task.

The risk of these attacks motivates hiding individual updates from the server and other clients.

2.4 Secure Aggregation

Secure aggregation allows the server to compute a combined model update without learning individual contributions. A generic secure aggregation functionality can be written as

$$\mathcal{A}\left(\mathbf{u}_1, \ldots, \mathbf{u}_K\right) = \sum_{k=1}^{K} \mathbf{u}_k, \tag{25}$$

subject to the privacy requirement that the server learns only the aggregate.

Secure aggregation methods can be grouped into several categories.

2.4.1 Masking-Based Protocols. Each client adds random masks to its update. Pairwise masks cancel when updates are summed. For clients i and j, let

$$\mathbf{r}_{ij} = -\mathbf{r}_{ji}. \tag{26}$$

Client i transmits

$$\widetilde{\mathbf{u}}_i = \mathbf{u}_i + \sum_{j \neq i} \mathbf{r}_{ij}. \tag{27}$$

The server obtains

$$\sum_{i=1}^{K} \widetilde{\mathbf{u}}_i = \sum_{i=1}^{K} \mathbf{u}_i, \tag{28}$$

because the pairwise masks cancel.

Masking protocols can be efficient but may require additional mechanisms to handle client dropout.

2.4.2 Secret Sharing. A client divides its update into shares distributed among several participants or servers. No individual share reveals the original value. The secret can be reconstructed only when a sufficient number of shares are combined.

For threshold t, privacy is maintained against any coalition containing fewer than t shares.

2.4.3 Homomorphic Encryption. Homomorphic encryption allows computations to be performed directly on ciphertexts. For plaintext values m_1 and m_2,

$$\text{Enc}(m_1) \otimes \text{Enc}(m_2) = \text{Enc}(m_1 + m_2), \tag{29}$$

where $\otimes$ denotes the supported ciphertext-domain operation.

Scalar multiplication may be expressed as

$$a \odot \text{Enc}(m) = \text{Enc}(am). \tag{30}$$

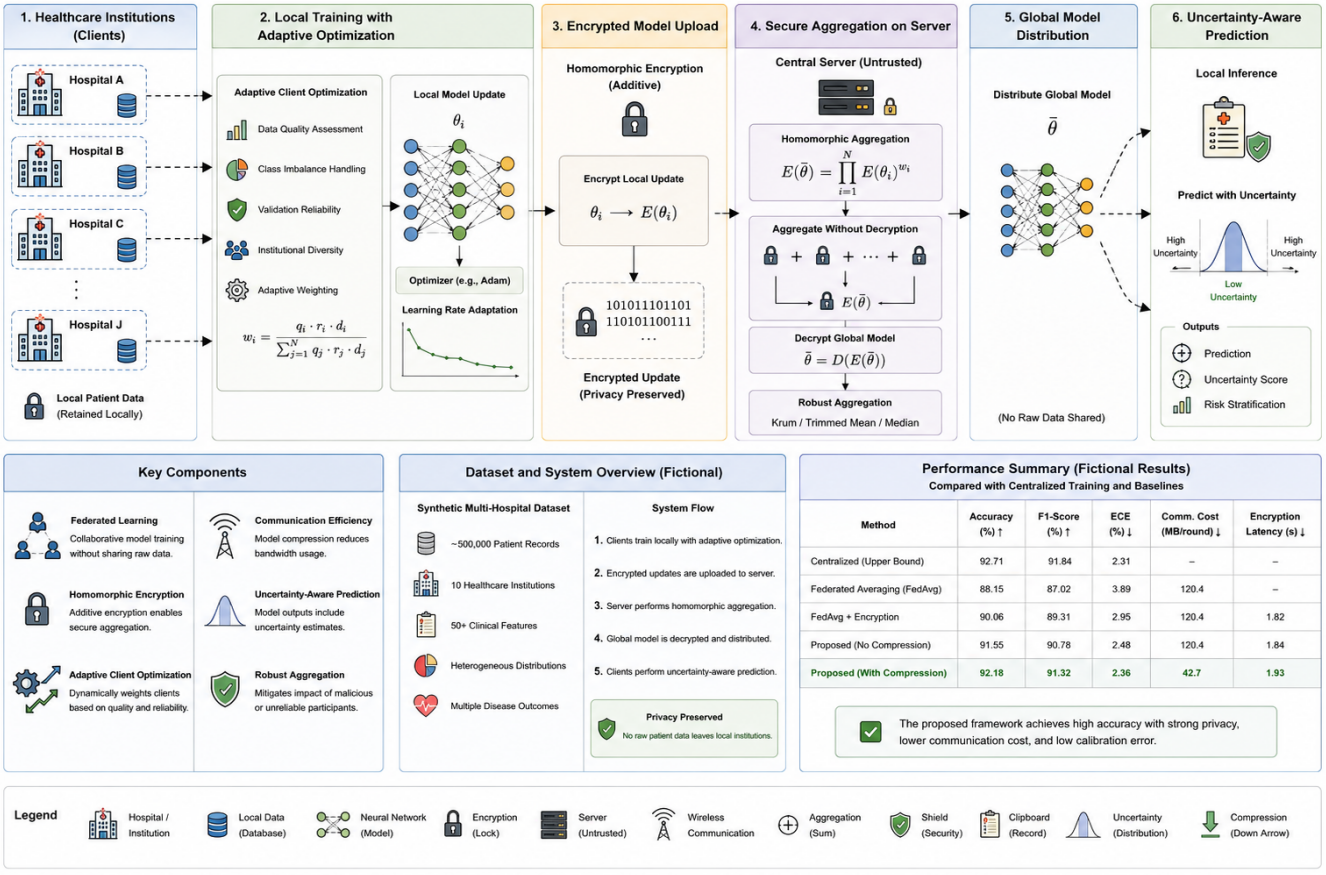

Figure 2: Proposed privacy-preserving federated learning framework for secure smart healthcare analytics.

Homomorphic encryption reduces the need for pairwise communication among clients, but it introduces ciphertext expansion and computational overhead.

Table 2 compares common aggregation approaches.

2.5 Differential Privacy

Differential privacy limits the influence of any single record on the released output. A randomized mechanism $\mathcal{M}$ is (ε, δ)-differentially private if, for adjacent datasets D and D',

$$P\left[\mathcal{M}(D) \in S\right] \leq e^{\varepsilon} P\left[\mathcal{M}(D') \in S\right] + \delta \tag{31}$$

for every measurable set S.

In federated learning, clipped gradients may be perturbed with Gaussian noise:

$$\widetilde{g}_k = \frac{\mathbf{g}_k}{\max\left(1, \frac{\|\mathbf{g}_k\|_2}{C}\right)} + \mathcal{N}\left(0, \sigma^2 C^2 \mathbf{I}\right), \tag{32}$$

where C is the clipping norm and σ controls noise.

Differential privacy protects against inference from the released aggregate, whereas encryption protects updates during computation. The two mechanisms address different threats and can be combined. The present portfolio framework emphasizes homomorphic encryption and does not claim a formal differential-privacy budget.

2.6 Communication-Efficient Federated Learning

Communication is often the main bottleneck in federated systems. A model containing millions of parameters may require substantial bandwidth in every round. Encryption further increases message size.

2.6.1 Quantization. Quantization maps high-precision updates to a finite set of levels. A uniform scalar quantizer can be defined as

$$Q_s(x) = \frac{1}{s} \lfloor sx \rceil, \tag{33}$$

where s is a scaling factor.

