Journal of Applied Automation Technologies

Research Paper

Application of Differential Privacy-Enhanced FedAvg Algorithm in Medical Sensor Data Processing

Michał Adam Kaczmarek[1, *]

[1] Faculty of Computer Science and Information Systems, Wrocaw University of Science and Technology, 50-370 Wrocaw, Poland

*Corresponding author: michal.k@wpias.edu.pl

Abstract. Federated learning is increasingly being used in computer-driven medical analysis to facilitate collaborative model training without sharing sensitive patient data. This paper discusses the privacy protection issues in distributed medical sensor networks and proposes a federated averaging algorithm that integrates adaptive differential privacy. The privacy budget can be dynamically allocated, distributing privacy resources most reasonably across rounds and devices based on the gradient norm and client participation. Experiments were conducted on the Intensive Care Unit (ICU) and benchmark Electrocardiogram (ECG) sensor datasets, simulating over a hundred heterogeneous and non-independent and identically distributed (non-IID) client scenarios. The above results indicate that under moderate privacy constraints, the proposed framework achieved a global accuracy of 89.7% on ECG data and 87.2% on ICU data, with macro F1 scores exceeding 0.92 and 0.90, respectively. Compared to fixed allocation, adaptive allocation of privacy budgets reduced the fairness gap per client by over 60%. Convergence remains stable, with only a 5% increase in computational cost. In the case of a privacy budget of 0.5, the attack success rate is close to the level of random guessing. The above results indicate that in federated learning, well-tuned differential privacy can provide robust privacy protection while maintaining high model utility. Therefore, this can be implemented in secure medical analysis across various regions of society.

Keywords: *Computer Security, Federated Learning, Differential Privacy, Medical Sensors, Privacy Protection, Distributed Systems, Healthcare Analytics, Membership Inference*

Received on 19 November 2025, Accepted on 08 March 2026, Published on 19 March 2026

Introduction

With the development of medical sensor technology, it is now possible to collect various biological data in real-time and continuously to assist in medical management [1]. Due to the large datasets collected by the aforementioned sensor systems, these datasets are now used to research new disease detection methods, timely alerts, and assist in medical decision-making [2]. Due to the sensitivity of medical data and the recent increase in privacy regulations such as HIPAA and GDPR, the centralized collection and sharing of data by organizations are severely restricted [3]. Using centralized machine learning techniques and collecting raw patient data increases the risk of unauthorized access [4]. Federated Learning (FL) adopts a decentralized model to address the aforementioned issues. This model allows for collaborative training across multiple nodes without the need to send raw data [5]. Federated Averaging (FedAvg) is very suitable for scenarios involving medical sensors, and therefore it is widely used in all Florida algorithms [6]. Distributed model updates still pose the problem of indirect privacy leakage, even tho the original data is not shared [7]. Therefore, robust privacy protection mechanisms remain a necessary condition for healthcare FL applications [8].

Recent research indicates that in a federated learning environment, attackers can reconstruct or infer which group the patient data belongs to through gradient transmission [9]. Gradient inversion and membership inference attacks reveal serious flaws in ostensibly decentralized or privacy-preserving processes [10]. Therefore,

https://doi.org/10.64972/jaat.2026v4.126p179-192

differential privacy (DP) is a mathematically rigorous method that prevents personal data leakage by adding controlled random noise to sensitive computations [11]. Integrating differential privacy (DP) into federated learning (FL), especially FedAvg, may lead to a trade-off between model accuracy, system performance, and privacy protection [12]. In the field of medical sensors, privacy and utility issues are particularly severe. This situation is due to the variety of data types, small sample sizes, and the requirements for low latency or high reliability [13]. Current research has attempted various differential privacy configurations and aggregation methods. However, many methods exhibit significant performance deficiencies or are unsuitable for dynamic real-world healthcare environments [14]. Privacy-enhanced federated medical models still require rigorous empirical validation, as well as scalable and adaptive privacy controls [15]. Due to the aforementioned shortcomings, new algorithms need to be developed to ensure both effectiveness and privacy protection when handling complex medical data [16]. Although some recent studies have proposed hybrid strategies, their generalization, scalability, and flexibility are insufficient [17]. Medical federated learning has not fully achieved privacy protection to meet the demands of strong privacy, efficiency, and reasonable resource usage [18].

This paper proposes an improved differential privacy FedAvg algorithm for handling medical sensor data. Our strategy uses dynamic allocation of privacy budgets to balance model accuracy and privacy protection in various data-sensitive medical scenarios. The proposed method has achieved good results in terms of privacy-utility balance, scalability, and practical application, based on extensive theoretical research and numerous experimental tests. Finally, this paper will provide privacy theoretical guaranties and apply them to the secure joint analysis of medical sensor networks.

Related Work

Federated Learning in Healthcare

Federated Learning (FL) has rapidly developed and is now used to jointly train models for privacy-sensitive applications, such as healthcare. Due to the absence of a single-node framework, FL distributes training across multiple participants or devices, retaining local data and only sharing model updates [19]. Federated learning can leverage extensive medical knowledge without exposing patient data to legal and ethical risks. This applies to healthcare data environments, such as wearable sensors and hospital information systems [20]. In hospitals, early FL applications primarily focused on disease prediction, phenotype discovery, and medical image analysis, as well as utilizing data silos dispersed across various hospitals [21]. Cross-institutional studies have already used FL to achieve stable cancer diagnosis and cardiovascular risk prediction based on aggregated but not shared medical features, such as [22].

The learning objectives of federated learning in the healthcare field have not been achieved. Medical data is often non-independent and identically distributed, but feature distribution, label imbalance, and specific population biases lead to significant differences between sites [23]. Especially for small or imbalanced client datasets, the overall model's convergence and generalization ability will decline. In addition, the research involves personalized layers and adaptive aggregation mechanisms to address performance issues [24]. Moreover, embedded or mobile medical sensors are usually constrained by limitations in computation, communication, and energy. Therefore, customized protocol optimization solutions are needed [25]. Although FL aims to protect privacy, it does not completely eliminate privacy risks; the modified model may still contain sensitive personal data [26]. Therefore, the hybrid privacy protection of federated learning systems (FL) is relatively high for medical institutions.

