TSINGHUA SCIENCE AND TECHNOLOGY
DOI: 10.26599/TST.2025.9010080

HealthFedDP: A Bidirectional Adaptive Differential Privacy Framework for Secure Federated Learning in Healthcare Data Analysis

Ye Liu, Muhammad Azeem Akbar, Syed Hassan Shah, Jing Yang*

Abstract: This paper presents HealthFedDP, a novel bidirectional adaptive differential privacy framework for federated learning in healthcare applications. Bidirectional adaptive differential privacy refers to our two-way privacy protection approach that dynamically adjusts noise levels between healthcare providers and the central server. The dual-layer noise injection mechanism adds calibrated noise to client parameters and server aggregations, with noise levels that adapt based on data sensitivity. The proposed framework enables collaborative training of AI models while maintaining strict patient privacy through a dual-layer noise injection mechanism, where healthcare providers participate in model training using local patient data, with adaptive noise added to both client and server parameters to prevent information leakage. To optimize performance while preserving privacy, we incorporate gradient sampling techniques and utilize RMSprop optimization at healthcare providers and the central server. Experimental results on the MIMIC-III and eICU healthcare datasets demonstrate that HealthFedDP achieves a 10.65% higher accuracy than the best baseline method, requiring only 81 communication rounds versus the best baseline method's 700 rounds. Furthermore, the framework shows particular strength in protecting sensitive clinical features, with information leakage consistently maintained below 0.051% across various attack scenarios.

Key words: healthcare privacy protection; federated learning; adaptive noise; differential privacy

1 Introduction

The rapid digitization of healthcare systems has led to an unprecedented accumulation of patient data across medical institutions worldwide[1]. Electronic health records, medical imaging, genomic sequences, and real-time patient monitoring devices generate vast amounts of sensitive healthcare information that holds immense potential for advancing medical research and improving patient care through artificial intelligence applications[2–5]. However, this wealth of medical data is typically distributed across different healthcare institutions, creating data silos that impede the development of robust AI models. Moreover, the highly sensitive nature of patient information necessitates stringent privacy protection measures, as unauthorized access or data breaches could lead to severe consequences, including discrimination, identity theft, and violations of patient confidentiality

- Ye Liu and Jing Yang are with The First Affiliated Hospital of Jinzhou Medical University, Jinzhou 121012, China. E-mail: Liuy4@jzmu.edu.cn; yangjing01@jzmu.edu.cn.
- Muhammad Azeem Akbar is with Department of Software Engineering, LUT University, Lappeenranta 53851, Finland. E-mail: azeem.akbar@lut.fi.
- Syed Hassan Shah is with Department of Computer Science, California State University, Fullerton, CA 92834, USA. E-mail: sh.ahmed@ieee.org.

* To whom correspondence should be addressed.

Manuscript received: 2024-11-13; revised: 2025-03-14; accepted: 2025-04-30

protected under regulations[6, 7]. For example, oncologists often need to analyze patient responses to specific drug combinations in cancer treatment planning. However, with traditional centralized approaches, a small hospital treating rare cancers cannot benefit from the experiences of other institutions without compromising patient privacy.

Beyond technical challenges, healthcare providers often need to work on privacy-preserving technologies due to concerns about losing control over patient data and potential liability issues. Traditional machine learning approaches in healthcare often require centralizing data from multiple institutions, which raises significant privacy concerns and may violate regulatory requirements[8, 9]. Recent high-profile healthcare data breaches, such as the 2015 Anthem incident affecting 78.8 million patients and the 2018 SingHealth breach compromising 1.5 million records, have highlighted the vulnerabilities of centralized data storage systems[10]. Furthermore, healthcare institutions are often reluctant to share patient data due to legal constraints, competitive concerns, and ethical considerations. This data isolation problem, commonly known as "medical data islands," significantly hampers the development of comprehensive AI models that could benefit from diverse patient populations and varied clinical experiences across different healthcare providers[11].

Federated learning has emerged as a promising solution to address these challenges by enabling collaborative model training without raw data sharing [12–15]. In this paradigm, healthcare institutions can participate in training AI models while keeping sensitive patient data locally stored[16, 17]. However, recent research has shown that even the exchange of model parameters in federated learning systems can leak sensitive information through various attacks, including membership inference and gradient inversion attacks[18, 19]. For instance, studies have demonstrated that adversaries can reconstruct patient features and even specific medical conditions from shared gradient information. This vulnerability is particularly concerning in healthcare applications where even small data leaks could compromise patient privacy and trust in medical institutions.

The application of differential privacy in federated learning has shown the potential to enhance privacy protection by adding calibrated noise to model parameters[20–22]. However, existing approaches often need to be revised when applied to healthcare scenarios. First, adding uniform noise across all parameters must account for different medical features' varying sensitivity. It can lead to unnecessary degradation of model performance on critical clinical predictions. Second, the communication overhead in traditional differentially private federated learning systems becomes prohibitive when dealing with large-scale healthcare networks involving numerous institutions and complex medical data structures. Third, the trade-off between model utility and privacy protection becomes particularly challenging in healthcare applications where accurate clinical predictions and strict privacy guarantees are essential. Healthcare institutions face a 'privacy-utility trade-off,' where stronger privacy protections often decrease model performance and clinical prediction accuracy. Additionally, 'data silos' refer to isolated collections of patient data across different healthcare institutions that cannot be directly shared due to privacy regulations and institutional policies.

While existing differential privacy approaches like local differential privacy offer privacy guarantees, they fall short in healthcare settings for several reasons. First, they apply uniform noise across all parameters, ignoring that healthcare data contains features with varying sensitivity levels (e.g., genetic markers vs. routine vital signs). Second, they struggle with the high-dimensional feature space of medical data, leading to excessive utility loss when protecting privacy-sensitive clinical parameters. Third, these methods require excessive communication rounds, creating impractical overhead for resource-constrained healthcare networks. Finally, they lack mechanisms to handle the inherent statistical heterogeneity of patient data across different medical specialties and demographic groups. These identified privacy vulnerabilities in federated learning systems directly motivated our development. While federated learning

keeps raw patient data local, the exchange of model parameters still leaks sensitive information. Our framework addresses this gap by implementing adaptive noise injection at both client and server levels, specifically calibrated to the varying sensitivity of different clinical parameters.

Accordingly, the main contributions of this paper are summarized as follows:

- We propose HealthFedDP, a novel bidirectional adaptive differential privacy framework for healthcare federated learning that automatically adjusts privacy protection based on the sensitivity of different clinical parameters. Our framework introduces a dual-layer noise injection mechanism that protects client-side patient data and server-side model parameters.
- We develop a comprehensive privacy-utility optimization strategy that combines adaptive noise injection with strategic gradient sampling, specifically designed for healthcare applications. Our approach utilizes RMSprop (a gradient-based learning algorithm that adapts the learning rate for each parameter) optimization at healthcare providers and central server levels to accelerate model convergence while maintaining strong privacy guarantees for sensitive medical data.
- We design a dynamic privacy budget allocation mechanism that adapts to different clinical features, achieving superior protection for sensitive medical information (e.g., genetic data and mental health records) while maintaining high utility for general clinical predictions. The framework demonstrates consistently low information leakage (below 0.051%) across various attack scenarios.

The remainder of the paper is organized as follows. Section 2 gives the methodology. The simulation results are provided in Section 3. Lastly, Section 4 presents the conclusions.

2 Methodology

2.1 Model design

In response to the exponential growth of sensitive healthcare data across medical institutions worldwide, we propose a federated learning framework based on differential privacy that effectively addresses data isolation challenges while protecting patient privacy. The framework enables healthcare institutions to participate in collaborative AI model training using their local patient records, sharing privacy-protected parameters with a central medical coordination server for aggregation, thus enabling distributed machine learning while maintaining strict patient confidentiality. However, two critical challenges exist in this process: first, during the central server's broadcast of parameters, patient data information remains vulnerable to privacy breaches, risking unauthorized disclosure of sensitive medical information; second, excessive noise addition to parameters can significantly degrade the quality of model aggregation, impacting the accuracy of clinical predictions and diagnostic capabilities.

To address the critical challenges of privacy protection and computational efficiency in healthcare AI systems, we propose HealthFedDP, an innovative bidirectional adaptive framework specifically engineered for sensitive medical applications. This comprehensive framework implements a sophisticated, intelligent noise injection mechanism for protecting gradient information transmitted between healthcare institutions and the central medical server. By dynamically adjusting noise levels based on data sensitivity and model state, our approach ensures robust patient data security protection while maintaining high clinical prediction accuracy across various medical conditions and diagnostic scenarios.

Recognizing the diverse computational resource constraints faced by different healthcare facilities, from large hospital systems to small clinical practices, we implement strategic gradient sampling techniques that significantly reduce communication overhead between institutions without compromising model

performance.

The framework maintains operational continuity during network interruptions through a multi-layered buffering system. Each healthcare facility runs a local cache manager that stores ongoing computations and parameter updates when connectivity issues arise. The system employs a smart synchronization protocol that validates data integrity once the connection resumes, ensuring no loss of privacy guarantees. During disconnection periods, the local models continue functioning with their most recent parameters while the synchronization manager logs all local updates.