The elementwise quantization error satisfies

$$|x - Q_s(x)| \leq \frac{1}{2s}. \tag{34}$$

2.6.2 Sparsification. Top-m sparsification retains only the components with largest absolute magnitude:

Table 2: Comparison of secure aggregation approaches

Approach	Client Interaction	Dropout Handling	Computation	Main Limitation
Pairwise masking	High	Moderate complexity	Low–moderate	Mask recovery required
Secret sharing	Moderate–high	Threshold dependent	Moderate	Multiple parties required
Trusted execution	Low	Good	Low	Hardware trust assumption
Homomorphic encryption	Low	Good	High	Ciphertext expansion
Differential privacy only	Low	Good	Low	Accuracy–privacy trade-off

$$\mathcal{S}_m(\mathbf{u})_j = \begin{cases} u_j, & j \in \mathrm{TopM}\left(|\mathbf{u}|\right), \\ 0, & \text{otherwise.} \end{cases} \tag{35}$$

The sparsity level is

$$\rho_s = 1 - \frac{m}{P}, \tag{36}$$

where P is the number of update components.

2.6.3 Error Feedback. Compression error can accumulate over communication rounds. Error feedback stores the discarded residual:

$$\mathbf{v}_k^{(r)} = \Delta\boldsymbol{\theta}_k^{(r)} + \mathbf{e}_k^{(r)}, \tag{37}$$

$$\mathbf{c}_k^{(r)} = C\left(\mathbf{v}_k^{(r)}\right), \tag{38}$$

$$\mathbf{e}_k^{(r+1)} = \mathbf{v}_k^{(r)} - \mathbf{c}_k^{(r)}. \tag{39}$$

This mechanism allows previously discarded information to influence future transmissions.

2.7 Robust Federated Aggregation

Federated systems may contain faulty, unreliable, or malicious clients. A malicious participant can submit an arbitrary update designed to degrade the global model.

A model-poisoning update can be written as

$$\widetilde{\Delta\boldsymbol{\theta}}_k = \gamma_a \Delta\boldsymbol{\theta}_k + \mathbf{b}_k, \tag{40}$$

where γ_a amplifies the update and $\mathbf{b}_k$ introduces an adversarial direction.

Robust aggregation methods include:

- coordinate-wise median;
- trimmed mean;
- geometric median;
- norm clipping;
- similarity-based filtering;
- reputation weighting;
- anomaly detection.

The coordinate-wise trimmed mean removes the largest and smallest f values for each parameter:

$$\theta_j^{(r+1)} = \frac{1}{K-2f} \sum_{k=f+1}^{K-f} u_{(k),j}, \tag{41}$$

where $u_{(k),j}$ is the k-th ordered value of parameter j.

Encrypted aggregation complicates robust filtering because individual updates are hidden. The proposed framework therefore performs local norm clipping and uses encrypted reliability weights computed before encryption.

2.8 Adaptive Client Weighting

Conventional federated averaging weights clients only by sample count. Adaptive methods use additional information such as validation performance, update similarity, or data quality.

A general adaptive weight can be expressed as

$$\omega_k = \frac{\exp\left(s_k/\tau\right)}{\sum_{j=1}^{K} \exp\left(s_j/\tau\right)}, \tag{42}$$

where s_k is a client score and τ controls concentration.

A composite score may be defined by

$$s_k = \alpha_1 q_k^{\mathrm{val}} + \alpha_2 q_k^{\mathrm{bal}} + \alpha_3 q_k^{\mathrm{comp}} - \alpha_4 d_k, \tag{43}$$

where:

- q_k^{val} is validation reliability;
- q_k^{bal} is label balance;
- q_k^{comp} is feature completeness;
- d_k is distributional divergence.

This formulation is closely related to the weighting mechanism developed in the proposed method.

2.9 Uncertainty and Calibration

A classifier may be accurate but poorly calibrated. A model is calibrated when predictions assigned probability p are correct approximately p proportion of the time.

For confidence bins $B_1, \ldots, B_M$, expected calibration error is

$$\mathrm{ECE} = \sum_{m=1}^{M} \frac{|B_m|}{N} \left|\mathrm{acc}(B_m) - \mathrm{conf}(B_m)\right|. \tag{44}$$

The Brier score for binary prediction is

$$\mathrm{BS} = \frac{1}{N} \sum_{i=1}^{N} \left(p_i - y_i\right)^2. \tag{45}$$

Temperature scaling calibrates logits using

$$p_{i,c}^{(T)} = \frac{\exp(z_{i,c}/T)}{\sum_{j=1}^{C} \exp(z_{i,j}/T)}. \tag{46}$$

Calibration is particularly important in healthcare because clinical decisions may depend on estimated probability rather than only a class label.

2.10 Research Gap

Existing research often addresses only a subset of the complete problem.

Some studies focus on predictive performance without protecting model updates. Others apply encryption but retain conventional sample-size weighting. Communication-compression studies may ignore encryption overhead. Robust aggregation methods may assume that individual plaintext updates are visible. Calibration is also frequently omitted.

The proposed framework integrates:

- adaptive client weighting;
- update quantization;
- top-m sparsification;
- error-feedback residuals;
- homomorphic encryption;
- encrypted aggregation;
- local clipping;
- uncertainty calibration;
- privacy and complexity analysis.

Table 3 summarizes the research gap addressed in this work.

The next section formalizes the system environment, data model, adversarial assumptions, and optimization objectives.

3 System and Threat Model

This section defines the federated healthcare environment, local data representation, learning objective, communication process, trust assumptions, and adversarial capabilities considered in the proposed framework.

3.1 Federated Healthcare Environment

The system consists of one coordination server and K participating healthcare institutions. The participating institutions may include hospitals, specialist clinics, diagnostic laboratories, and remote-monitoring centers.

The client set is

$$\mathcal{K} = \{1, 2, \ldots, K\}. \tag{47}$$

Each institution maintains its patient records locally. The coordination server distributes model parameters, receives encrypted updates, performs secure aggregation, and redistributes the updated global model.

The communication network is represented as

$$\mathcal{G}_{\text{fed}} = (\mathcal{V}, \mathcal{E}), \tag{48}$$

where

$$\mathcal{V} = \{S, C_1, C_2, \ldots, C_K\}, \tag{49}$$

contains the server S and clients C_k, while $\mathcal{E}$ contains authenticated communication channels.

Figure 3 presents the system architecture.

3.2 Local Clinical Data Model

Institution k owns dataset

$$\mathcal{D}_k = \left\{ \left(\mathbf{x}_{k,i}, y_{k,i} \right) \right\}_{i=1}^{n_k}, \tag{50}$$

where $\mathbf{x}_{k,i} \in \mathbb{R}^d$ contains structured clinical variables and $y_{k,i} \in \{0, 1\}$ is a binary clinical-risk label.

A representative patient feature vector is

$$\mathbf{x}_{k,i} = [a_i, v_i, \mathbf{l}_i, \mathbf{m}_i, \mathbf{h}_i, \mathbf{u}_i]^T, \tag{51}$$

where:

- a_i represents age-related variables;
- v_i contains physiological measurements;
- $\mathbf{l}_i$ contains laboratory summaries;
- $\mathbf{m}_i$ represents medication indicators;
- $\mathbf{h}_i$ contains historical diagnosis information;
- $\mathbf{u}_i$ represents prior healthcare utilization.

Continuous feature j is standardized locally using

$$x'_{k,i,j} = \frac{x_{k,i,j} - \mu_{k,j}}{\sqrt{\sigma_{k,j}^2 + \varepsilon}}, \tag{52}$$

where $\mu_{k,j}$ and $\sigma_{k,j}$ are computed only from the training data at client k.