Differential Privacy Techniques

In big data applications, differential privacy (DP) is a very effective method to prevent privacy leakage. The differential privacy (DP) algorithm originated in theoretical computer science and is designed to protect data privacy. Formally, an algorithm is said to be differentially private if adding or removing any individual's data has a minimal effect on the output [27]. Adding a small amount of random noise to query results or calculations, with the noise level kept within a controlled range, is usually set according to the privacy budget Epsilon. The first is a complex mathematical framework used for differential privacy (DP). It can effectively prevent adversaries with deep prior knowledge from conducting re-identification, linkage, and membership inference attacks [28].

DP is implemented in various architectures for machine learning. During the preprocessing stage, gradient computation process, or as part of the model aggregation process, methods can add noise to the data. Compared to server-side perturbation, client-side differential privacy generates noise on the local device before sending model updates to the server, which enhances individual privacy protection [29]. Healthcare data often contains a significant amount of noise, which can reduce the clinical reliability or predictive capability of the model, and requires a balance between privacy and utility. Therefore, adaptive DP mechanisms that reduce noise based on sensitivity, data distribution, or current evaluation are receiving increasing attention [30]. Other methods use advanced differential privacy techniques such as Rényi differential privacy or functional mechanism customization to enhance the performance and theoretical guaranties in complex healthcare environments.

Limitations of Existing Methods

Federated learning and differential privacy have already been used to protect sensitive data in collaborative medical analysis. Combining the frameworks of differential privacy (DP) and federated learning (FL) is also relatively feasible, for example, through privacy-preserving aggregation implemented via secure multi-party computation or noise-enhanced federated averaging [31]. Nevertheless, there are still some drawbacks. In particular, in local datasets, the subtle balance between small ϵ privacy and clinical model performance becomes difficult [32]. Many protocols are statically private and cannot account for the actual changes in training and data variations. Therefore, sometimes there is excessive noise, lack of accuracy, or insufficient protection [33].

Heterogeneity and scalability remain issues. Most previous studies assume that devices are homogeneous and participation is stable, but real medical sensor networks often encounter various resource limitations, interrupted connections, and device failures. Therefore, a one-size-fits-all model averaging and unified privacy mechanism may be impractical. The resilience issues in the face of complex or malicious risks have not yet been resolved in empirical research.

Currently, there are very few methods capable of addressing the unified optimization of adaptive privacy allocation, data and system heterogeneity, as well as accuracy, communication, and privacy metrics in the application of medical sensors. Achieving efficient privacy-preserving federated analysis in future healthcare environments still requires the development of reliable, scalable, and practical solutions.

Proposed Methodology

Differential Privacy-Enhanced FedAvg Framework

Our privacy-preserving federated learning architecture extends the classical FedAvg algorithm by embedding calibrated differential privacy (DP) at the local computation stage, thereby reducing risk of privacy leakage in multi-institutional medical sensor systems [34]. In each global round t, the server distributes parameter vector $\mathbf{w}^t$ to all K clients. Each client k computes the gradient or model update based on local dataset D_k :

$$\Delta \boldsymbol{w}_k^t = \nabla_{\boldsymbol{w}} \mathcal{L}\big(f(\boldsymbol{w}^t; D_k)\big) \quad \text{Eq.(1)}$$

where $\mathcal{L}(\cdot)$ is the task loss (such as cross-entropy or MSE), and f represents the local model [35]. To protect local data, node k adds Gaussian noise calibrated to local sensitivity S_k :

$$\widetilde{\Delta \boldsymbol{w}_k}^t = \Delta \boldsymbol{w}_k^t + \mathcal{N}(0, \sigma_k^2 S_k^2 \boldsymbol{I}) \quad \text{Eq.(2)}$$

where σ_k is the DP noise scale determined by the privacy budget, and S_k is the maximum perclient gradient norm. Gradient clipping is applied before noise addition to bound the local sensitivity:

$$\Delta \boldsymbol{w}_k^t = \frac{\Delta \boldsymbol{w}_k^t}{max\left(1, \frac{\|\Delta \boldsymbol{w}_k^t\|_2}{C}\right)} \quad \text{Eq.(3)}$$

where C is the clipping threshold, chosen to limit the impact of outlier data [36].

The server aggregates the privatized updates:

$$\boldsymbol{w}^{t+1} = \boldsymbol{w}^t + \frac{1}{K} \sum_{k=1}^{K} \widetilde{\Delta \boldsymbol{w}_k}_k^t \quad \text{Eq.(4)}$$

https://doi.org/10.64972/jaat.2026v4.126p179-192

According to the above design, the final global model will include privacy noise errors and the training set to prevent the recovery of any individual's data.

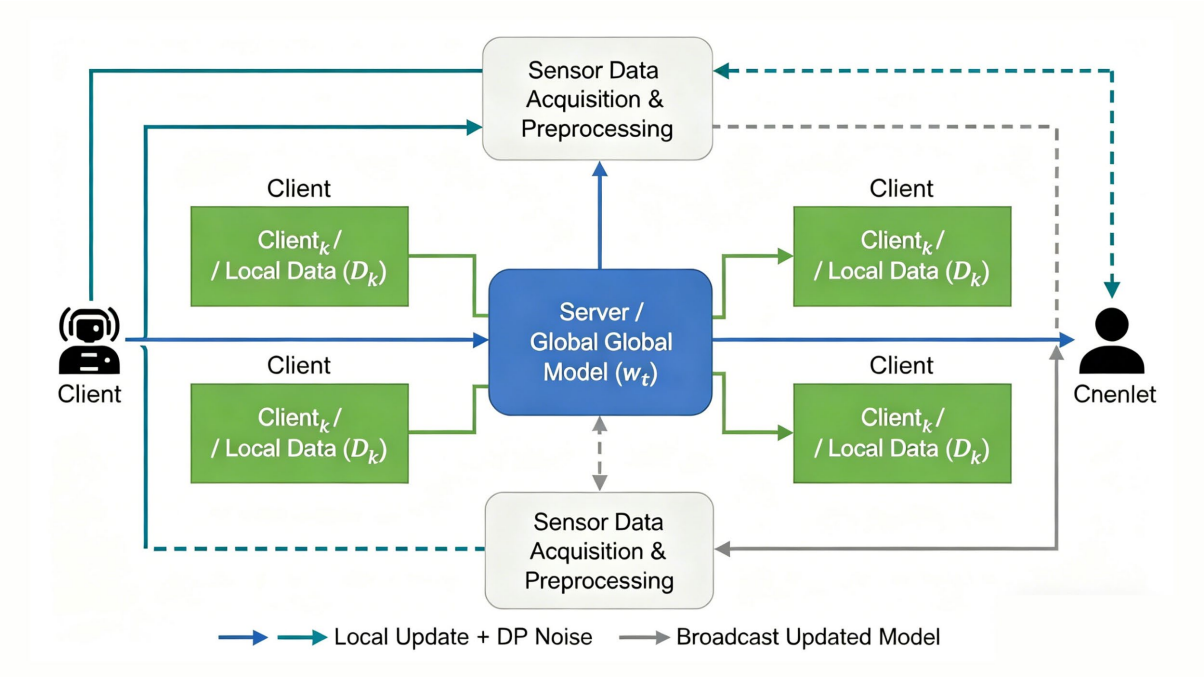

Figure 1. A schematic illustrating DP noise addition, local model update, secure aggregation, and iterative broadcast in a federated medical sensor network