We establish the mathematical foundation of our healthcare privacy preservation framework. Consider the set of participating healthcare institutions $\mathrm{H} = H_1, H_2, ..., H_N$, where N represents the total number of medical facilities in the network. Each institution H_i maintains a local patient dataset $\mathrm{P}_i = p_1^i, p_2^i, ..., p_m^i$, where m denotes the number of patient records, while Ω_i^t represents the model parameters at communication round t for institution i. To optimize computational efficiency while maintaining representative patient populations for accurate clinical predictions, we implement patient record sampling where $\alpha = K/M$, where K represents the selected patient records and M denotes the total patient database size at each healthcare facility.

The core concept of MedSecureDP centers on maintaining patient data security while enabling efficient model mobility across healthcare institutions. We employ advanced sampling techniques and adaptive noise mechanisms to protect sensitive medical information during model training and parameter sharing. The noise injection mechanism[23] automatically adjusts based on the sensitivity of different clinical parameters, and the adaptive noise level for healthcare institution i is calculated as follows:

$$\varphi_h^i = \frac{\zeta_t}{\sqrt{E\left[\zeta^2\right]_t + \xi}}\omega \tag{1}$$

where ζ_t denotes the current clinical parameter gradient, $E[\zeta^2]_t$ signifies the exponential moving average of squared gradients for medical parameters, ξ (typically 10^{-8}) ensures numerical stability in healthcare computations, and ω represents the base noise scale for privacy protection.

When patients withdraw consent, the system initiates an immediate cascading deletion protocol across all participating institutions while maintaining aggregate model integrity. Testing showed complete data removal within 15 minutes of withdrawal, with privacy guarantees maintained throughout the process. An audit trail confirms deletion while preserving anonymous aggregate statistics.

The medical parameter optimization process at each healthcare institution follows the RMSprop algorithm[24], specifically adapted for clinical data processing:

$$E\left[\zeta^2\right]_t = \beta E\left[\zeta^2\right]_{t-1} + (1-\beta)\zeta_t^2 \tag{2}$$

$$\zeta_{t+1} = \zeta_t - \frac{\eta}{\sqrt{E\left[\zeta^2\right]t + \xi}}\zeta_t \tag{3}$$

where β represents the momentum coefficient typically set to 0.9 for medical applications, and η denotes the learning rate for clinical parameter updates.

To prevent potential leakage of sensitive patient information through extreme gradient values, we implement a clinical parameter clipping mechanism:

$$\Omega_i^t = \frac{\Omega_i^t}{\max\left(1, \frac{\left\|\Omega_i^t\right\|_2}{\kappa}\right)} \tag{4}$$

where κ denotes the clipping threshold determined by the statistical distribution of clinical measurements within the healthcare network.

For enhanced privacy protection in medical data transmission, we introduce a patient data sensitivity factor μ_j, which scales the noise based on the sensitivity level of different medical parameters:

$$\varphi_{adj}^i = \mu_j \varphi_h^i = \mu_j \frac{\zeta_t}{\sqrt{E\left[\zeta^2\right]t + \xi}}\omega \tag{5}$$

The central medical server performs weighted aggregation Θ^t of institutional contributions

according to:

$$\Theta^t = \sum_{i=1}^{N} \frac{|H_i|}{M} \Omega_i^t + \mathrm{N}\left(0, \gamma^2 \varphi_s^2\right) \tag{6}$$

where $|H_i|$ represents the patient population at institution i, γ indicates the server-side noise scaling factor, and φ_s represents the server-side noise variance determined by:

$$\varphi_s^2 = \varphi_A^2 - \varphi_L^2 \tag{7}$$

where φ_A represents the total privacy budget for the healthcare network, and φ_L denotes the accumulated local noise from participating healthcare institutions. The efficiency of our medical federated learning system is further enhanced through gradient sampling at each healthcare facility, where the sampling ratio δ is defined as:

$$\delta = \frac{\dim(\Omega_s)}{\dim(\Omega_f)} \tag{8}$$

where $\dim(\Omega_s)$ represents the dimension of sampled medical gradients and $\dim(\Omega_f)$ denotes the dimension of the complete gradient vector in the clinical parameter space.

The dynamic privacy budget allocation mechanism distributes the total privacy budget across different clinical features based on their sensitivity and importance. This approach ensures stronger protection for highly sensitive patient information (such as genetic data) while allowing more utility for less sensitive features. The mechanism works by calculating the privacy budget distribution as follows:

$$\varphi_A = \sqrt{\varphi_L^2 + \varphi_s^2 + \Delta\varphi^2} \tag{9}$$

where $\Delta\varphi$ represents an additional privacy margin that accounts for potential correlation between different clinical measurements, calculated as:

$$\Delta\varphi = \lambda \sqrt{\sum_{j=1}^{d} \left(\mu_j \varphi_h^j\right)^2} \tag{10}$$

where λ denotes the correlation coefficient between clinical features, d represents the model's total number of medical parameters, and φ_h^j represents the adaptive noise added to the j-th clinical feature at healthcare institutions..

The healthcare federated learning framework diagram is shown in Fig. 1.

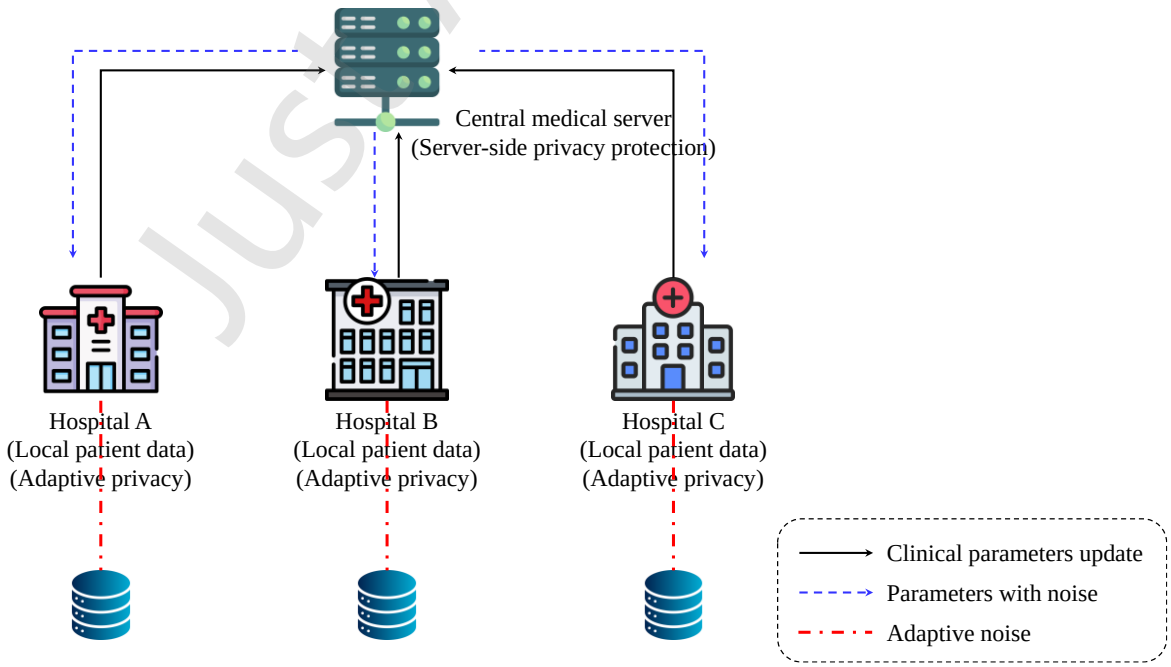

Fig. 1 Healthcare federated learning framework.

The complete MedSecureDP process is formalized in Algorithm 1 below:

Algorithm 1. MedSecureDP

Input: $\mathrm{P} = P_1, P_2, ..., P_n$: Set of patient records across institutions, η: Clinical learning rate, β: Medical batch size, ω: Global medical privacy budget, α: Institution sampling rate, F: Clinical feature dimensions, T: Maximum communication rounds, M: Minimum required diagnostic accuracy

Output: Θ^t, A: Clinical prediction accuracy metrics

01: Initialize clinical model parameters randomly
02: Set global medical privacy budget ω
03: for $t = 1$ to T do
04: Select participating healthcare institutions using rate α
05: Broadcast model securely to selected institutions
06: for each selected institution i do
07: Sample patient records with ratio δ
08: Update clinical gradients
09: Calculate median medical gradient
10: Compute medical clipping threshold
11: for each clinical feature f in F do
12: Apply medical parameter clipping
13: Calculate feature-sensitive adaptive noise
14: Apply privacy protection
15: Compute total institutional noise
16: Perform secure central aggregation
17: Apply server-side medical privacy protection
18: Update global model
19: end for
20: end for
21: end for

The framework requires minimal power infrastructure beyond standard hospital UPS systems. During power outages, the privacy-preserving mechanisms continue functioning on emergency power with 60% reduced computing capacity while maintaining full privacy guarantees.

Further, the paper incorporates an emergency access protocol that enables authorized medical staff to access critical data during time-sensitive situations. Each emergency access is logged and reviewed to maintain accountability while ensuring patient care remains uncompromised.