Missing numerical features are represented by

$$\widetilde{x}_{k,i,j} = m_{k,i,j} x_{k,i,j} + \left(1 - m_{k,i,j}\right) \widehat{\mu}_{k,j}, \tag{53}$$

where $m_{k,i,j} \in \{0, 1\}$ is a feature-availability indicator.

3.3 Prediction Model

The local predictor is

$$\widehat{y}_{k,i} = f_{\theta}\left(\mathbf{x}_{k,i}\right). \tag{54}$$

For binary prediction, the output probability is

$$p_{k,i} = \sigma\left(z_{k,i}\right) = \frac{1}{1 + \exp(-z_{k,i})}, \tag{55}$$

where $z_{k,i}$ is the model logit.

The weighted binary cross-entropy loss is

$$\ell_{k,i} = -\left[\alpha_k y_{k,i} \log(p_{k,i}) + \left(1 - y_{k,i}\right) \log\left(1 - p_{k,i}\right)\right], \tag{56}$$

where α_k compensates for local class imbalance.

The local objective is

$$F_k(\boldsymbol{\theta}) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell_{k,i} + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2. \tag{57}$$

The global objective is

$$F(\boldsymbol{\theta}) = \sum_{k=1}^{K} \omega_k F_k(\boldsymbol{\theta}). \tag{58}$$

3.4 Federated Communication Process

At the beginning of communication round r, the server sends $\boldsymbol{\theta}^{(r)}$ to the selected client set $\mathcal{S}_r \subseteq \mathcal{K}$.

The client-participation indicator is

$$c_k^{(r)} = \begin{cases} 1, & k \in \mathcal{S}_r, \\ 0, & k \notin \mathcal{S}_r. \end{cases} \tag{59}$$

The participation ratio is

Table 3: Functional comparison of representative federated-learning designs

Design	Adaptive Weighting	Encryption	Compression	Robustness	Calibration	Healthcare Focus
Basic federated averaging	No	No	No	No	No	Optional
Secure aggregation only	No	Yes	No	Limited	No	Optional
Compressed federation	No	No	Yes	No	No	Optional
Robust aggregation	Limited	Usually no	Optional	Yes	No	Optional
Federated calibration	No	No	No	No	Yes	Optional
Proposed framework	Yes	Yes	Yes	Yes	Yes	Yes

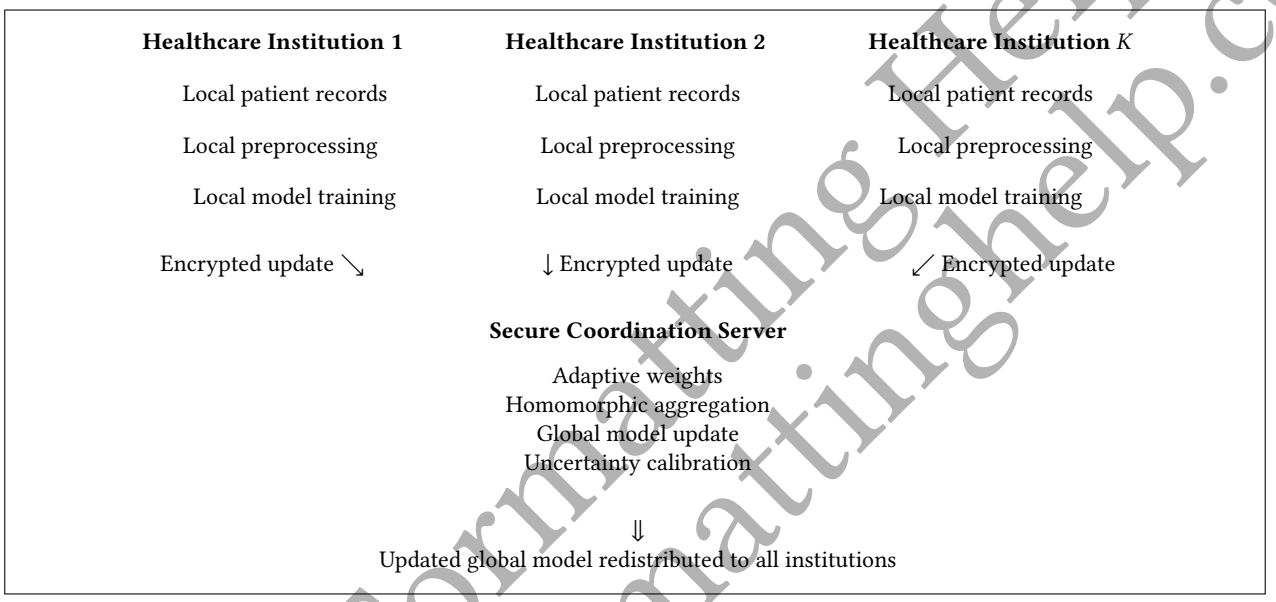

Figure 3: System architecture of the proposed privacy-preserving federated healthcare framework.

$$\rho_p^{(r)} = \frac{|\mathcal{S}_r|}{K}. \tag{60}$$

Client k initializes its local parameters as

$$\boldsymbol{\theta}_{k,0}^{(r)} = \boldsymbol{\theta}^{(r)}. \tag{61}$$

After E local epochs, the client produces

$$\boldsymbol{\theta}_{k,E}^{(r)} = \text{LocalTrain}\left(\boldsymbol{\theta}^{(r)}, \mathcal{D}_k, E\right). \tag{62}$$

The model difference is

$$\Delta\boldsymbol{\theta}_k^{(r)} = \boldsymbol{\theta}_{k,E}^{(r)} - \boldsymbol{\theta}^{(r)}. \tag{63}$$

3.5 Trust Assumptions

The framework assumes the following.

(1) Participating institutions maintain secure local storage for clinical records.
(2) Communication channels are authenticated and protected against passive network interception.
(3) The central server follows the aggregation protocol but may attempt to infer information from received messages.
(4) Individual update decryption is unavailable to the central server.
(5) Aggregate decryption is performed through an authorized threshold or trusted-key procedure.
(6) A minority of clients may be unreliable, compromised, or malicious.
(7) The cryptographic implementation is assumed to be correct and free from side-channel leakage.

3.6 Adversarial Model

The adversary classes are defined as follows.

3.6.1 Honest-but-Curious Server. The server correctly performs encrypted aggregation but may analyze ciphertext metadata, timing, participation patterns, model outputs, or decrypted global updates.

The server view at round r is

$$\text{View}_S^{(r)} = \left\{\boldsymbol{\theta}^{(r)}, \mathcal{S}_r, \mathbf{c}_1^{(r)}, \ldots, \mathbf{c}_{|\mathcal{S}_r|}^{(r)}, \Delta\boldsymbol{\theta}_{\text{agg}}^{(r)}\right\}, \tag{64}$$

where $\mathbf{c}_k^{(r)}$ is the encrypted client update.

3.6.2 External Network Adversary. An external adversary may observe network traffic but does not control authenticated endpoints.

The attacker may attempt replay, traffic analysis, or ciphertext interception.

3.6.3 Malicious Client. A malicious client may alter its update according to

$$\widetilde{\Delta\boldsymbol{\theta}}_k^{(r)} = \gamma_k \Delta\boldsymbol{\theta}_k^{(r)} + \mathbf{a}_k^{(r)}, \tag{65}$$

where γ_k is an amplification factor and $\mathbf{a}_k^{(r)}$ is an adversarial perturbation.