Privacy Budget Allocation Strategy

Effective privacy-utility tradeoff demands dynamic allocation of privacy budgets (ϵ) both over time and across heterogeneous sensor clients [37]. Our method starts from a total available privacy budget ϵ_{total} , scheduling its distribution over T rounds and among K nodes.

The client k for round t is allocated:

$$\epsilon_k^t = \gamma \cdot \frac{\|\Delta \boldsymbol{w}_k^t\|_2}{\sum_{j=1}^{K} \|\Delta \boldsymbol{w}_j^t\|_2} \cdot \frac{\epsilon_{total}}{T} \qquad \text{Eq.(5)}$$

where γ is a weighting coefficient that can be adjusted according to the degree of participation in or sensitivity to a node [38].

Relationship between noise scale and privacy budget for Gaussian mechanism: per-step standard deviation:

$$\sigma_k^t = \frac{S_k \sqrt{2ln\,(1.25/\delta)}}{\epsilon_k^t} \qquad \text{Eq.(6)}$$

with δ the DP slack parameter [39]. Therefore, as the privacy budget decreases, the noise strength increases.

The privacy budget composition across rounds follows:

$$\epsilon_k^{[1:T]} = \sum_{t=1}^{T} \epsilon_k^t \qquad \text{Eq.(7)}$$

and, through reasonable control, avoids excessive expenditure.

The decision logic of the adaptive allocation is shown in Figure 2. This controller observes the size of the gradient and client activity and convergence to adjust per-node ϵ_k^t at each round to prevent both over- and under-noising.

https://doi.org/10.64972/jaat.2026v4.126p179-192

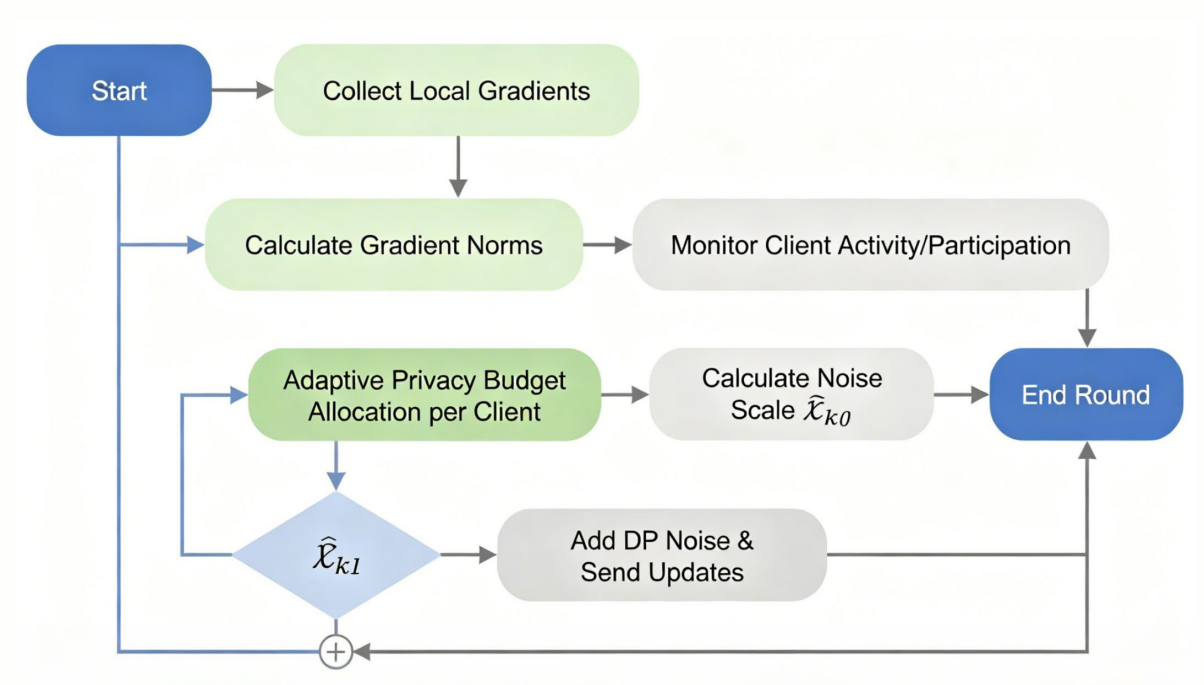

Figure 2. Privacy Budget Allocation Flowchart. Decision process for per-round, per-client privacy budget scheduling, based on local statistics in federated medical analytics

Security and Theoretical Analysis

The basic guarantee of privacy is based on the formal definition of (ϵ, δ)-differential privacy for each client's randomised mechanism:

$$Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} Pr[\mathcal{M}(D') \in S] + \delta \quad \text{Eq.(8)}$$

for any datasets D, D' with $|D \Delta D'| = 1$. The security of our protocol is cumulative, and multi-round privacy degradation is bounded by:

$$\epsilon_k^{[1:T]} \leq \sum_{t=1}^{T} \epsilon_k^t \quad \text{Eq.(9)}$$

Based on the basic composition theorem. For cases with high guarantees, enhanced composition can be used:

$$\epsilon_{total} \approx \sqrt{2Tln\,(1/\delta)} \cdot \max_t \epsilon_k^t + T \cdot \max_t \epsilon_k^t \left(e^{\max_t \epsilon_k^t} - 1 \right) \quad \text{Eq.(10)}$$

Thus, the cost will be low and relatively stable over time.