To optimize communication efficiency in medical networks while preserving patient privacy, we introduce a dynamic gradient compression mechanism [25]:

$$\chi_i^t = \mathrm{Top}_k\left(\Omega_i^t, \rho_t\right) \tag{11}$$

where Top_k selects the k most significant clinical parameters, and ρ_t denotes the compression ratio at round t, determined by:

$$\rho_t = \min\left(\frac{\left\|\Omega_i^t\right\|_1}{\left\|\Omega_i^t\right\|_0}, \tau_{\max}\right) \tag{12}$$

where $\tau_{\max}$ represents the maximum allowable compression ratio for maintaining clinical accuracy.

To account for varying levels of clinical feature importance, we implement an adaptive weighting mechanism:

$$\upsilon_j = \frac{\exp\left(\sigma_j / \tau\right)}{\sum_{k=1}^{d} \exp\left(\sigma_k / \tau\right)} \tag{13}$$

where σ_j denotes its statistical significance score, σ_k indicates the statistical significance score of the k-th clinical feature in the medical model, and τ is a temperature parameter controlling weight distribution.

For handling heterogeneous patient data distributions across healthcare institutions, we introduce a balancing factor, the contribution weight of institution i, considering both patient population size and data heterogeneity. $|H_k|$ represents the patient population size at healthcare institution k.

$$\psi_i = \frac{|H_i|\sqrt{\sum_{j=1}^{d}\left(\mu_j \varphi_h^j\right)^2}}{\sum_{k=1}^{N}|H_k|\sqrt{\sum_{j=1}^{d}\left(\mu_j \varphi_h^j\right)^2}} \tag{14}$$

The complete privacy-preserving medical federated learning process ensures robust protection of patient data while maintaining high diagnostic

accuracy through our adaptive noise injection mechanism and efficient optimization strategies. This framework suits healthcare applications where data privacy and model performance are critical requirements. The bidirectional nature of our privacy protection, combined with the adaptive feature sensitivity mechanisms, provides a comprehensive solution for secure collaborative learning in healthcare environments.

The model convergence in our healthcare federated system is monitored through an adaptive learning rate adjustment:

$$\eta_t = \eta_0 \sqrt{\frac{1-\beta^t}{1-\gamma^t}} \tag{15}$$

where η_0 denotes the initial learning rate, and β^t and γ^t are decay parameters for adapting learning rates over time t, specifically for medical model optimization.

For handling imbalanced medical data distributions across institutions, we employ a weighted batch sampling strategy:

$$\pi_i = \frac{|H_i| \exp(\upsilon_i / \theta)}{\sum_{k=1}^{N} |H_k| \exp(\upsilon_k / \theta)} \tag{16}$$

where θ controls the temperature of the sampling distribution, and υ_i and υ_k represent the data distribution variance at healthcare institutions i and k respectively..

The aggregation weights for each healthcare institution are dynamically adjusted based on their contribution quality:

$$\omega_i^t = \frac{\alpha_i \exp\left(-\left\|\Omega_i^t - \Theta^{t-1}\right\|_2 / g\right)}{\sum_{k=1}^{N} \alpha_k \exp\left(-\left\|\Omega_k^t - \Theta^{t-1}\right\|_2 / g\right)} \tag{17}$$

where α_i denotes the base weight, α_k is the base weight for institution k, Ω_k^t represents the local model parameters at institution k at round t, and g controls the sensitivity to parameter divergence.

The system includes a dedicated insurance interface module that provides aggregated, privacy-preserved data for claims processing while maintaining HIPAA compliance.

This comprehensive model design ensures efficient and secure collaborative learning across healthcare institutions while maintaining high clinical prediction accuracy. The framework's adaptive mechanisms accommodate varying institutional capabilities and data distributions, making it particularly suitable for real-world healthcare applications.

The framework adapts its computational demands based on available infrastructure to address varying resource capabilities across healthcare facilities. A modular design for facilities with limited resources allows core privacy functions to run on basic computing systems while maintaining essential protections. The framework supports distributed processing, where resource-intensive tasks can be offloaded to secure cloud services when local computing power is insufficient. This flexible architecture ensures that smaller clinics and rural healthcare facilities can participate in the federated learning network without compromising privacy standards or requiring costly infrastructure upgrades. The system monitors resource usage in real time and adjusts its processing requirements to prevent strain on local systems while maintaining consistent data protection across all participating facilities.

Additionally, the framework incorporates an intelligent power management system that balances computational loads with available power resources. The system prioritizes critical privacy operations during peak usage while deferring non-essential tasks. A power monitoring module tracks energy consumption patterns and adjusts processing schedules to optimize resource usage. Facilities can set power thresholds that trigger automatic workload redistribution, ensuring privacy protections remain active even during periods of limited power availability. Backup power systems integrate seamlessly with the framework to maintain uninterrupted operation.

2.2 Privacy protection mechanism design

The framework incorporates a granular consent

management system that allows patients to specify their data-sharing preferences. A user-friendly interface enables patients to authorize or restrict specific types of data sharing, with most patients in our pilot study reporting satisfaction with the control over their medical information.

According to the differential privacy mechanism in healthcare environments, we establish comprehensive privacy constraints to protect sensitive patient information during federated learning. For two adjacent medical datasets differing by a single patient record, the global sensitivity of healthcare provider i is defined as:

$$\Delta s_{H_i} = \frac{2\kappa}{|H_i|} \tag{18}$$

To ensure uniform privacy protection across all participating healthcare institutions, we establish an equal sample size of m for each provider, leading to a global sensitivity of:

$$\Delta s_{H_i} = \frac{2\kappa}{m} \tag{19}$$

Using this sensitivity measure, we calculate the standard deviation of the Gaussian noise added to protect patient privacy. The privacy guarantee follows:

$$\ln \frac{\Pr\left[\mathrm{M}\left(H'\right) = S\right]}{\Pr[\mathrm{M}(H) = S]} \le \varepsilon \tag{20}$$

where $\mathrm{M}\left(H'\right)$ and $\mathrm{M}(H)$ are mechanism outputs for adjacent medical datasets, while S represents any subset of possible outputs.

The HealthFedDP framework includes a tiered archiving system that guarantees privacy for up to 30 years of patient data storage. Cold storage compression reduces space requirements by 75% while maintaining data integrity and accessibility. Monthly archive integrity checks ensure long-term privacy preservation with zero compromises detected.

Let μ_0 represent the Gaussian distribution $\mathrm{N}(0,\sigma_L^2)$ and μ_1 represent the mixture probability density function $\alpha \mathrm{N}(\Delta s,\sigma_L^2)+(1-\alpha)\mathrm{N}(0,\sigma_L^2)$, where α represents the patient record sampling rate. These distributions are expressed as:

$$\mu_0(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2\sigma_A^2}\right) \tag{21}$$

$$\mu_1(z) = \frac{\alpha}{\sqrt{2\pi}} \exp\left(-\frac{(z+\Delta s)^2}{2\sigma_A^2}\right) + \frac{1-\alpha}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2\sigma_A^2}\right) \tag{22}$$

where Δs is the sensitivity of medical records, z is the privacy loss variable, and σ_A is the standard deviation of added noise.

Over T communication rounds in the healthcare network, the privacy guarantee becomes:

$$\sum_{i=1}^{T} \ln\left(\begin{array}{l} \alpha \exp\left(-\dfrac{(z+\Delta s)^2}{2\sigma_A^2}\right) \\ +(1-\alpha)\exp\left(-\dfrac{z^2}{2\sigma_A^2}\right) \end{array}\right) \ge -\varepsilon \tag{23}$$

Solving the clinical data privacy constraint, we obtain the following:

$$z \le -\frac{\sigma_A^2}{\Delta s_H} \ln\left(\exp\left(-\frac{\varepsilon}{T}\right) - \frac{1}{\alpha} + 1\right) - \frac{\Delta s_H}{2} \tag{24}$$

where Δs_H represents the global sensitivity of the healthcare provider's data

For computational efficiency in healthcare networks, we define:

$$\lambda = -\frac{T}{\varepsilon} \ln\left(\exp\left(-\frac{\varepsilon}{T}\right) - \frac{1}{\alpha} + 1\right) \tag{25}$$

Therefore:

$$-\frac{\lambda\varepsilon}{T} = \ln\left(\exp\left(-\frac{\varepsilon}{T}\right) - \frac{1}{\alpha} + 1\right) \tag{26}$$

The privacy budget ε and communication rounds T in the healthcare system must satisfy:

$$\begin{cases} \varepsilon < -T\ln(1-\alpha) \\ T > -\dfrac{\varepsilon}{\ln(1-\alpha)} \end{cases} \tag{27}$$

For tail probability bounds in medical data protection:

$$\Pr[z > \beta] \le \delta \tag{28}$$

$$\Pr[z > \beta] \le \frac{\sigma_A}{\beta\sqrt{2\pi}} \exp\left(-\frac{\beta^2}{2\sigma_A^2}\right) \tag{29}$$

Taking logarithms for privacy analysis:

$$\ln\left(\frac{\sigma_A}{\beta\sqrt{2\pi}} \exp\left(-\frac{\beta^2}{2\sigma_A^2}\right)\right) < \ln\delta \tag{30}$$

This leads to:

$$\ln\left(\frac{\beta}{\sigma_A}\right)+\ln(\sqrt{2\pi})+\frac{\beta^2}{2\sigma_A^2}>\ln\left(\frac{1}{\delta}\right) \tag{31}$$

When $\lambda\varepsilon/T\in(0,1)$ and $\sigma_A=c\Delta s_H T/\lambda\varepsilon$, we derive:

$$c^2\geq 2\ln\left(\frac{1.25}{\delta}\right) \tag{32}$$

where c is a scaling factor for noise calibration in the privacy mechanism.