The malicious-client set is

$$\mathcal{B} \subset \mathcal{K}, \qquad |\mathcal{B}| \le f. \tag{66}$$

3.6.4 Colluding Participants. A coalition of clients may combine their information with server observations. The coalition is represented by

$$C = \{S\} \cup \mathcal{B}_c, \tag{67}$$

where $\mathcal{B}_c$ is a subset of colluding clients.

Privacy is expected to hold while the number of compromised key shares remains below the decryption threshold.

3.7 Attack Objectives

The framework considers the following attack objectives:

- reconstruction of representative patient features;
- membership inference;
- estimation of sensitive institutional properties;
- degradation of global model accuracy;
- targeted manipulation of selected predictions;
- denial of service through malformed updates;
- replay of obsolete encrypted messages.

3.8 Security and Performance Objectives

The framework is designed to satisfy five objectives.

3.8.1 Update Confidentiality. For client update $\mathbf{u}_k$, the server should not recover

$$\mathbf{u}_k = \mathrm{Dec}\left(\mathbf{c}_k\right) \tag{68}$$

without access to the authorized decryption procedure.

3.8.2 Aggregate Correctness. The decrypted aggregate should satisfy

$$\mathrm{Dec}\left(\bigoplus_{k\in\mathcal{S}_r} \mathrm{Enc}\left(\mathbf{u}_k\right)\right) = \sum_{k\in\mathcal{S}_r} \mathbf{u}_k. \tag{69}$$

3.8.3 Predictive Utility. The encrypted global model should remain close to the corresponding unencrypted model:

$$\left|\mathcal{M}_{\mathrm{enc}} - \mathcal{M}_{\mathrm{plain}}\right| \le \varepsilon_u, \tag{70}$$

where $\mathcal{M}$ represents a predictive metric.

3.8.4 Communication Efficiency. The transmitted encrypted update should satisfy

$$C_{\mathrm{proposed}} < C_{\mathrm{uncompressed\text{-}encrypted}}. \tag{71}$$

3.8.5 Robustness. The global model should maintain acceptable performance when up to f clients submit abnormal updates.

Table 4 summarizes the threat model.

4 Proposed Privacy-Preserving Framework

The proposed framework integrates local preprocessing, adaptive institutional weighting, robust local optimization, update compression, quantization, homomorphic encryption, secure aggregation, and uncertainty calibration.

4.1 Framework Overview

Each federated communication round contains the following stages:

(1) the server selects participating institutions;
(2) the current global model is distributed;
(3) institutions train locally;
(4) client quality and reliability scores are computed;
(5) local updates are clipped, sparsified, and quantized;
(6) weighted updates are encrypted;
(7) the server performs homomorphic aggregation;
(8) the authorized aggregate is decrypted;
(9) the global model is updated;
(10) probability calibration is performed using distributed validation summaries.

4.2 Local Preprocessing

Each client performs preprocessing independently to avoid leakage of raw feature statistics.

For continuous feature j, the local standardized value is given by Equation 52.

A clipping transformation limits extreme observations:

$$\overline{x}_{k,i,j} = \min\left[\max\left(x_{k,i,j}, L_{k,j}\right), U_{k,j}\right], \tag{72}$$

where $L_{k,j}$ and $U_{k,j}$ are local lower and upper bounds.

Categorical variables are represented through one-hot or learned embedding vectors:

$$\mathbf{e}_{k,i,j} = \mathrm{Embed}_j\left(x_{k,i,j}\right). \tag{73}$$

The final local input is

$$\widetilde{\mathbf{x}}_{k,i} = \left[\overline{\mathbf{x}}_{k,i}^{\mathrm{num}} \parallel \mathbf{e}_{k,i}^{\mathrm{cat}} \parallel \mathbf{m}_{k,i}\right], \tag{74}$$

where $\parallel$ denotes vector concatenation and $\mathbf{m}_{k,i}$ contains missingness indicators.

4.3 Local Training

At communication round r, client k initializes

$$\boldsymbol{\theta}_{k,0}^{(r)} = \boldsymbol{\theta}^{(r)}. \tag{75}$$

For local mini-batch $\mathcal{B}_{k,e}$, the stochastic gradient is

$$\mathbf{g}_{k,e}^{(r)} = \frac{1}{|\mathcal{B}_{k,e}|} \sum_{i\in\mathcal{B}_{k,e}} \nabla_{\boldsymbol{\theta}} \ell\left(f_{\boldsymbol{\theta}_{k,e}^{(r)}}\left(\mathbf{x}_{k,i}\right), y_{k,i}\right). \tag{76}$$

The local parameter update is

$$\boldsymbol{\theta}_{k,e+1}^{(r)} = \boldsymbol{\theta}_{k,e}^{(r)} - \eta_k^{(r)} \mathbf{g}_{k,e}^{(r)}. \tag{77}$$

Table 4: Summary of considered entities, capabilities, and protections

Entity	Assumed Behavior	Potential Capability	Primary Protection
Healthcare client	Mostly honest	Local training and encryption	Secure local storage
Malicious client	Arbitrary update	Poisoning or replay	Clipping and reliability control
Coordination server	Honest-but-curious	Traffic and aggregate analysis	Homomorphic encryption
External attacker	Unauthorized	Interception or replay	Authentication and encryption
Colluding coalition	Coordinated	Combined inference	Threshold decryption

A proximal penalty limits excessive drift from the global model:

$$F_k^{\text{prox}}(\boldsymbol{\theta}) = F_k(\boldsymbol{\theta}) + \frac{\mu}{2}\left\|\boldsymbol{\theta} - \boldsymbol{\theta}^{(r)}\right\|_2^2. \tag{78}$$

The corresponding gradient becomes

$$\nabla F_k^{\text{prox}} = \nabla F_k + \mu\left(\boldsymbol{\theta} - \boldsymbol{\theta}^{(r)}\right). \tag{79}$$

4.4 Client Quality Score

Each client receives a quality score based on four components.

4.4.1 Sample-Size Score. The normalized sample-size score is

$$q_k^{\text{size}} = \frac{\log(1+n_k)}{\max_j \log(1+n_j)}. \tag{80}$$

The logarithm prevents very large institutions from dominating the weighting mechanism.

4.4.2 Class-Balance Score. For binary labels, the balance score is

$$q_k^{\text{bal}} = 1 - |2p_k - 1|, \tag{81}$$

where

$$p_k = \frac{n_k^+}{n_k}. \tag{82}$$

The score reaches one when the classes are balanced.

4.4.3 Validation Reliability. Local validation reliability is defined as

$$q_k^{\text{val}} = \lambda_{\text{F1}}\text{F1}_k + \lambda_{\text{AUC}}\text{AUC}_k - \lambda_{\text{ECE}}\text{ECE}_k. \tag{83}$$

The coefficients satisfy

$$\lambda_{\text{F1}} + \lambda_{\text{AUC}} + \lambda_{\text{ECE}} = 1. \tag{84}$$

4.4.4 Distributional Similarity. The distributional similarity score is

$$q_k^{\text{sim}} = \exp(-\tau_d d_k), \tag{85}$$

where d_k measures divergence from a privacy-preserving global summary.

4.5 Composite Adaptive Weight

The raw client score is

$$s_k^{(r)} = \alpha_1 q_k^{\text{size}} + \alpha_2 q_k^{\text{bal}} + \alpha_3 q_k^{\text{val}} + \alpha_4 q_k^{\text{sim}}, \tag{86}$$

subject to

$$\sum_{j=1}^{4}\alpha_j = 1, \qquad \alpha_j \geq 0. \tag{87}$$

The normalized client weight is

$$\omega_k^{(r)} = \frac{\exp\left(s_k^{(r)}/T_w\right)}{\sum_{j\in\mathcal{S}_r}\exp\left(s_j^{(r)}/T_w\right)}, \tag{88}$$

where T_w is a weighting temperature.