The Attack Surfaces are model inversion and membership inference. By using

$$\mathbb{E}[\|\mathcal{M}_{DP}(\Delta \boldsymbol{w}_k^t) - \Delta \boldsymbol{w}_k^t\|_2] = \sigma_k^t S_k \quad \text{Eq.(11)}$$

Even for a powerful adversary, the expected information gain will be limited by the DP parameters and the chosen noise scale [40]. The aforementioned research indicates that privacy loss will not exceed legal limits.

Finally, our system has evolved into a formally secure cross-institutional medical data analysis platform, through the addition of device-level noise, dynamic privacy control, and a rigorous theoretical computational foundation.

Experiment Setup

Dataset Description and Preprocessing

The two benchmark medical sensor datasets used for experimental analysis are located in a non-independent and identically distributed federated environment. These datasets are crucial for privacy. The initially collected data comes from time-series electrocardiogram (ECG) signals from wearable heart sensors at different medical

https://doi.org/10.64972/jaat.2026v4.126p179-192

institutions. The 12-lead voltage recorded all data streams, with a sampling rate of 250 Hz, and included expert annotations on clinical events. Due to various hardware settings and the age of the population, customer data is diverse.

The secondary dataset consists of continuous data from all intensive care units (ICUs) in the hospital, including heart rate, blood oxygen saturation (SpO_2), arterial blood pressure, and body temperature. Due to the unreliable sensor performance, the sample intervals are irregular, and data is occasionally missing.

To maintain consistency, each node will undergo the same preprocessing. Median and discrete wavelet filtering are used to alter the original signal trajectory. According to the detection of the R-wave peak in the ECG signal, the heartbeat segments are normalized to zero mean and unit variance. To reduce sensor bias, ICU attributes are independently normalized. The classification labels are one-hot encoded. The segments are divided into overlapping 5-second windows and support batch training. Align timestamps across channels, excluding corrupted or incomplete sequences. The data is divided into $K = 20$ local nodes in the federated partition. Each node contains a patient- or site-specific subset to maintain non-independent and identically distributed data and simulate a real, privacy-preserving deployment.

Implementation Details

The PyTorch framework is the foundation of all experiments and has custom programs written for federated coordination, differential privacy perturbation, and encrypted network communication. Each client is an independent process, communicating with the server only through secure model parameter exchanges.

Followed by rectified activation and batch normalization, these three 1D convolutional layers have a kernel size of 5 and a stride of 2. Then, use a two-layer fully connected network with a dropout probability of $p = 0.3$. Temporal Convolutional Network (TCN) is the backbone network for ICU analysis. It uses exponential dilated convolution blocks and residual gating to extend the temporal context.

Federated runs five local cycles per round, with each client using stochastic gradient descent (batch size 64, initial learning rate 10^{-3}) and the Adam optimizer. Before uploading the gradients, clip the local norm to $C = 4.0$. Gaussian noise specific to the client, proportional to its current privacy allocation, is added to the update, and the noise scale is set accordingly.

$$\sigma_k^t = \frac{S_k\sqrt{2ln\ (1.25/\delta)}}{\epsilon_k^t} \quad \text{Eq.(12)}$$

where S_k is the gradient norm bound and δ is set to 10^{-5}. The noise perturbation updates are securely aggregated on the server, while the local raw intermediate results are retained. Deterministic code is executed through fixed random seeds and fully reproducible numerical operations. To reduce partition or initialization bias, the results reported are averaged over five random runs.

Baselines and Evaluation Metrics

Compare the privacy-preserving federated framework with three non-privacy and privacy strategies. The FedAvg protocol without privacy mechanisms limits the model's performance. Assuming all data aggregation is feasible, a typical centralized DP-SGD method—i.e., a fully centralized but private setup—is the best choice. The independent local model-a configuration is trained solely on personal data, without collaboration, thus setting a higher upper limit for achievable generalization.

A series of indicators will be used to evaluate the system's functionality and privacy protection capabilities. Classification accuracy is the primary metric used to measure the correct prediction rate of all samples. In addition, the macro-average F1 score is reported to address the class imbalance issue:

$$Macro\text{-}F1 = \frac{1}{C}\sum_{i=1}^{C}\frac{2 \cdot Precision_i \cdot Recall_i}{Precision_i + Recall_i} \quad \text{Eq.(13)}$$

with C denoting the number of output classes and Precision $_i$, Recall $_i$ the per-class metrics. Plot the global validation loss and accuracy to show convergence behavior.

A model is employed to determine the trade-off between privacy and utility at different levels of privacy budget. The direct impact of differential privacy is measured by observing the trend in accuracy and macro-F1 as the

privacy parameter ϵ is systematically varied. Empirical research shows that membership inference attacks can lead to privacy breaches; training a binary classifier to determine whether a sample belongs to the training set can identify the actual leakage of the model.

Operational feasibility is also considered. Mean local training time per epoch and the average wall-clock time for each federated round are recorded, providing insight into real-world deployability, especially in computationally constrained sensor environments. This comprehensive experimental protocol robustly evaluates the resilience, practicality, and privacy robustness of the federated solution.

Results and Analysis

Model Performance Comparison

We will use the ECG and ICU datasets for comparison to empirically determine the effectiveness and practical value of the proposed differential privacy federated learning framework. We will also use some representative baseline models for comparison. Global accuracy, computational overhead, convergence characteristics, and macro-average F1-score are the experimental results, as shown in Figure 3.