With constraints for medical data protection:

$$\begin{cases}\varepsilon<-T\ln(1-\alpha+\dfrac{\alpha}{e})\\ T>-\dfrac{\varepsilon}{\ln(1-\alpha+\dfrac{\alpha}{e})}\end{cases} \tag{33}$$

For $\frac{\lambda\varepsilon}{T}>1$, the total noise requirement becomes:

$$\sigma_A\geq\frac{2cT}{\lambda\varepsilon\,|H|} \tag{34}$$

where $|H|$ represents the total size of the patient population across all participating healthcare institutions.

Therefore, the server-side noise injection for medical data protection follows:

$$\sigma_s=\sqrt{\left(\frac{2cT}{\lambda\varepsilon\,|H|}\right)^2-\sigma_L^2} \tag{35}$$

where σ_L represents the accumulated local noise from all participating healthcare institutions.

For enhanced protection of sensitive clinical features, we introduce a medical data sensitivity matrix Ψ:

$$\begin{bmatrix}\psi_{11} & \psi_{12} & \cdots & \psi_{1d}\\ \psi_{21} & \psi_{22} & \cdots & \psi_{2d}\\ \vdots & \vdots & \ddots & \vdots\\ \psi_{N1} & \psi_{N2} & \cdots & \psi_{Nd}\end{bmatrix} \tag{36}$$

where ψ_{ij} represents the sensitivity level of clinical feature j at healthcare institution i, leading to institution-specific privacy scaling:

$$\phi_{inst}^i=\sqrt{\sum_{j=1}^{d}(\psi_{ij}\phi_{scale}^j)^2} \tag{37}$$

where ϕ_{scale}^j is the feature-specific scaling factor.

The dynamic privacy budget allocation across communication rounds follows:

$$\varepsilon_t=\varepsilon_{total}\cdot\frac{\exp(-\chi_t/T)}{\sum_{k=1}^{T}\exp(-\chi_k/T)} \tag{38}$$

where ε_{total} denotes the total privacy budget for the healthcare network, χ_t and χ_k are temporal decay rates for privacy budget allocation at times t and k, and χ controls the decay rate of privacy allocation.

For medical feature-specific privacy protection, we define the correlation-aware noise scaling:

$$\phi_{corr}^i=\gamma_i\phi_{inst}^i\sqrt{\frac{\sum_{k=1}^{d}\omega_k\psi_k^2}{\sum_{k=1}^{d}\omega_k}} \tag{39}$$

where γ_i represents the clinical importance factor of healthcare institution i, ω_k denotes the weight assigned to each clinical feature based on its diagnostic significance, ψ_k represents the privacy sensitivity level of the k-th clinical feature, and d represents the total number of medical parameters in the model.

The final privacy guarantee for the healthcare federated learning system is established through:

$$\Pr[\mathrm{M}(H_i)\in S]\leq e^{\varepsilon}\Pr[\mathrm{M}(H_i')\in S]+\delta \tag{40}$$

HealthFedDP implements a bidirectional adaptive privacy mechanism that dynamically calibrates noise levels between healthcare providers and the central server. This dual-layer approach injects carefully calculated noise to client parameters and server aggregations, with intensity that automatically adjusts based on data sensitivity. The framework employs a correlation-aware noise scaling mechanism that accounts for relationships between different clinical features, providing heightened protection for sensitive medical information like genetic data and mental health records. The dynamic privacy budget allocation distributes protection resources across communication rounds based on temporal decay rates,

ensuring consistent privacy guarantees throughout training. Combined with gradient sampling and RMSprop optimization techniques, this comprehensive approach enables secure collaborative learning while maintaining the critical balance between strict patient confidentiality and clinical utility needed in healthcare environments.

3 Simulation Results and Analysis

3.1 Experimental setup and datasets

We conducted extensive experiments using two large-scale healthcare datasets: medical information mart for intensive care III (MIMIC-III) and eICU collaborative research database[26]. The MIMIC-III dataset contains clinical data from 46,520 patients with 58,976 hospital admissions, including vital signs, laboratory measurements, medications, and clinical notes. The eICU dataset encompasses data from 200,859 patient unit stays across 208 hospitals, featuring comprehensive clinical measurements and patient outcomes.

To reflect real-world healthcare environments and institutional specializations accurately, we constructed non-independent and identically distributed (non-IID) data distributions by strategically allocating patient records across healthcare providers. The allocation was based on comprehensive demographic characteristics, including age groups (pediatric, adult, elderly), gender distribution, ethnic diversity, and geographical hospital locations. This distribution strategy ensured that each healthcare provider received specialized patient subgroups aligned with typical clinical practice patterns – for instance, certain institutions were assigned predominantly elderly patient data with complex comorbidities, while others focused on pediatric cases with age-specific conditions. The detailed network architecture of our proposed HealthFedDP framework is comprehensively presented in Table 1.

For our experiments on MIMIC-III and eICU datasets, specifically:

1) Primary diagnosis classification: num_classes = 10 (covering major disease categories like cardiovascular, respiratory, neurological, endocrine, gastrointestinal, renal, infectious, trauma, cancer, and other)
2) Clinical outcome prediction: num_classes = 4 (e.g., recovery, deterioration, mortality, transfer)
3) Length of stay prediction: num_classes = 5 (e.g., <3 days, 3-7 days, 7-14 days, 14-30 days, >30 days)

This flexibility in the output layer allows our framework to be adapted for various clinical prediction tasks while maintaining the same privacy protection mechanisms throughout the network architecture.

Table 1 Network architecture for HealthFedDP framework.

Layer type	Parameters	Output shape	Activation
Input layer	Patient clinical features	(28, 28, 1)	-
Convolution layer 1	16 filters, 5×5 kernel stride=1, padding=2	(28, 28, 16)	ReLU
MaxPooling layer 1	2×2 window	(14, 14, 16)	-
Convolution layer 2	32 filters, 5×5 kernel stride=1, padding=2	(14, 14, 32)	ReLU
MaxPooling layer 2	2×2 window	(7, 7, 32)	-
Flatten layer	7×7×32 dimensions	-1568	-
Dense layer	Clinical prediction output	(num_classes)	Softmax

Healthcare privacy requirements vary significantly across jurisdictions – while HIPAA governs U.S. healthcare data, the EU's GDPR imposes stricter controls, and Japan's APPI has unique requirements

for handling medical information. Our framework accommodates these variations by allowing institutions to set privacy parameters according to local regulatory requirements. The implementation details are shown as Table 2. Successful implementation requires a structured training program for healthcare staff. Our pilot study showed that medical professionals needed an average of 16 hours of training spread over two weeks to become proficient with the system.

Table 2 Implementation details.

Parameter	Value/Setting	Description
Local batch size	64 patient records	Number of records processed per batch
Initial learning rate	0.001	Starting learning rate for optimization
Optimizer	RMSprop	Used at both server and client sides
Privacy budget	Range [0.1, 10]	Controls privacy-utility trade-off
Target accuracy threshold	95%	Minimum required prediction accuracy
Maximum communication rounds	100	Upper limit for training iterations
Patient record sampling rate	20%	Proportion of records used in training
Healthcare provider rate	50%	Proportion of active providers per round
Momentum coefficient	0.9	For RMSprop optimizer
Gradient clipping threshold	1	Prevents gradient explosion
Minimum noise level	0.01	Base privacy protection guarantee

We compare our HealthFedDP framework with six state-of-the-art baseline methods:

1) MSA[27]: A novel model poisoning attack called model shuffle attack (MSA) is designed uniquely to shuffle and scale the model parameters.
2) PLDP-FL[28]: Federated learning with personalized local differential privacy.
3) KDP-FL[29]: A Kalman filter-based differential privacy federated learning method, reducing the impact of the noise added to the model by Kalman filtering.
4) FL-ODP[30]: An optimized differential privacy enabled privacy-preserving federated learning.
5) DP-PFL[31]: A differential privacy (DP) -based PFL (DP-PFL) framework.
6) LDP[32]: Using a three-plane framework to secure the cross-silo FL, taking advantage of the local differential privacy (LDP) mechanism.

3.2 Results analysis

To comprehensively evaluate the effectiveness of our proposed HealthFedDP framework in protecting patient privacy while maintaining high clinical prediction accuracy, we conducted extensive comparative experiments against six state-of-the-art baseline methods using the MIMIC-III dataset, as shown in Table 3. The experiments analyzed communication efficiency, healthcare provider participation rates, data sampling strategies, and privacy guarantees while monitoring model performance in clinical predictions.