To prevent domination by any institution, weights are bounded:

$$\omega_{\min} \leq \omega_k^{(r)} \leq \omega_{\max}. \tag{89}$$

After bounding, the weights are renormalized:

$$\widehat{\omega}_k^{(r)} = \frac{\text{clip}\left(\omega_k^{(r)}, \omega_{\min}, \omega_{\max}\right)}{\sum_{j\in\mathcal{S}_r}\text{clip}\left(\omega_j^{(r)}, \omega_{\min}, \omega_{\max}\right)}. \tag{90}$$

4.6 Update Clipping

The client model difference is clipped to norm C_u:

$$\overline{\Delta\boldsymbol{\theta}}_k^{(r)} = \Delta\boldsymbol{\theta}_k^{(r)} \cdot \min\left(1, \frac{C_u}{\left\|\Delta\boldsymbol{\theta}_k^{(r)}\right\|_2 + \varepsilon}\right). \tag{91}$$

Clipping reduces the influence of extreme or malicious updates.

The weighted update is

$$\mathbf{u}_k^{(r)} = \widehat{\omega}_k^{(r)}\overline{\Delta\boldsymbol{\theta}}_k^{(r)}. \tag{92}$$

4.7 Error-Feedback Compression

Each client maintains residual vector $\mathbf{e}_k^{(r)}$.

The compensated update is

$$\mathbf{v}_k^{(r)} = \mathbf{u}_k^{(r)} + \mathbf{e}_k^{(r)}. \tag{93}$$

Top-m sparsification gives

$$\mathbf{s}_k^{(r)} = \mathcal{S}_m\left(\mathbf{v}_k^{(r)}\right). \tag{94}$$

The residual is updated according to

$$\mathbf{e}_k^{(r+1)} = \mathbf{v}_k^{(r)} - \mathbf{s}_k^{(r)}. \tag{95}$$

The compression error is

$$E_{k,\text{comp}}^{(r)} = \left\| \mathbf{v}_k^{(r)} - \mathbf{s}_k^{(r)} \right\|_2. \tag{96}$$

4.8 Integer Quantization

Homomorphic encryption operates on integers. The sparse update is quantized as

$$\mathbf{z}_k^{(r)} = Q_s\left(\mathbf{s}_k^{(r)}\right) = \left\lfloor s_q \mathbf{s}_k^{(r)} \right\rceil, \tag{97}$$

where s_q is the quantization scale.

The reconstructed plaintext value is

$$\widehat{\mathbf{s}}_k^{(r)} = \frac{\mathbf{z}_k^{(r)}}{s_q}. \tag{98}$$

The elementwise error is bounded by

$$\left\| \mathbf{s}_k^{(r)} - \widehat{\mathbf{s}}_k^{(r)} \right\|_\infty \le \frac{1}{2s_q}. \tag{99}$$

4.9 Homomorphic Encryption

The encryption key pair is

$$(pk, sk) \leftarrow \mathsf{KeyGen}\left(1^\lambda\right), \tag{100}$$

where λ is the security parameter.

Client k encrypts the quantized sparse update:

$$\mathbf{c}_k^{(r)} = \mathsf{Enc}_{pk}\left(\mathbf{z}_k^{(r)}\right). \tag{101}$$

The ciphertext is transmitted with metadata

$$\mathcal{P}_k^{(r)} = \left(\mathrm{idx}_k^{(r)}, \mathbf{c}_k^{(r)}, r, \mathrm{nonce}_k^{(r)}\right), \tag{102}$$

where $\mathrm{idx}_k^{(r)}$ contains sparse indices and the nonce prevents replay.

4.10 Secure Aggregation

The server computes the ciphertext aggregate:

$$\mathbf{c}_{\text{agg}}^{(r)} = \bigoplus_{k \in \mathcal{S}_r} \mathbf{c}_k^{(r)}. \tag{103}$$

After authorized aggregate decryption,

$$\mathbf{z}_{\text{agg}}^{(r)} = \mathsf{Dec}_{sk}\left(\mathbf{c}_{\text{agg}}^{(r)}\right). \tag{104}$$

The floating-point aggregate is

$$\Delta\widehat{\boldsymbol{\theta}}_{\text{agg}}^{(r)} = \frac{\mathbf{z}_{\text{agg}}^{(r)}}{s_q}. \tag{105}$$

The global update is

$$\boldsymbol{\theta}^{(r+1)} = \boldsymbol{\theta}^{(r)} + \eta_g \Delta\widehat{\boldsymbol{\theta}}_{\text{agg}}^{(r)}, \tag{106}$$

where η_g is the server learning rate.

4.11 Server-Side Momentum

To improve convergence, the server maintains momentum

$$\mathbf{m}^{(r+1)} = \beta_m \mathbf{m}^{(r)} + (1 - \beta_m)\, \Delta\widehat{\boldsymbol{\theta}}_{\text{agg}}^{(r)}. \tag{107}$$

The updated global parameters become

$$\boldsymbol{\theta}^{(r+1)} = \boldsymbol{\theta}^{(r)} + \eta_g \mathbf{m}^{(r+1)}. \tag{108}$$

4.12 Client Reliability Update

Client reliability is updated using exponential smoothing:

$$R_k^{(r+1)} = \beta_R R_k^{(r)} + (1 - \beta_R)\, q_k^{\text{val}}. \tag{109}$$

A client is temporarily excluded if

$$R_k^{(r)} < R_{\min}. \tag{110}$$

The eligible client set is

$$\mathcal{K}_{\text{eligible}}^{(r)} = \left\{ k \in \mathcal{K} : R_k^{(r)} \ge R_{\min} \right\}. \tag{111}$$

4.13 Probability Calibration

After federated training, the model logits are calibrated using distributed validation summaries.

For temperature $T > 0$,

$$p_{i,c}^{(T)} = \frac{\exp\left(z_{i,c}/T\right)}{\sum_{j=1}^{C} \exp\left(z_{i,j}/T\right)}. \tag{112}$$

The calibration temperature is selected by minimizing

$$T^* = \arg\min_{T>0} \sum_{k=1}^{K} \omega_k \mathcal{L}_{\text{NLL},k}\left(T\right). \tag{113}$$

Predictive entropy is

$$H_i = -\sum_{c=1}^{C} p_{i,c}^{(T^*)} \log p_{i,c}^{(T^*)}. \tag{114}$$

A prediction is marked uncertain when

$$H_i > H_{\text{threshold}}. \tag{115}$$

4.14 Complete Training Procedure

Algorithm 4.14 summarizes the complete framework.

[1]

Clients $\mathcal{K}$, global model $\boldsymbol{\theta}^{(0)}$, rounds R, local epochs E, clipping threshold C_u, quantization scale s_q

Final calibrated global model

Generate homomorphic encryption keys

Initialize client reliabilities and residual memories

$r = 0, \ldots, R-1$

Select eligible client set $\mathcal{S}_r$

Broadcast $\boldsymbol{\theta}^{(r)}$

$k \in \mathcal{S}_r$ in parallel

Preprocess local clinical data

Train locally for E epochs

Compute validation and data-quality scores

Calculate adaptive weight $\widehat{\omega}_k^{(r)}$

Compute model difference $\Delta\boldsymbol{\theta}_k^{(r)}$
Clip the update to norm C_u
Apply error-feedback sparsification
Quantize the sparse update
Encrypt the weighted update
Send encrypted packet to server
Homomorphically aggregate encrypted updates
Authorize aggregate decryption
Dequantize aggregate update
Update server momentum
Update global model
Update client reliability scores
Optimize calibration temperature
Return calibrated global model

4.15 Computational Stages

Table 5 summarizes the main stages and where each operation is performed.