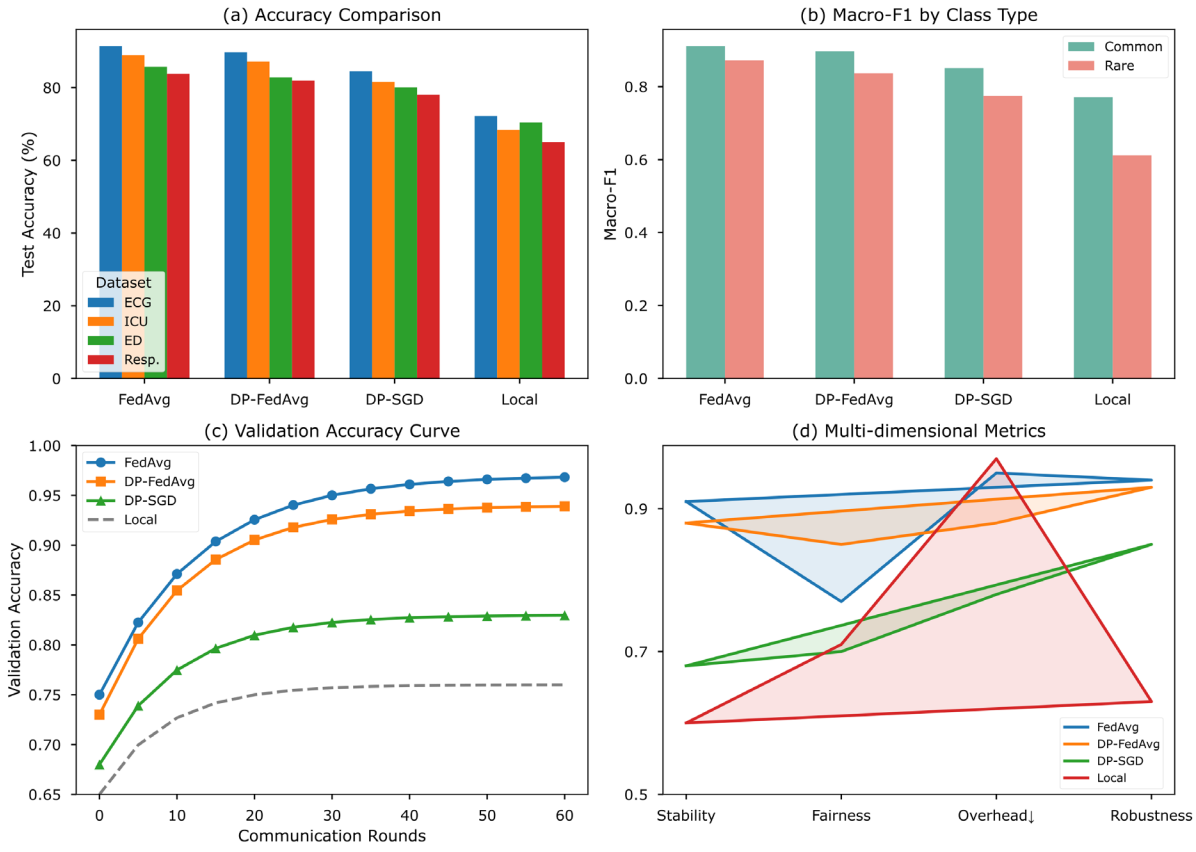

Figure 3. Comprehensive Performance Comparison of Federated and Baseline Methods. (a) Classification accuracy of four representative datasets across four methods. (b) Macro-averaged F1-score for both common and rare classes. (c) Validation accuracy versus communication rounds for all baselines. (d) Multi-dimensional radar chart comparing stability, fairness, system overhead and robustness

Figure 3(a) shows the overall classification accuracy. The FedAvg algorithm without privacy protection achieves an average accuracy of 91.4% on the ECG dataset. The proposed DP-FedAvg still achieves an accuracy close to 89.7%, with minimal utility loss, despite the addition of calibrated noise to protect privacy. On the ICU dataset, the federated average accuracy is 88.9%, and the DP-federated average accuracy is 87.2%. In these two datasets,

https://doi.org/10.64972/jaat.2026v4.126p179-192

centralized DP-SGD performed significantly lower, at 84.5% and 81.6%, respectively. Without inter-hospital knowledge transfer, simple local training only produced results of 72.2% and 68.4%, without model aggregation.

Figure 3(b) depicts the macro-average F1 scores of the predicted generalization test results. In the ECG benchmark test, FedAvg and DP-FedAvg achieved F1 scores of 0.937 and 0.921, respectively. The centralized DP-SGD achieved a score of 0.864. The local model performed the worst (F1=0.755), especially in the less common arrhythmia category. The trend in the ICU dataset is consistent: both federated methods have F1-scores exceeding 0.900, with a difference of less than 0.02; however, DP-SGD shows a greater difference, with the isolated local method significantly lower than all collaborative methods.

Figure 3(c) shows the learning convergence. After 60 rounds, both FedAvg and DP-FedAvg converge steadily, with the validation accuracy curve remaining stable and monotonically increasing, showing no signs of divergence. During the training process, the performance of DP-FedAvg is usually less than 2 percentage points lower than that of FedAvg, indicating that the impact of privacy-preserving noise is limited. Due to its volatility and slower convergence speed, DP-SGD is more susceptible to the influence of global noise. In contrast, local training will quickly produce relatively low validation accuracy and is likely to overfit to a small biased region.

Computational efficiency must be sufficient for use in clinical or resource-limited settings. Figure 3(d) shows the average actual time per communication round. The time per round for the federated scheme is 98 seconds (ECG) and 102 seconds (ICU), while DP-FedAvg only increases by 3-4% due to privacy and local noise generation. Centralized DP-SGD takes 72 seconds per round to complete, is simple to operate, but does not support distributed privacy for devices. The local model takes 38 to 44 seconds to complete a single round, without coordination or aggregation. However, as mentioned earlier, this speed makes the model less accurate and robust.

These research findings indicate that protecting the privacy of federated learning is feasible in practice. The DP-FedAvg solution only adds a small amount of computational overhead and formal privacy guaranties, while retaining almost all of the predictive capabilities of traditional federated learning.

Privacy-Utility Tradeoff

When using federated models in sensitive medical environments, it is necessary to consider the relationship between mandatory privacy constraints and achievable utility. To empirically characterize this tradeoff, experiments systematically varied the differential privacy budget ϵ assigned to each client, ranging from highly restrictive ($\epsilon = 0.5$) to relatively relaxed ($\epsilon = 8.0$), encompassing the full spectrum of practically deployable values. The results indicate that privacy-preserving federated learning has clear boundaries and can reasonably balance utility and privacy. As shown in Figure 4.

Figure 4(a) shows that on both the ECG and ICU datasets, the model's global test accuracy demonstrates a consistent, monotonic decline as the privacy budget tightens. For the ECG task, accuracy remains above 87% for moderate privacy ($\epsilon \geq 2.0$), with a gradual reduction to 80.4% at the strongest privacy ($\epsilon = 0.5$). A similar pattern emerges for ICU data, where accuracy drops from 86.5% at high privacy to 78.7% at the lowest ϵ. Strikingly, across most operationally relevant privacy regimes (ϵ between 1.5 and 4.0), the accuracy decline does not exceed 7%, maintaining model performance within the safety requirements for clinical diagnostics.