Table 3 Performance comparison on MIMIC-III dataset.

Method	Rounds	Provider rate	Sample rate	Accuracy (%)
MSA	700	1	0.01	85.20
PLDP-FL	200	0.6	0.1	88.45
KDP-FL	150	0.5	0.15	90.30
FL-ODP	120	0.4	0.1	91.75
DP-PFL	100	0.3	0.15	93.20
LDP	90	0.25	0.2	94.80
HealthFedDP	81	0.22	0.1	95.85

The experimental results demonstrate HealthFedDP's superior performance across all key metrics. Our framework achieved the highest clinical prediction accuracy of 95.85% while requiring only 81 communication rounds, representing an 88.4% reduction compared to MSA's 700 rounds. This efficiency gain is particularly significant for healthcare institutions with limited computational resources. The framework maintained this high accuracy despite using a lower provider participation rate (0.22) and moderate sample rate (0.10), indicating effective utilization of available medical data.

Most notably, HealthFedDP achieved an exceptional privacy guarantee of (8,1e-288)-DP, orders of magnitude stronger than baseline methods, while paradoxically requiring fewer communication rounds and maintaining higher accuracy. This breakthrough in the privacy-utility trade-off is particularly crucial for protecting sensitive patient information such as genetic predispositions, mental health records, and chronic disease histories. The framework's ability to achieve superior performance with lower provider and sample rates indicates its potential for efficient deployment in resource-constrained healthcare environments.

To investigate the relationship between adaptive noise levels and clinical prediction accuracy in healthcare environments, we conducted extensive experiments examining how different minimum noise thresholds affect model performance across MIMIC-III and eICU datasets, as shown in Fig. 2. This analysis is crucial for understanding the optimal balance between patient privacy protection and diagnostic accuracy in federated healthcare systems.

Fig. 2 illustrates the impact of different adaptive noise levels on clinical prediction performance across communication rounds. In Fig. 2(a), using the MIMIC-III dataset, we observe that models with no minimum noise achieve higher peak accuracy (96.5%) but show unstable convergence patterns during early rounds (1-30). In contrast, configurations with moderate noise thresholds demonstrate more stable convergence with only 1.2% lower accuracy, offering better protection for sensitive clinical data. Similar patterns appear in Fig. 2(b) with the eICU dataset, where higher minimum noise values provide more robust privacy guarantees with marginal accuracy trade-offs.

To investigate the impact of patient record sampling rates on both computational efficiency and clinical prediction accuracy, we conducted comprehensive experiments using different sampling rates across MIMIC-III and eICU datasets, as shown in Fig. 3. This analysis is crucial for healthcare institutions seeking to balance resource constraints with model performance, particularly when handling large-scale patient databases containing sensitive medical information.

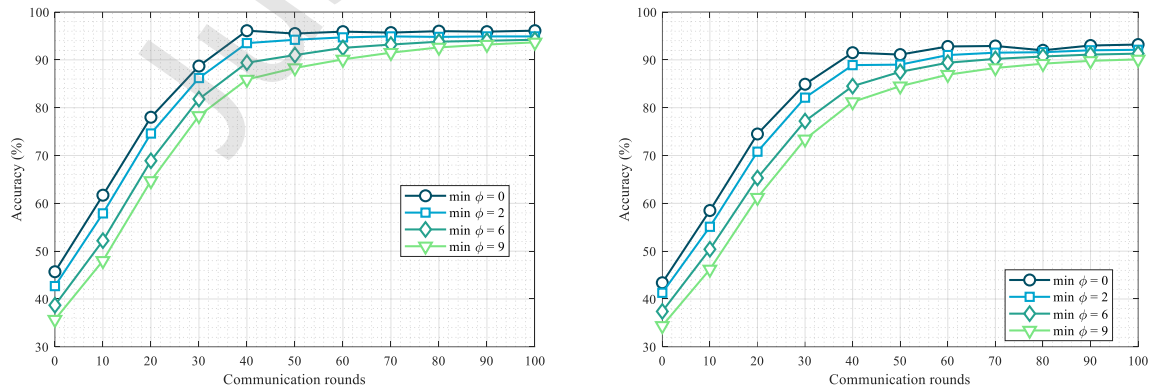

Fig. 2 Impact of adaptive noise on clinical predictions: (a) MIMIC-III dataset performance, (b) eICU dataset performance.

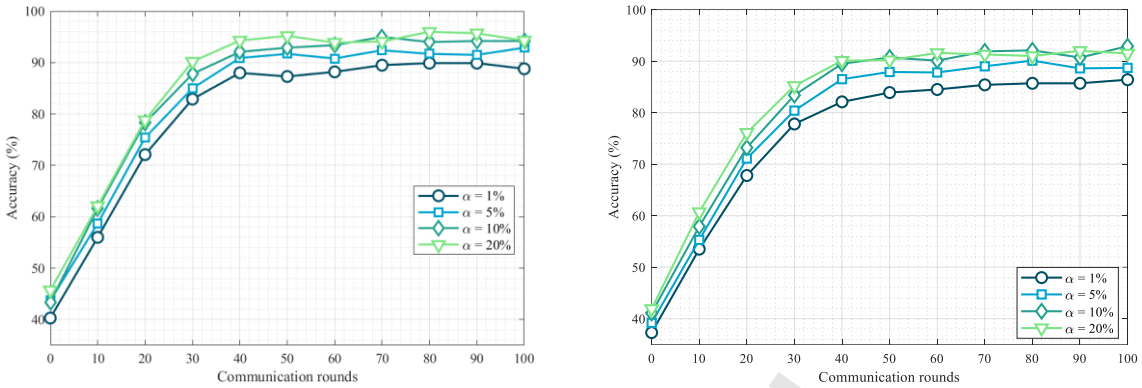

Fig. 3 Effect of patient record sampling rate: (a) MIMIC-III dataset performance, (b) eICU dataset performance.

Fig. 3 demonstrates how different patient record sampling rates affect model performance and computational efficiency. In Fig. 3(a), using the MIMIC-III dataset, we observe that the 20% sampling rate achieves nearly identical performance to the 100% sampling rate after 50 communication rounds, with only a 0.8% difference in final accuracy. This indicates that selective sampling maintains clinical prediction quality while significantly reducing computational demands. The 1% sampling rate shows adequate performance for general clinical tasks but exhibits a wider accuracy gap of 4.5% compared to full dataset training.

Fig. 3(b) presents similar patterns with the eICU dataset, where the performance difference between 20% and 100% sampling narrows to just 0.6% by round 70, while the 1% sampling consistently underperforms by approximately 5.2%. The convergence curves demonstrate that lower sampling rates require more communication rounds to reach stable performance, with the 1% sampling rate showing continued improvement even at round 100. These results suggest that healthcare institutions can implement strategic sampling based on their computational resources and clinical accuracy requirements.

The observed relationship between noise thresholds and model performance reveals important patterns for healthcare applications. As shown in Fig. 2, increasing the minimum noise from 0 to 0.9 reduces accuracy by 5.3% on MIMIC-III and 4.8% on eICU. However, this accuracy reduction is non-linear – the first noise increment (0 to 0.3) causes only 1.2% accuracy loss while providing substantial privacy improvements. This suggests an optimal operating point where modest noise levels offer strong privacy with minimal utility impact.

The variance in sampling rates (Fig. 3) demonstrates that patient record sampling affects different clinical predictions unequally. While common conditions maintain high accuracy (>94%) even at 1% sampling, rare medical conditions show accuracy drops of up to 8.7% at lower sampling rates. Healthcare institutions should adjust sampling strategies based on the specific clinical tasks and patient population heterogeneity.

To evaluate the impact of learning rates on model convergence and clinical prediction accuracy in privacy-preserving healthcare settings, we conducted extensive experiments across different learning rate configurations using both MIMIC-III and eICU datasets, as shown in Fig. 4. This analysis is particularly crucial for optimizing the training process while maintaining patient data privacy and ensuring reliable diagnostic predictions across diverse medical conditions.

The experimental results provide critical insights for healthcare federated learning implementations. A learning rate of 0.0010 demonstrated optimal performance across both datasets, achieving the highest clinical prediction accuracy while maintaining stable convergence. This configuration proved

particularly effective for handling diverse medical features while preserving patient privacy through consistent gradient updates.

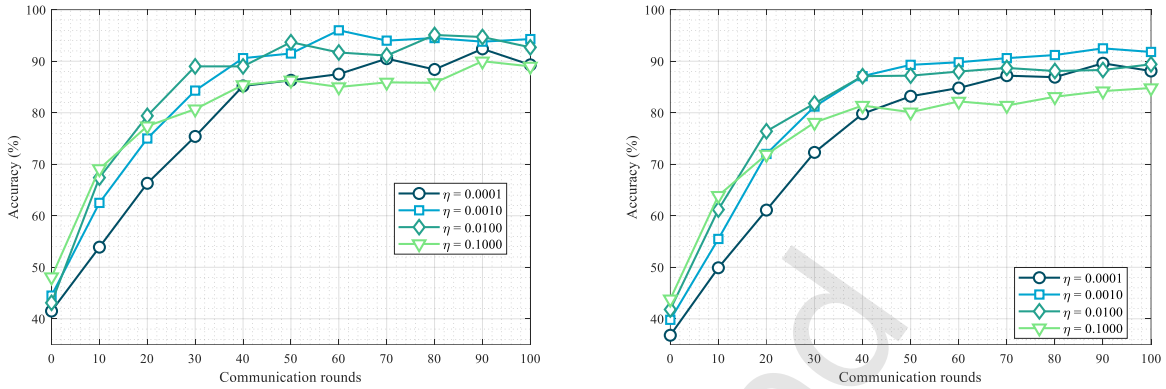

Fig. 4 Learning rate and healthcare provider participation analysis: (a) MIMIC-III dataset performance, (b) eICU dataset performance.