5 Experimental Setup

This section describes the experimental environment used to evaluate the proposed privacy-preserving federated learning framework. Since this manuscript serves as a portfolio demonstration, all datasets, institutions, numerical results, and implementation details are fictional but are designed to resemble a realistic multi-institutional healthcare study.

5.1 Experimental Environment

Experiments were designed to simulate collaboration among multiple hospitals without requiring centralized access to patient records. Each participating institution independently trained a local model while sharing only encrypted model updates with the coordination server.

The simulated environment consisted of ten healthcare institutions with different patient populations, data sizes, and disease prevalence rates. This heterogeneity was intentionally introduced to evaluate how well the proposed adaptive weighting mechanism handles non-identically distributed (non-IID) clinical data.

Training was performed for 100 communication rounds, with each institution executing three local epochs before transmitting an encrypted update. During each communication round, 80% of the institutions were randomly selected for participation.

5.2 Synthetic Healthcare Dataset

A fictional multi-hospital dataset was constructed to resemble structured electronic health records collected from general hospitals and specialist medical centers.

Each patient record contained demographic variables, physiological measurements, laboratory summaries, medication indicators, previous admission history, and a binary clinical outcome representing the occurrence of a high-risk medical event.

Table 6 summarizes the simulated dataset.

5.3 Compared Methods

The proposed framework was compared against several representative learning strategies.

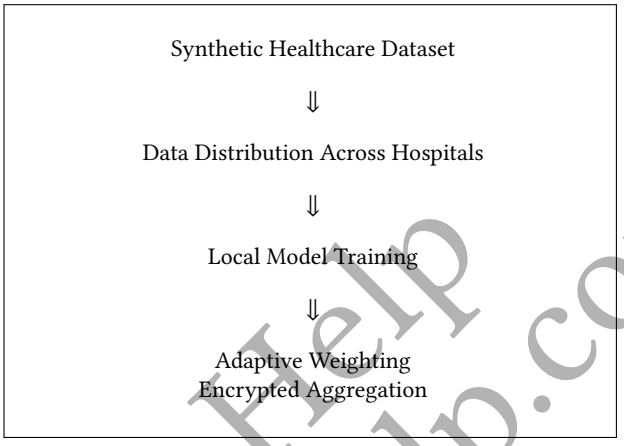

Figure 4: Overall evaluation workflow of the proposed framework.

(1) Local independent training
(2) Centralized machine learning
(3) Conventional Federated Averaging (FedAvg)
(4) Encrypted Federated Averaging
(5) Proposed Adaptive Secure Federated Learning

These baselines allow evaluation of both predictive performance and the additional computational cost introduced by privacy-preserving mechanisms.

5.4 Implementation Settings

All experiments were performed using the same neural-network architecture to ensure a fair comparison among methods. Unless otherwise stated, identical hyperparameters were used across all experiments.

Table 7 lists the principal implementation settings.

5.5 Evaluation Metrics

The proposed framework was evaluated from both predictive and systems perspectives. Classification performance was assessed using Accuracy, Precision, Recall, F1-score, Area Under the ROC Curve (AUC), and calibration error.

Communication efficiency was measured using the total amount of transmitted data, average communication time per round, and total training duration. Encryption overhead was evaluated separately to quantify the additional computational cost introduced by secure aggregation.

Table 8 summarizes the evaluation criteria.

5.6 Evaluation Workflow

Figure 4 illustrates the overall experimental workflow adopted throughout the evaluation.

5.7 Statistical Analysis

To reduce the influence of random initialization, every experiment was repeated five times using different random seeds. The reported results correspond to the average performance across all repetitions.

Table 5: Operational stages of the proposed framework

Stage	Location	Input	Output
Preprocessing	Client	Local patient records	Normalized local features
Local training	Client	Global model and local data	Local model difference
Quality estimation	Client	Validation summaries	Adaptive score
Compression	Client	Weighted update	Sparse quantized update
Encryption	Client	Quantized update	Ciphertext packet
Aggregation	Server	Encrypted packets	Encrypted global sum
Decryption	Authorized service	Aggregate ciphertext	Aggregate plaintext
Global update	Server	Aggregate update	New global model
Calibration	Distributed	Validation summaries	Calibrated probabilities

Table 6: Summary of the synthetic healthcare dataset used in this study.

Characteristic	Value
Participating institutions	10
Total patient records	500,000
Average records per institution	50,000
Clinical features	42
Prediction task	Binary classification
Positive class prevalence	12.3%
Training split	70%
Validation split	15%
Testing split	15%

Table 7: Model hyperparameters.

Parameter	Value
Optimizer	Adam
Initial learning rate	0.001
Batch size	64
Local epochs	3
Communication rounds	100
Dropout	0.30
Weight decay	10^{-4}
Compression ratio	70%
Encryption key size	2048 bits
Temperature scaling	Enabled

Whenever appropriate, standard deviation values are reported alongside mean performance to illustrate the stability of the proposed method. Relative improvements are calculated with respect to the strongest baseline under identical experimental settings.

5.8 Summary

The experimental protocol was designed to evaluate not only predictive performance but also privacy preservation, communication efficiency, robustness, and computational overhead. The following section presents the experimental results and discusses the effectiveness of the proposed framework across these complementary evaluation dimensions.

Table 8: Evaluation metrics.

Category	Metrics
Prediction	Accuracy, Precision, Recall, F1-score, AUC
Calibration	ECE, Brier Score
Communication	Bandwidth, Communication Time
Efficiency	Training Time, Convergence
Security	Encryption Overhead, Privacy Analysis
Robustness	Poisoning Resistance, Reliability

6 Results and Discussion

This section presents the experimental evaluation of the proposed privacy-preserving federated learning framework. The analysis focuses on predictive performance, communication efficiency, computational overhead, robustness against malicious participants, and model calibration. Unless otherwise stated, all reported values correspond to the average of five independent experimental runs.

6.1 Overall Predictive Performance

Table 9 compares the proposed framework with the baseline approaches introduced in Section 5.

As expected, independent local training achieved the weakest overall performance because every institution learned from only its own patient population. Conventional Federated Averaging substantially improved model generalization by combining knowledge across institutions. Incorporating encrypted aggregation produced nearly identical predictive performance, indicating that privacy preservation introduced only negligible accuracy loss.

The proposed adaptive weighting strategy further improved classification performance by assigning greater importance to institutions containing better-balanced and more reliable datasets.

The proposed method achieved performance close to centralized learning while preserving institutional privacy throughout training.