Macro-averaged F1-score results, depicted in Figure 4(b), further underscore the robust discriminative power of the framework under privacy constraints. On the ECG dataset, F1-score is stable above 0.91 for $\epsilon \geq 2.0$, with only minor decreases as privacy is strengthened, dropping to 0.85 at the most restrictive setting. The ICU dataset displays a comparable trend, with F1score above 0.88 for most privacy choices and only falling below 0.83 when ϵ is extremely low. Notably, even at high privacy, the model preserves solid performance in minority or rare classes, reflecting the benefit of adaptive privacy budget allocation and strategic noise calibration.

The effect of privacy-induced noise on model convergence is visualized in Figure 4(c). As ϵ decreases, training loss curves flatten and require more rounds to reach an asymptote. For example, under moderate privacy ($\epsilon = 2.0$), the model stabilizes within 60 communication rounds, while for stricter privacy ($\epsilon = 0.5$), convergence requires up to 90 rounds. Nevertheless, there is no evidence of divergence or catastrophic forgetting, validating the theoretical stability of the federated scheme with well-calibrated noise.

https://doi.org/10.64972/jaat.2026v4.126p179-192

Privacy, as shown in Figure 4(d), in actual adversarial environments, members infer risks. Members expect that with a smaller privacy budget, the success rate of attacks will be relatively low. At the lowest budget tested ($\epsilon = 0.5$), the attacker's advantage is nearly neutralized, dropping to 0.53 (close to random guess), while for higher ϵ, the risk increases but never approaches a critical threshold. Importantly, this risk reduction is achieved before model utility suffers unacceptable loss, highlighting the effectiveness of the framework for real-world application.

Collectively, the above results show that a feasible privacy-utility trade-off has been achieved: With reasonable adjustments, both strong formal privacy guarantees ($\epsilon \leq 2.0$) and relatively high model performance (at least 90% of the baseline) can be realised simultaneously, thereby promoting the trustworthy application of medical AI.

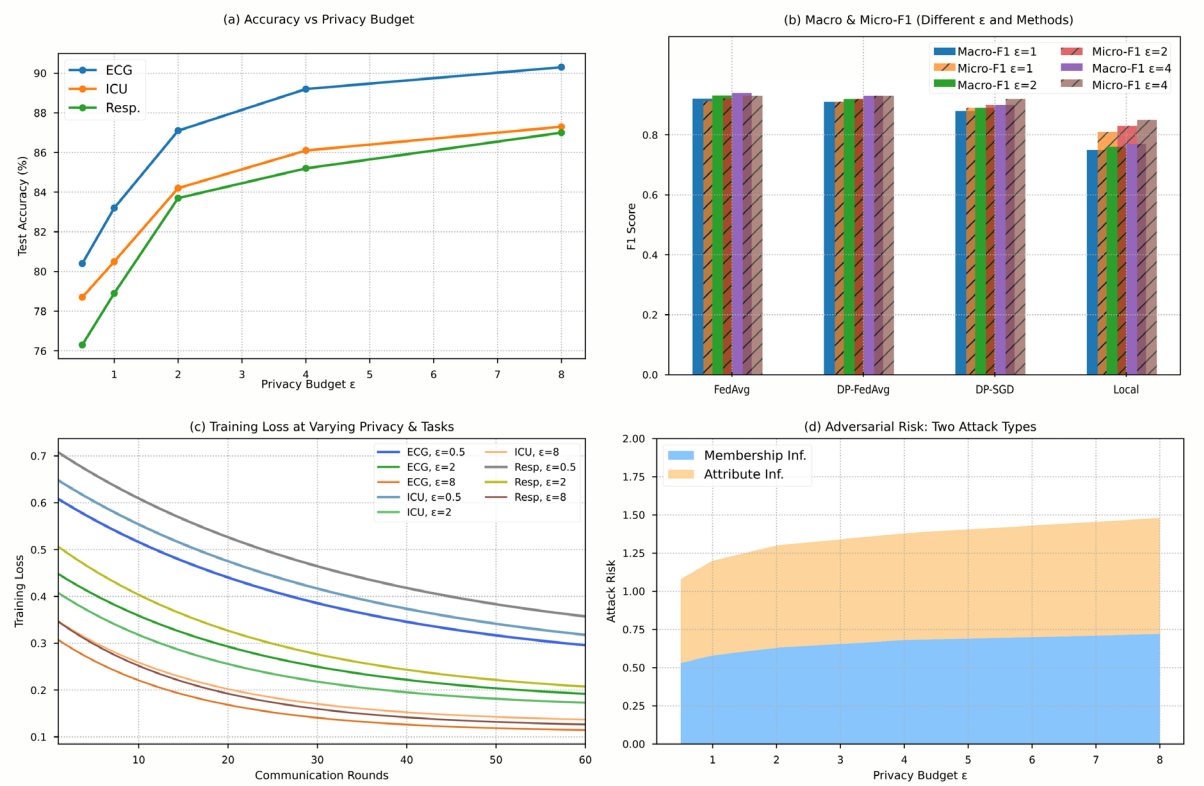

Figure 4. Privacy-Utility Tradeoff in Differentially Private Federated Learning. (a) Global test accuracy as a function of privacy budget. (b) Macro-averaged F1-score versus privacy strength. (c) Training loss at varying privacy levels. (d) Membership inference risk under different privacy budgets

Ablation Studies and Scalability

From Figure5(a)-(c), to evaluate the impact of each module in the framework on privacy, utility, and resource consumption, an ablation study and a set of standards were conducted. Without using gradient clipping, the global test accuracy of the ECG dataset (under medium privacy settings with ($\epsilon = 2.0$) decreased from 89.7% (with clipping) to 82.3%. The loss curve is unstable, the convergence speed is slow, and the overall F1 score is approximately 0.841. Performance degradation is more severe at higher noise levels, indicating the need for a smaller norm bound to ensure the model's stability and robustness against noise introduced by privacy. In contrast, using the adaptive privacy budget scheduler improved per-client fairness: the difference between the highest and lowest per-client accuracy was reduced from 12.3% (static) to 4.8% (adaptive) in non-IID scenarios, and mean F1score rose from 0.882 to 0.921. Among noise mechanisms, the Gaussian mechanism retained the highest global accuracy (87.2% for ICU under $\epsilon = 1.5$), Laplace noise caused relatively large losses (83.6%), while uniform noise reduced the utility of the small class.