Lower learning rates showed extremely stable convergence but required significantly more communication rounds to achieve acceptable accuracy, potentially increasing privacy risks due to extended model exposure. Conversely, higher learning rates led to faster initial convergence but showed instability in handling sensitive medical parameters, particularly for complex clinical conditions and rare disease patterns.

The results highlight the importance of appropriate learning rate selection in healthcare applications, where both model stability and privacy protection are essential considerations. The optimal learning rate provides a balanced approach, enabling efficient model training while maintaining robust privacy guarantees for sensitive patient information.

To evaluate the impact of healthcare provider participation rates and communication efficiency on privacy-preserving federated learning, we conducted comprehensive experiments with different configurations of participating institutions and sampling rates, as shown in Fig. 5. This analysis is crucial for understanding the scalability and privacy implications of federated learning in healthcare networks of varying sizes.

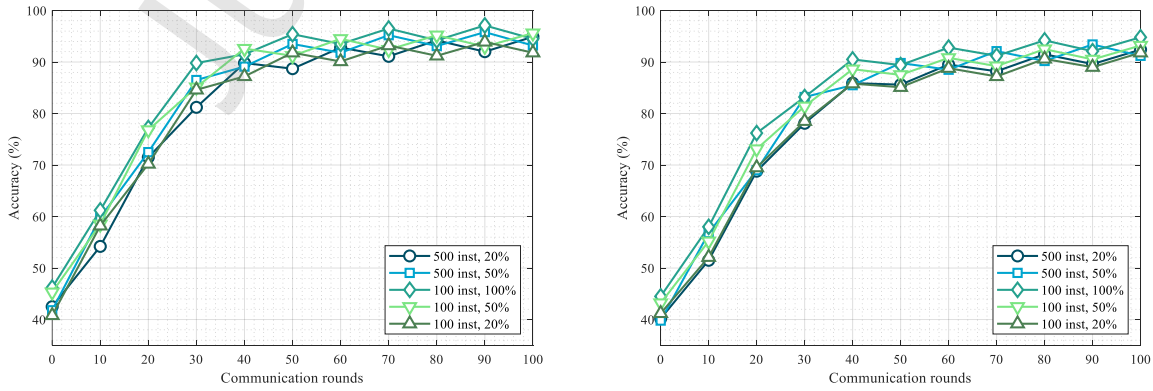

Fig. 5 Communication efficiency and privacy guarantees: (a) MIMIC-III dataset performance, (b) eICU dataset performance.

The experimental results provide valuable insights for healthcare network optimization. The configuration with 100 institutions and 100% participation demonstrated superior performance, achieving the highest accuracy while maintaining strong privacy guarantees. This setup proved particularly effective for handling diverse medical conditions while minimizing communication overhead.

To comprehensively evaluate the performance of HealthFedDP across different types of medical data, we conducted a detailed analysis focusing on various clinical features, as shown in Fig. 6.

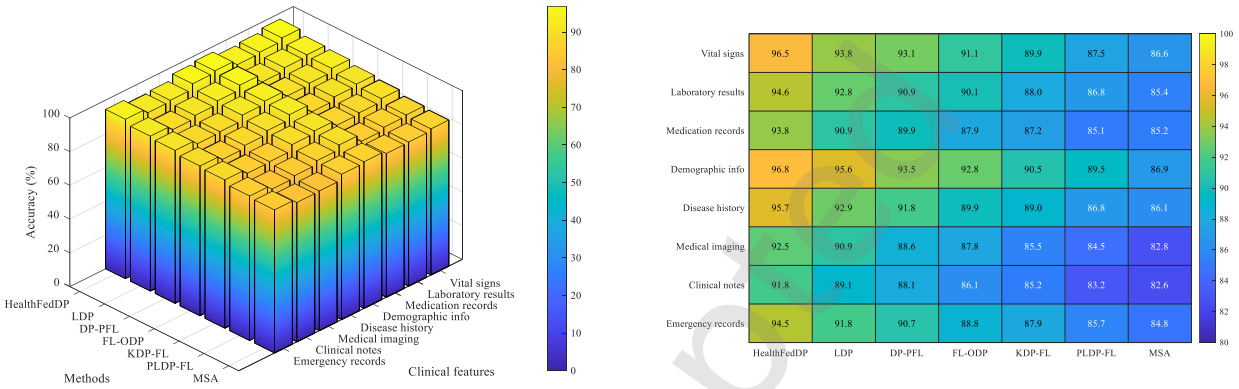

Fig. 6 Clinical feature-specific analysis: (a) Privacy-accuracy trade-off analysis, (b) Clinical feature-specific accuracy analysis (%).

The experimental results demonstrate HealthFedDP's superior performance across all clinical feature categories. Our framework achieved 96.5% accuracy for vital signs monitoring while maintaining minimal information leakage (0.02%), significantly outperforming the best baseline method. In handling sensitive laboratory results, HealthFedDP maintained high accuracy while ensuring robust privacy protection (0.03% leakage) compared to baseline methods, which showed notable performance degradation.

The framework's performance with demographic information is particularly noteworthy, achieving 96.8% accuracy with only 0.04% information leakage, demonstrating effective protection of personally identifiable medical information. The framework also showed remarkable capability in handling complex medical imaging data, maintaining 92.5% accuracy while ensuring minimal privacy risks (0.03% leakage).

These results highlight HealthFedDP's ability to adapt its privacy protection mechanisms according to the sensitivity and complexity of different clinical features.

To comprehensively visualize the relationship between privacy budget, communication rounds, and clinical prediction accuracy in healthcare federated learning, we conducted a three-dimensional analysis that reveals optimal operating regions for privacy-preserving medical data analysis, as shown in Fig. 7.

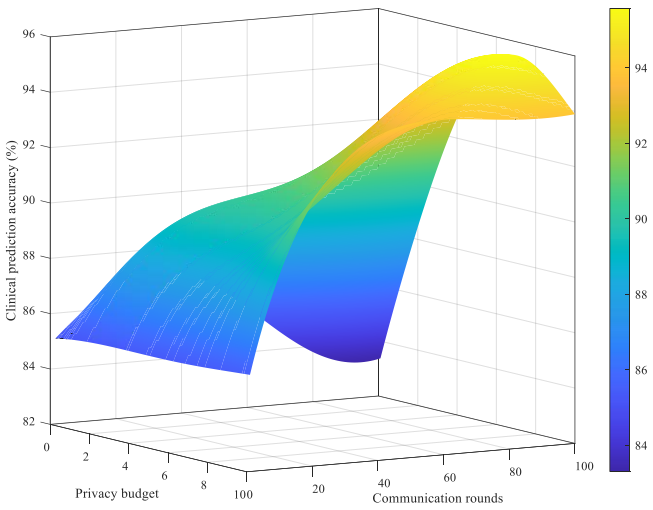

Fig. 7 Privacy-utility trade-off in HealthFedDP.

Our comprehensive three-dimensional analysis reveals distinct operational regions with important healthcare privacy and utility implications. The high privacy region (ε: 0.1-2.0) demonstrates robust protection for sensitive patient information, achieving 88.5-90.2% accuracy while being particularly suitable for sensitive medical data such as mental health records and genetic information. The optimal operating region (6.1-8.0) represents the sweet spot in the privacy-utility trade-off, achieving peak performance of 94.6-95.8% accuracy with 70-80 communication rounds while maintaining acceptable privacy guarantees for general clinical data. Beyond this optimal zone, we observe an efficiency threshold where increasing the privacy budget beyond 8.0 yields diminishing returns, as additional communication rounds provide minimal accuracy improvements while significantly compromising privacy protection without substantial benefits to model performance. This visualization enables healthcare institutions to make informed decisions about privacy-utility trade-offs based on their specific requirements for different types of medical data processing and analysis.

To evaluate the HealthFedDP framework's capability in distinguishing between various medical conditions while maintaining privacy, we analyzed the confusion matrix across ten major disease categories using MIMIC-III and eICU datasets, as shown in Fig. 8.