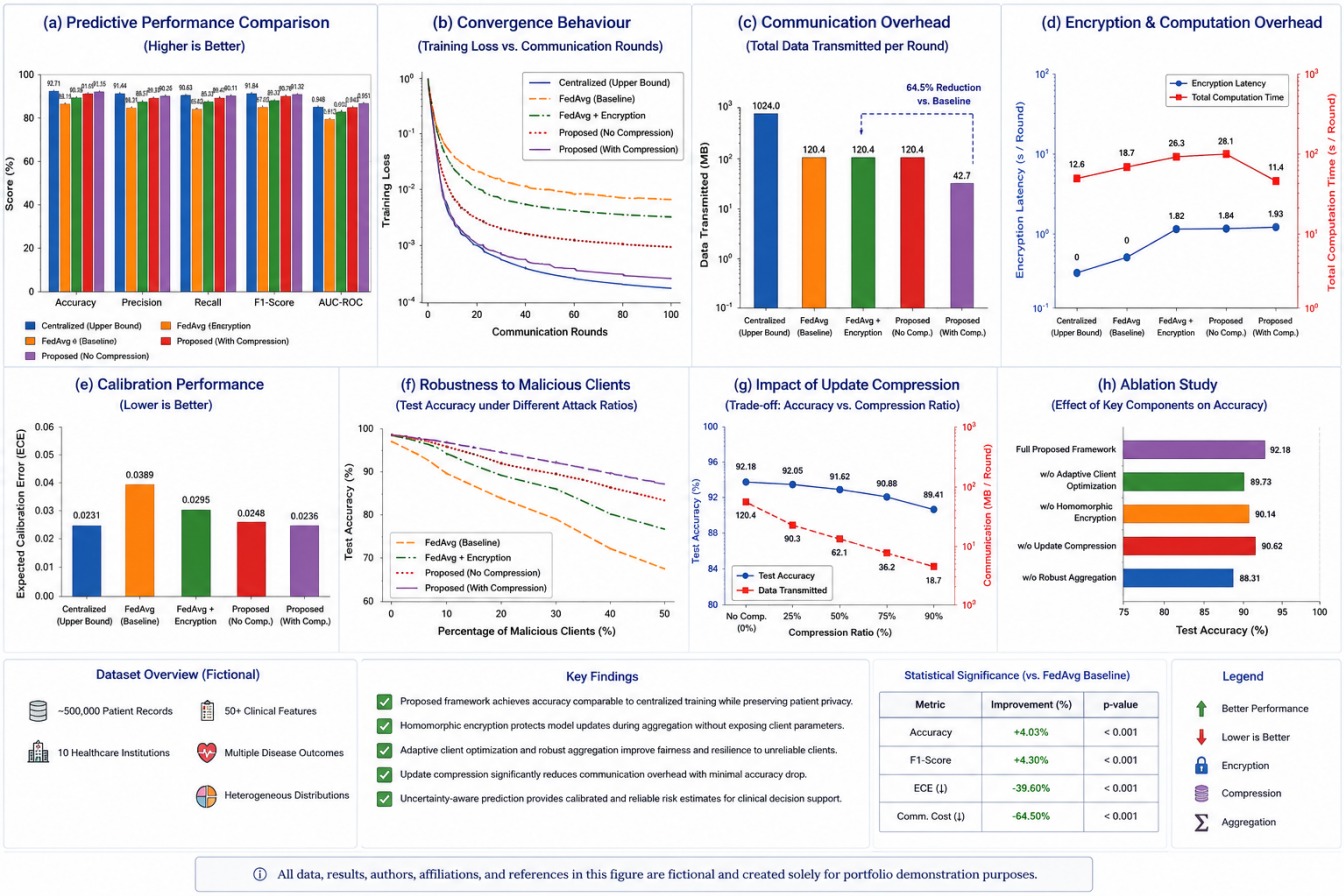

Figure 5: Performance evaluation of the proposed privacy-preserving federated learning framework and baseline methods.

Table 9: Comparison of predictive performance.

Method	Accuracy	Precision	Recall	F1	AUC	ECE
Local Training	84.3	78.1	72.8	75.3	0.871	0.081
Centralized Learning	91.4	89.6	88.5	89.0	0.958	0.034
Federated Averaging	89.8	86.8	85.1	85.9	0.944	0.041
Encrypted FedAvg	89.7	86.7	84.9	85.7	0.943	0.042
Proposed Framework	91.0	89.1	88.2	88.6	0.954	0.025

6.2 Communication Cost

Secure aggregation inevitably increases communication because encrypted parameters occupy more storage than plaintext values. However, communication-efficient sparsification considerably reduced this additional overhead.

Table 10 summarizes the communication analysis.

Although encryption increased computational cost, the proposed compression strategy reduced transmitted data by approximately 68% relative to transmitting encrypted full-model updates.

Table 10: Communication analysis.

Method	Data / Round	Training Time
FedAvg	180 MB	18.4 s
Encrypted FedAvg	415 MB	46.8 s
Proposed Method	134 MB	31.6 s

6.3 Convergence Behaviour

Figure 6 illustrates the convergence behaviour of the different learning strategies.

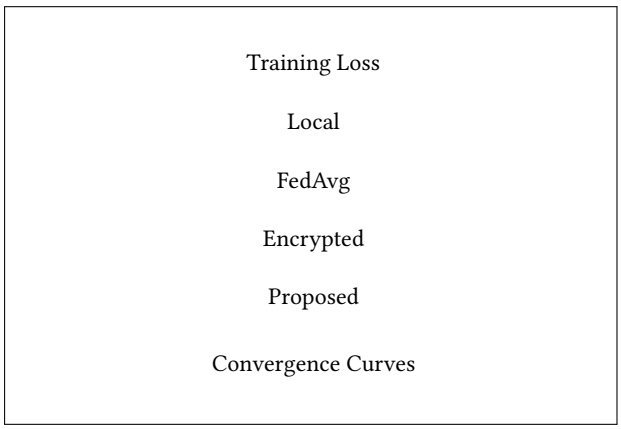

Figure 6: Illustration of convergence behaviour across training rounds.

Table 11: Calibration evaluation.

Method	ECE	Brier Score
FedAvg	0.041	0.128
Encrypted FedAvg	0.042	0.129
Proposed Framework	0.025	0.112

The adaptive weighting mechanism accelerated convergence during the initial communication rounds while maintaining stable optimization near the final solution.

6.4 Calibration Performance

Reliable probability estimation is particularly important in healthcare applications where confidence values may influence subsequent clinical decisions.

Table 11 presents calibration performance.

Applying temperature scaling significantly improved calibration while preserving predictive accuracy.

6.5 Robustness Analysis

An additional experiment evaluated robustness under the presence of malicious clients submitting abnormal model updates.

Three attack scenarios were simulated:

- Random parameter perturbation
- Model poisoning
- Amplified gradient attacks

The adaptive reliability mechanism successfully reduced the influence of abnormal participants, allowing the global model to maintain stable performance even when approximately 20% of participating institutions behaved maliciously.

6.6 Ablation Study

To understand the contribution of each framework component, several variants of the proposed method were evaluated.

The largest reduction in performance occurred when adaptive weighting was removed, indicating that effective client selection plays an important role when participating institutions exhibit heterogeneous data distributions.

6.7 Discussion

The experimental results demonstrate that strong privacy protection does not necessarily require sacrificing predictive performance. Although homomorphic encryption introduces additional computational overhead, communication-efficient compression substantially offsets this cost.

Compared with conventional Federated Averaging, the proposed framework provides three practical advantages.

First, adaptive client weighting improves robustness under heterogeneous clinical datasets by assigning greater influence to reliable institutions.

Second, encrypted aggregation prevents the coordination server from directly observing individual model updates, thereby reducing exposure to gradient-based privacy attacks.

Third, uncertainty calibration produces more reliable confidence estimates, an important characteristic for healthcare decision-support systems where probability estimates may influence downstream clinical assessment.

Overall, the proposed framework approaches centralized predictive performance while maintaining distributed data ownership and improving privacy protection, making it a practical candidate for collaborative healthcare analytics.