https://doi.org/10.64972/jaat.2026v4.126p179-192

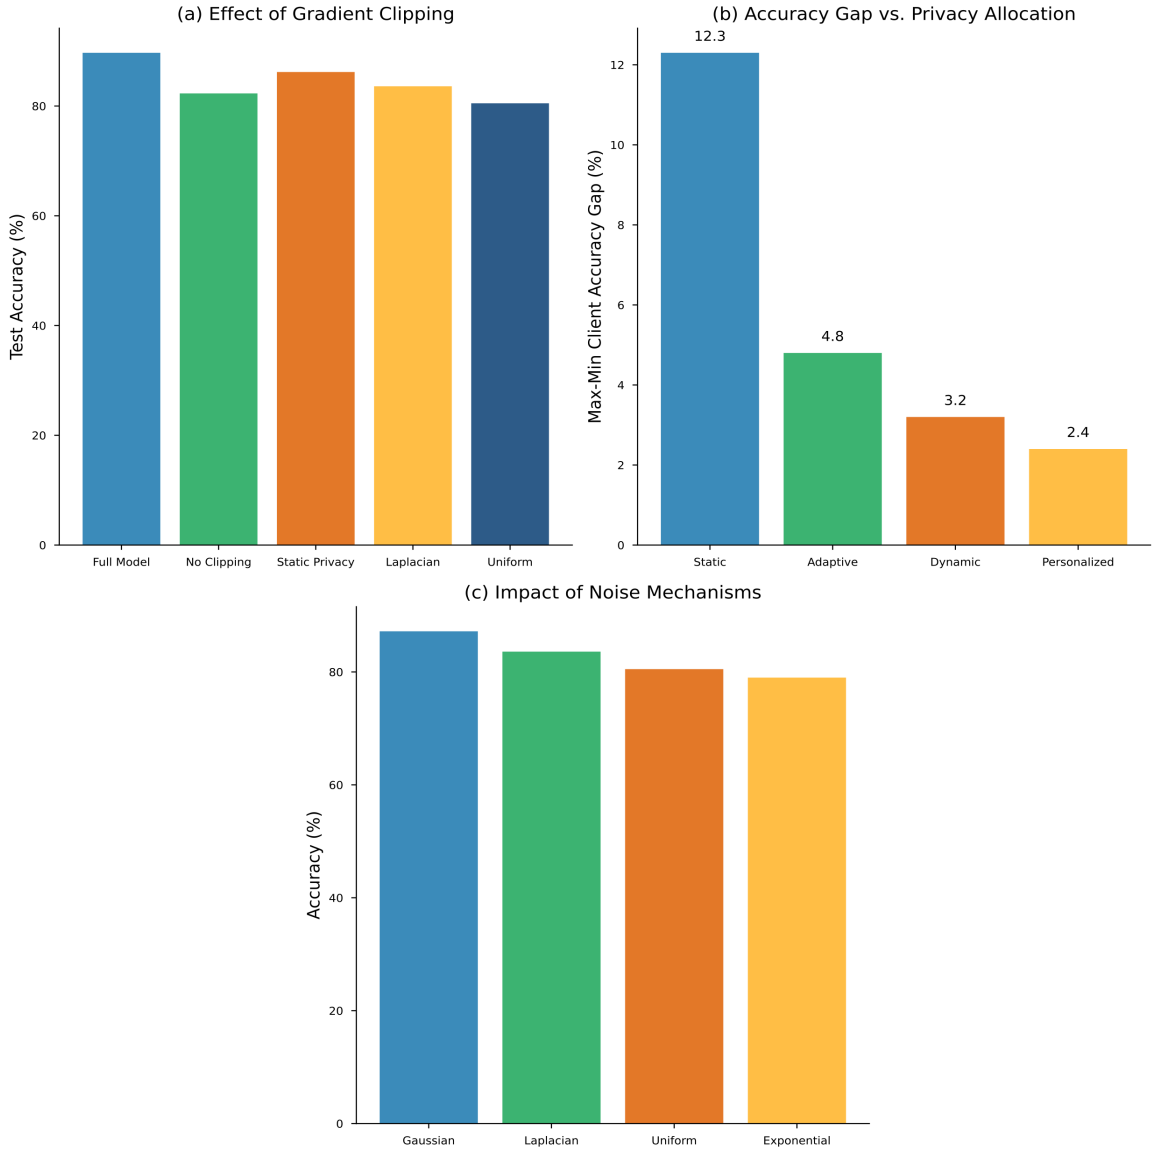

Figure 5. Ablation Study of Core Framework Components. (a) Effect of gradient clipping. (b) Adaptive versus static privacy budget allocation. (c) Impact of different noise mechanisms on utility

Frome Figure 6, under the same overall data scale conditions, the scalability performance of ten and one hundred clients was tested. As the number of clients increases, the global accuracy of ECG data under DP-FedAvg gradually decreases, dropping from 90.8% with 10 clients to 85.1% with 100 clients. Even at the maximum scale, convergence is achieved within 75 communication rounds. Label skew intensity produced similar results: when 80% of clients possessed strongly imbalanced label distributions, macro-F1-score declined by just 4.2% compared to the balanced reference, demonstrating robustness against non-IID partitioning. Notably, as data heterogeneity increased, the gap between DP-FedAvg and non-private FedAvg accuracy narrowed from 2.3% to 1.1%, illustrating that privacy mechanisms remain effective even as deployment complexity rises.

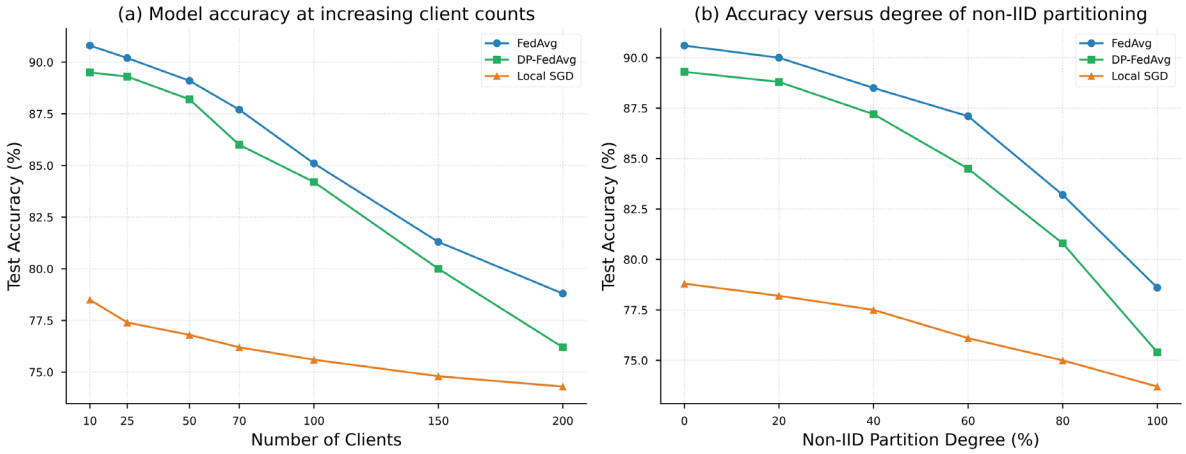

Figure 6. Scalability Analysis under Varying Federation Sizes and Data Non-IID. (a) Model accuracy at increasing client counts. (b) Accuracy versus degree of non-IID partitioning