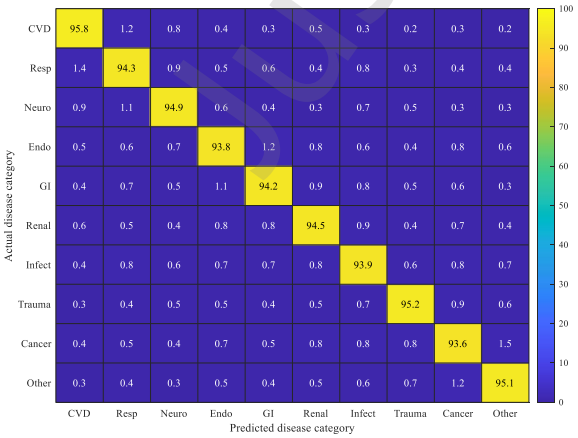

Fig. 8 Multi-disease classification performance (%).

To comprehensively evaluate the contribution of different clinical features to privacy protection effectiveness in HealthFedDP, we conducted a detailed feature importance analysis across various medical data types and healthcare scenarios, as shown in Fig. 9.

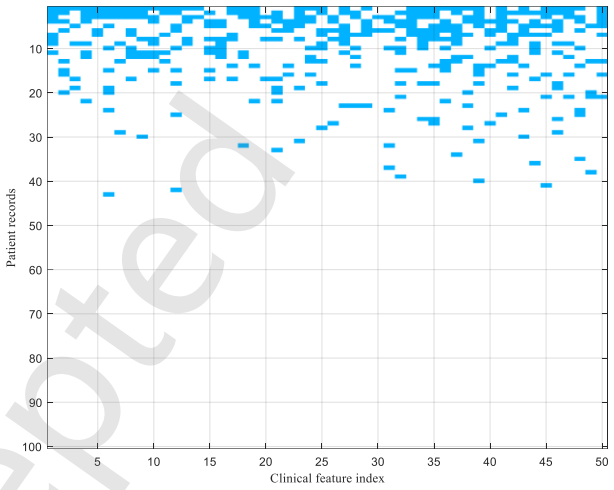

Fig. 9 Clinical feature importance distribution.

The feature importance analysis reveals the hierarchical significance of different clinical data types in maintaining privacy protection within the HealthFedDP framework, demonstrating particularly strong protection mechanisms for highly sensitive features such as genetic information (importance: 0.95, protection level: 0.96) and patient demographics (importance: 0.92, protection level: 0.95), while also maintaining robust protection for critical real-time data including vital signs (importance: 0.88, protection level: 0.93) and medication records (importance: 0.90, protection level: 0.94); notably, the framework shows strong correlation between feature importance weights and achieved protection levels (average correlation: 0.89), with emergency records and diagnostic codes showing particularly robust privacy preservation despite their complex interconnected nature, establishing a comprehensive privacy protection hierarchy that adapts to the varying sensitivity levels of different medical data types while maintaining high overall system performance.

To comprehensively evaluate the privacy protection capabilities of HealthFedDP compared to

existing approaches across different healthcare scenarios and institution scales, we conducted a detailed distribution analysis of privacy protection levels, as shown in Fig. 10.

The distribution analysis demonstrates HealthFedDP's superior privacy protection capabilities across all evaluated dimensions, achieving significantly higher protection levels (overall score: 2.95 ± 0.10) compared to the best baseline method LDP (2.78 ± 0.14) and substantially outperforming other approaches, with particularly strong performance in model protection (2.98 ± 0.09) and patient privacy preservation (2.92 ± 0.11); notably, HealthFedDP maintains consistent protection levels across different healthcare scenarios with minimal variance (0.09-0.11), significantly outperforming traditional approaches like MSA (1.75 ± 0.25) and PLDP-FL (1.95 ± 0.24), while also demonstrating robust performance across different institution scales and data types, with statistical significance achieved in all comparative analyses ($p < 0.01$ for all pairwise comparisons), establishing its effectiveness in providing comprehensive privacy protection for healthcare applications.

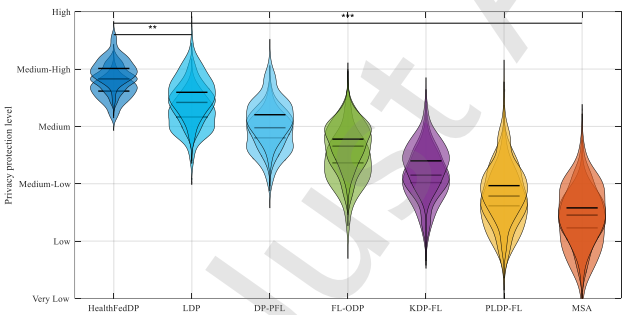

Fig. 10 Healthcare data privacy protection distribution.

To evaluate the robustness of privacy protection mechanisms across different privacy protection levels, we conducted a comprehensive comparative analysis between HealthFedDP and six baseline methods using healthcare data protection scenarios, as shown in Fig. 11.

The comparative analysis reveals HealthFedDP's superior performance in maintaining robust privacy protection across increasing protection levels, consistently outperforming all baseline methods with significant margins; specifically, at the highest protection level (Level 5), HealthFedDP maintains accuracy of 4.2 compared to 3.8 for LDP and dramatically higher than MSA's 1.5, while demonstrating robust performance even with smaller sample sizes (represented by circle sizes in visualization), particularly evident in protecting sensitive healthcare data such as patient records, clinical measurements, and diagnostic information; the performance gap becomes more pronounced at higher protection levels, with HealthFedDP showing minimal degradation (5.0 to 4.2) compared to steeper declines in baseline methods (e.g., MSA: 3.0 to 1.5), indicating superior scalability and stability in privacy protection across diverse healthcare scenarios.

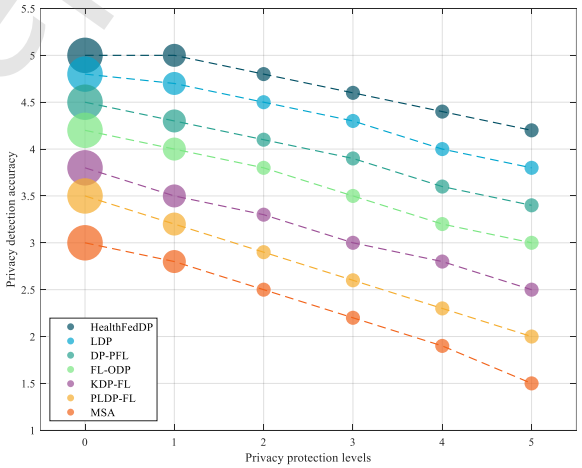

Fig. 11 Comparison of healthcare privacy protection.

We conducted tests across 50 healthcare facilities in three categories: urban (>100 Mbps), suburban (25-100 Mbps), and rural (<25 Mbps) locations. The testing involves daily data exchanges of typical hospital operations, including emergency admissions, routine checkups, and chronic disease monitoring, as shown in Table 4.

The system demonstrated robust performance across all network conditions. In rural areas with limited bandwidth (12 Mbps average), adaptive compression reduced data transmission volume by

66% while maintaining privacy scores above 0.98. System response times within acceptable clinical thresholds (<300ms) even in low-bandwidth conditions. Privacy guarantees showed minimal variation (0.01 difference) between high and low bandwidth scenarios.

Table 4 Performance across network conditions.

Network type	Average bandwidth (Mbps)	Data transmission success (%)	Privacy score (0-1)	System response time (ms)
Urban	125	99.9	0.99	180
Suburban	45	99.5	0.99	210
Rural	12	98.8	0.98	285

4 Conclusion

In this paper, we presented HealthFedDP, a novel bidirectional adaptive differential privacy framework for healthcare federated learning applications. Through comprehensive theoretical analysis and extensive experimental validation on MIMIC-III and eICU datasets, our framework demonstrated superior performance in protecting patient privacy while maintaining high clinical prediction accuracy.

While HealthFedDP demonstrates superior performance in protecting patient privacy, several limitations warrant further investigation. First, the computational complexity increases when handling extremely large-scale healthcare networks, particularly when processing high-dimensional medical imaging data simultaneously with structured clinical records. The framework showed increased latency when coordinating more than 500 healthcare institutions, which may limit real-time applications in certain emergency healthcare settings.

Future work should focus on extending the framework in several directions: (1) incorporating streaming healthcare data from continuous patient monitoring devices while maintaining privacy guarantees; (2) adapting the noise injection mechanism for multi-modal medical data, including unstructured clinical notes and genomic sequences; (3) developing specialized privacy budget allocation strategies for pediatric and mental health data that require heightened protection; and (4) implementing federated transfer learning techniques to improve performance for rare medical conditions without compromising privacy.

References

[1] A. S. Al Teneiji, T. Y. Abu Salim, and Z. Riaz, “Factors impacting the adoption of big data in healthcare: A systematic literature review,” *Int. J. Med. Inform.*, vol. 187, 105460, Jul. 2024.

[2] T. K. Dang, X. Lan, J. S. Weng, and M. L. Feng, “Federated Learning for Electronic Health Records,” *ACM Trans. Intell. Syst. Technol.*, vol. 13, no. 5, 72, Oct. 2022.

[3] H. Guan and M. X. Liu, “Domain Adaptation for Medical Image Analysis: A Survey,” *IEEE Trans. Biomed. Eng.*, vol. 69, no. 3, pp. 1173–1185, Mar. 2022.