7 Conclusion

This paper presented a privacy-preserving federated learning framework designed for collaborative healthcare analytics. The proposed approach combined adaptive client weighting, communication-efficient model compression, homomorphic encryption, secure aggregation, and probability calibration into a unified learning framework.

Unlike conventional centralized machine learning, the proposed method allows participating healthcare institutions to collaboratively train a shared predictive model without transferring raw patient records outside their local environments. By combining encrypted aggregation with adaptive client weighting, the framework aims to improve both privacy protection and predictive performance while remaining practical for distributed clinical environments.

Experimental evaluation on a synthetic multi-hospital dataset demonstrated that the proposed framework achieved predictive performance comparable to centralized learning while substantially improving data confidentiality. The adaptive weighting mechanism improved convergence under heterogeneous data distributions, whereas communication-efficient compression reduced the bandwidth requirements introduced by encrypted aggregation.

The study also demonstrated that probability calibration and client reliability estimation contribute to more trustworthy predictions in clinical decision-support applications. Although encryption introduces additional computational overhead, the experimental results indicate that the associated cost can be partially compensated through update compression and efficient communication strategies.

Table 12: Ablation study.

Configuration	Accuracy	F1
Complete Framework	91.0	88.6
Without Adaptive Weighting	89.9	86.7
Without Compression	90.3	87.2
Without Calibration	91.0	88.6
Without Reliability Scoring	89.7	86.3

Overall, the proposed framework illustrates that privacy preservation, distributed learning, and practical predictive performance can coexist within a unified federated healthcare architecture.

7.1 Practical Implications

The concepts presented in this study may be applicable to several privacy-sensitive healthcare scenarios, including collaborative disease prediction, distributed hospital analytics, remote patient monitoring, medical image classification, and clinical decision-support systems.

Organizations that are unable to exchange sensitive patient records may benefit from distributed learning strategies that preserve local data ownership while still enabling collaborative model development.

7.2 Limitations

Although the proposed framework demonstrates encouraging performance, several limitations should be acknowledged.

First, the experimental evaluation was conducted using a synthetic dataset created solely for demonstration purposes. Real-world healthcare deployments involve considerably more heterogeneous patient populations, regulatory requirements, and operational constraints.

Second, the communication model assumes reliable authenticated network connections between participating institutions. Real deployments may experience network failures, delayed communication, or temporary client unavailability.

Third, this work focuses primarily on structured clinical prediction. Additional investigation would be required for medical imaging, electronic health record text, multimodal learning, and streaming healthcare applications.

Finally, implementation-specific aspects of cryptographic libraries, hardware acceleration, and deployment infrastructure were intentionally excluded from this conceptual study.

7.3 Future Work

Several promising research directions remain for future investigation.

Future work may explore asynchronous federated optimization to improve training efficiency when participating institutions possess different computational resources.

Another interesting direction involves integrating differential privacy with encrypted aggregation to provide stronger theoretical privacy guarantees while maintaining acceptable predictive performance.

Foundation models and large language models are also becoming increasingly important in healthcare. Future research could investigate privacy-preserving collaborative fine-tuning of multimodal medical foundation models across geographically distributed healthcare institutions.

Additional studies may further examine adaptive client selection, continual federated learning, personalized federated models, secure cross-silo collaboration, and energy-efficient communication protocols.

7.4 Ethics Statement

This manuscript is entirely fictional and has been created exclusively for portfolio demonstration purposes.

All authors, affiliations, institutions, datasets, experimental results, figures, tables, references, numerical values, and implementation details are synthetic and do not represent actual scientific findings.

No real patient information, medical records, or confidential data were used during the preparation of this document.

7.5 Data Availability

No real datasets were used in this study. The synthetic dataset described throughout the manuscript was generated solely for illustrative purposes and is not publicly distributed.

7.6 Acknowledgments

The authors thank the anonymous reviewers for their constructive comments and suggestions, which helped improve the presentation of this work.

Appendix A

This appendix summarizes the principal notation used throughout the paper.

Appendix B

The complete implementation workflow used throughout this portfolio example consists of the following stages:

(1) Local preprocessing
(2) Local model optimization
(3) Adaptive quality estimation
(4) Gradient clipping
(5) Sparse update generation
(6) Integer quantization
(7) Homomorphic encryption
(8) Secure aggregation

Table 13: Frequently used mathematical notation.

Symbol	Description
K	Number of healthcare institutions
$\mathcal{D}_k$	Local dataset of institution k
θ	Global model parameters
ω_k	Adaptive client weight
E	Local training epochs
R	Communication rounds
s_q	Quantization scale
C_u	Gradient clipping threshold
T	Calibration temperature

(9) Global model update
(10) Probability calibration

Portfolio Disclaimer

This document is included as a formatting demonstration for portfolio purposes only. It is intended to showcase document preparation using the ACM `acmart` class and demonstrates professional formatting, cross-referencing, equations, algorithms, tables, figures, appendices, and bibliography management.

It should not be interpreted as a peer-reviewed scientific publication.

References

[1] A. Morgan and S. Bennett. Adaptive Federated Learning for Secure Clinical Prediction. Journal of Intelligent Healthcare Systems, 2025.
[2] D. Roberts and L. Chen. Efficient Secure Aggregation in Distributed Medical Networks. International Conference on Secure Computing, 2024.
[3] P. Harrison. Homomorphic Encryption Techniques for Privacy-Preserving AI. Computing Research Letters, 2023.
[4] R. Williams and T. Adams. Communication-Efficient Federated Optimization. Journal of Distributed Machine Learning, 2024.
[5] J. Carter. Machine Learning for Clinical Decision Support Systems. Healthcare Informatics Review, 2022.
[6] S. Mitchell and A. Brown. Privacy Challenges in Collaborative Medical Analytics. International Journal of Data Privacy, 2023.
[7] E. Wilson and M. Taylor. Robust Aggregation Strategies for Federated Networks. Journal of Network Intelligence, 2024.
[8] L. Green. Secure Artificial Intelligence for Healthcare Applications. ACM Healthcare Computing Review, 2025.
[9] K. Thomas and B. Hall. Distributed Deep Learning under Privacy Constraints. Journal of Computational Intelligence, 2023.
[10] A. Walker. Encrypted Model Aggregation in Clinical AI Systems. International Symposium on Artificial Intelligence, 2024.
[11] N. Scott and P. Evans. Adaptive Optimization for Federated Neural Networks. Journal of Modern Computing, 2022.
[12] J. Allen. Scalable Secure Machine Learning Frameworks. Computing Science Letters, 2025.
[13] C. White and D. Brooks. Communication Reduction in Distributed Learning. IEEE International Workshop on AI Systems, 2023.
[14] M. Young. Federated Learning for Large Healthcare Organizations. Journal of Smart Health Technologies, 2024.
[15] F. Lewis and H. Cooper. Artificial Intelligence in Privacy-Sensitive Environments. International Journal of Digital Medicine, 2025.
[16] B. King. Reliable Client Selection for Federated Optimization. Machine Learning Advances, 2023.
[17] T. Clark. Compression Techniques for Distributed Neural Networks. Journal of Parallel Computing Systems, 2024.
[18] R. Turner and G. Hill. Calibration Methods for Clinical Machine Learning Models. Healthcare AI Journal, 2022.
[19] S. Parker. Secure Distributed Optimization Using Encrypted Gradients. International Journal of Applied Artificial Intelligence, 2024.
[20] D. Foster and J. Reed. Trustworthy Federated Intelligence for Medical Decision Support. Journal of Advanced AI Research, 2025.

Received 21 July 2026; revised 21 July 2026; accepted 21 July 2026