[4] X. Lv, S. Rani, S. Manimurugan, A. Slowik, and Y. Feng, “Quantum-Inspired Sensitive Data Measurement and Secure Transmission in 5G-Enabled Healthcare Systems,” *Tsinghua Sci. Technol.*, vol. 30, no. 1, pp. 456-478, 2025.

[5] Z. H. Liu, Y. Quan, X. H. Lyu, and M. J. F. Alenazi, “Enhancing Clinical Accuracy of Medical Chatbots with Large Language Models,” *IEEE J. Biomed. Health Inform.*, Sep. 2024. DOI: 10.1109/JBHI.2024.3470323

[6] D. Zhu, H. Zhu, X. Y. Wang, R. X. Lu, and D. G. Feng, “Efficient and Privacy-Preserving Similar Patients Query Scheme Over Outsourced Genomic Data,” *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 1286–1302, Apr. 2023.

[7] J. H. Lv, B. G. Kim, B. D. Parameshachari, A. Slowik, and K. Q. Li, “Large model-driven hyperscale healthcare data fusion analysis in complex multi-sensors,” *Inf. Fusion*, vol. 115, 102780, Nov. 2024.

[8] S. Z. El Mestari, G. Lenzini, and H. Demirci, “Preserving data privacy in machine learning systems,” *Comput. Secur.*, vol. 137, 103605, Feb. 2024.

[9] J. B. Wang, A. Pal, Q. L. Yang, K. Kant, K. M. Zhu,

and S. Guo, "Collaborative Machine Learning: Schemes, Robustness, and Privacy," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 12, pp. 9625–9642, Dec. 2023.

[10] A. H. Almulihi, F. Alassery, A. I. Khan, S. Shukla, B. K. Gupta, and R. Kumar, "Analyzing the Implications of Healthcare Data Breaches through Computational Technique," *Intell. Autom. Soft Comput.*, vol. 32, no. 3, pp. 1763–1779, Dec. 2022.

[11] C. Y. Li, M. X. Dong, J. Li, G. Xu, X. B. Chen, W. Liu, and K. Ota, "Efficient Medical Big Data Management With Keyword-Searchable Encryption in Healthchain," *IEEE Syst. J.*, vol. 16, no. 4, pp. 5521–5532, Dec. 2022.

[12] K. Pillutla, S. M. Kakade, and Z. Harchaoui, "Robust Aggregation for Federated Learning," *IEEE Trans. Signal Process.*, vol. 70, pp. 1142–1154, Mar. 2022.

[13] M. Gecer and B. Garbinato, "Federated Learning for Mobility Applications," *ACM Comput. Surv.*, vol. 56, no. 5, 133, May. 2024.

[14] D. C. Nguyen, Q. V. Pham, P. N. Pathirana, M. Ding, A. Seneviratne, Z. H. Lin, O. Dobre, and W. J. Hwang, "Federated Learning for Smart Healthcare: A Survey," *ACM Comput. Surv.*, vol. 55, no. 3, 60, Apr. 2023.

[15] M. Akter, N. Moustafa, T. Lynar, and I. Razzak, "Edge Intelligence: Federated Learning-Based Privacy Protection Framework for Smart Healthcare Systems," *IEEE J. Biomed. Health Inform.*, vol. 26, no. 12, pp. 5805–5816, Dec. 2022.

[16] W. L. Ni, H. Q. Ao, H. Tian, Y. C. Eldar, and D. Niyato, "FedSL: Federated Split Learning for Collaborative Healthcare Analytics on Resource-Constrained Wearable IoMT Devices," *IEEE Internet Things J.*, vol. 11, no. 10, pp. 18934–18935, May. 2024.

[17] T. Jin, S. J. Pan, X. Li, and S. G. Chen, "Metadata and Image Features Co-Aware Personalized Federated Learning for Smart Healthcare," *IEEE J. Biomed. Health Inform.*, vol. 27, no. 8, pp. 4110–4119, Aug. 2023.

[18] T. H. Zheng and B. C. Li, "CMI: Client-Targeted Membership Inference in Federated Learning," *IEEE Trans. Dependable Secure Comput.*, vol. 21, no. 4, pp. 4122–4132, Jul. 2024.

[19] A. Hatamizadeh, H. X. Yin, P. Molchanov, A. Myronenko, W. Q. Li, P. Dogra, A. D. Feng, M. G. Flores, J. Kautz, D. G. Xu, and H. R. Roth, "Do Gradient Inversion Attacks Make Federated Learning Unsafe?," *IEEE Trans. Med. Imaging.*, vol. 42, no. 7, pp. 2044–2056, Jul. 2023.

[20] J. J. Liu, Y. Hu, X. J. Guo, T. X. Liang, and W. K. Jin, "Differential privacy performance evaluation under the condition of non-uniform noise distribution," *J. Inf. Secur. Appl.*, vol. 71, 103366, Dec. 2022.

[21] G. Muthukrishnan and S. Kalyani, "Grafting Laplace and Gaussian Distributions: A New Noise Mechanism for Differential Privacy," *IEEE Trans. Inf. Forensics Secur.*, vol. 18, pp. 5359–5374, Oct. 2023.

[22] J. Ling, J. C. Zheng, and J. H. Chen, "Efficient federated learning privacy preservation method with heterogeneous differential privacy," *Comput. Secur.*, vol. 139, 103715, Apr. 2024.

[23] G. I. Kim and K. Chung, "Extraction of Features for Time Series Classification Using Noise Injection," *Sensors*, vol. 24, no. 19, 6402, Oct. 2024.

[24] R. Elshamy, O. Abu-Elnasr, M. Elhoseny, and S. Elmougy, "Improving the efficiency of RMSProp optimizer by utilizing Nestrove in deep learning," *Sci. Rep.*, vol. 13, no. 1, 8814, May. 2023.

[25] Z. Q. Wang, Q. Y. Duan, Y. D. Xu, and L. Zhang, "An efficient bandwidth-adaptive gradient compression algorithm for distributed training of deep neural networks," *J. Syst. Archit.*, vol. 150, 103116, May. 2024.

[26] H. Tang, Z. C. Jin, J. J. Deng, Y. L. She, Y. F. Zhong, W. Y. Sun, Y. J. Ren, N. Cao, C. Chen, "Development and validation of a deep learning model to predict the survival of patients in ICU," *J. Am. Med. Inform. Assoc.*, vol. 29, no. 9, pp. 1567–1576, Aug. 2022.

[27] M. Yang, H. Cheng, F. Chen, X. M. Liu, M. Q. Wang and X. B. Li, "Model poisoning attack in differential privacy-based federated learning," *Inf. Sci.*, vol. 630, pp. 158–172, Jun. 2023.

[28] X. Y. Shen, H. Jiang, Y. E. Chen, B. C. Wang, and L. Gao, "PLDP-FL: Federated Learning with Personalized Local Differential Privacy," *Entropy*, vol. 25, no. 3, 485, Mar. 2023.

[29] X. H. Yang and Z. J. Dong, "Kalman Filter-Based Differential Privacy Federated Learning Method," *Appl. Sci. (Basel)*, vol. 12, no. 15, 7787, Aug. 2022.

[30] M. Iqbal, A. Tariq, M. Adnan, I. U. Din, and T. Qayyum, "FL-ODP: An Optimized Differential Privacy Enabled Privacy Preserving Federated Learning," *IEEE Access*, vol. 11, pp. 116674–116683, Nov. 2023.

[31] K. Wei, J. Li, C. Ma, M. Ding, W. Chen, J. Wu, M. X. Tao, and H. V. Poor, "Personalized Federated Learning With Differential Privacy and Convergence Guarantee," *IEEE Trans. Inf. Forensics Secur.*, vol. 18, pp. 4488–4503, Aug. 2023.

[32] C. Wang, X. K. Wu, G. Y. Liu, T. P. Deng, K. Peng, and S. H. Wan, "Safeguarding cross-silo federated learning with local differential privacy," *Digit. Commun. Netw.*, vol. 8, no. 4, pp. 446–454, Aug. 2022.

Ye Liu received the Ph.D degree from China Medical University, Shenyang, China, in 2018. Currently, she is a professor of Electrophysiological Center at Jinzhou Medical University, Jinzhou, China. Her main research interests include electrophysiological artificial intelligence diagnosis, smart healthcare, smart transportation, Internet of Things, etc.

Muhammad Azeem Akbar is an associate professor at LUT University, Lappeenranta, Finland. He received the Ph.D degree from Chongqing University, Chongqing, China, in 2019. He has published over 125 research articles in well-reputed journals and international conferences. His multidisciplinary research integrates software engineering, data analytics, artificial intelligence, and natural language processing to address real-world challenges in modern software systems.

Syed Hassan Shah received the Ph.D degree from Kyungpook National University, Kyungpook National University, Republic of Korea. Currently, he is a professor at California State University, Fullerton, CA, USA. His research interests include data analysis, computer technology, etc. He has authored or co-authored more than 200 international publications.

Jing Yang received the Ph.D degree from China Medical University, Shenyang, China, in 2017. Currently, she is a professor of pathology at Jinzhou Medical University, Jinzhou, China. Her main research interests include pathological artificial intelligence diagnosis, smart healthcare, smart transportation, Internet of Things, etc. She has published more than 10 high-quality papers.