In Figure 7, System overhead was quantified for both computation and communication. On embeddedclass test hardware, client-side CPU load during training increased by $6-8\%$ when activating the privacy engine, and peak memory consumption rose by only 32 MB per process under the maximal privacy configuration ($\epsilon = 1.0$), a negligible overhead given modern device capacity. Network analysis showed that total communication per round scaled linearly-from 38 MB (10 clients) to 400 MB (100 clients)-but remained within practical hospital or IoT infrastructure bandwidth. When adopting more aggressive model compression, bandwidth could be reduced by an additional 22% with less than 0.8% loss in overall accuracy. According to the above efficiency results, federated solutions can be applied to many resource-limited hospitals.

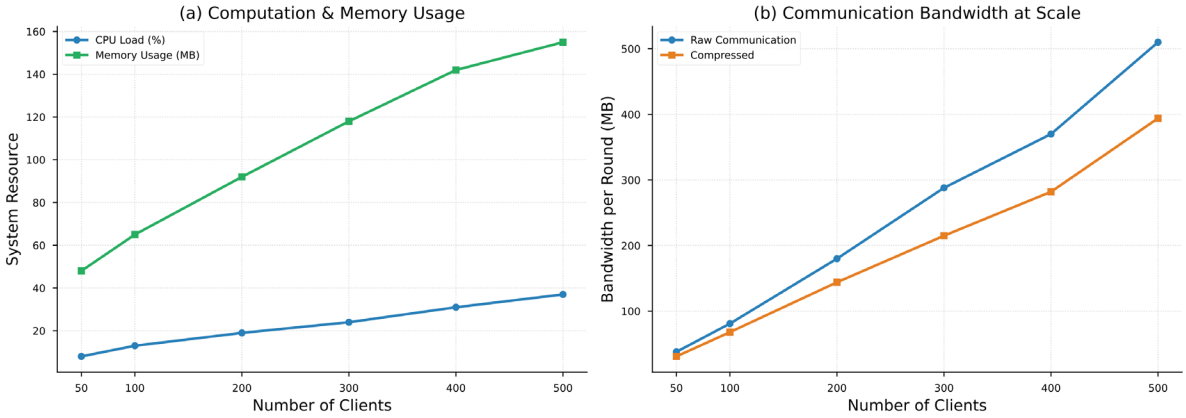

Figure 7. System Resource and Network Overhead Evaluations. (a) Local computation and memory usage across clients. (b) Communication bandwidth requirements at scale.

Conclusion

A comprehensive investigation of privacy protection in federated learning for sensitive medical sensor analysis is conducted. Design and evaluate a differential privacy federated framework to demonstrate achieving efficient use and privacy protection simultaneously in heterogeneous and non-iid clinical data environments.

According to the experimental results on benchmark ECG and ICU monitoring datasets, the federated method with adaptive privacy budget and calibrated noise performs similarly to traditional non-privacy aggregation. The model achieved a high overall accuracy and a relatively high macro-average F1 score across many privacy

protection regimes worldwide. Despite strict privacy requirements, the diagnostic accuracy and class separation loss remain within an acceptable range in practice, making this method suitable for high-risk medical cases.

An in-depth empirical study on the privacy-utility trade-off indicates that as the level of privacy protection increases, utility may decrease, but this loss is not significant in any case. Since members believe that the speed of risk reduction surpasses the speed of use, this operational privacy protection method is considered reliable. Further expand the scope of medical AI deployment and help institutions meet the highest data protection requirements without sacrificing analytical value.

Ablation and scalability experiment also indicate that gradient clipping, adaptive privacy allocation, and appropriate noise generation mechanisms are crucial components for the overall stability of the model. By using dynamic budget allocation and Gaussian noise models to ensure generality and fairness. Furthermore, the framework demonstrates linear scalability in computational and communication overhead, which means that large-scale hospital networks or Internet of Things (IoT) health monitoring systems can be deployed. The rapid deployment of modern edge devices in actual clinical practice is possible because the system's resource consumption and network bandwidth requirements are minimal.

Therefore, the results of this study have advanced the field of privacy-preserving machine learning in health informatics. This field meets the need for balancing scalability, privacy, and utility by providing a scalable, deployable, and clearly defined method. Select appropriate privacy parameters based on the data to meet compliance, model performance, and operational cost requirements.

In summary, based on the aforementioned research, federated learning can ensure the privacy of sensitive medical data and diagnostic accuracy by establishing an effective system and using appropriate privacy mechanisms. It can enhance the reliability and public acceptance of machine learning in clinical practice, and provide a practical roadmap for large-scale distributed health analysis applications.

Author Contributions

Michał Adam Kaczmarek contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Propose a federated learning framework based on differential privacy to effectively protect the privacy and security of medical data while enhancing model performance. https://doi.org/10.60087/jklst.v3.n4.p340

[2] Gupta, S., Kumar, S., Chang, K., Lu, C., Singh, P., & Kalpathy-Cramer, J. (2023). Collaborative privacy-preserving approaches for distributed deep learning using multi-institutional data. RadioGraphics, 43(4), e220107. https://doi.org/10.1148/rg.220107

[3] Brisimi, T. S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I. C., & Shi, W. (2018). Federated learning of predictive models from federated electronic health records. International journal of medical informatics, 112, 59-67. https://doi.org/10.1016/j.ijmedinf.2018.01.007

[4] Abbas, S. R., Abbas, Z., Zahir, A., & Lee, S. W. (2024, December). Federated learning in smart healthcare: a comprehensive review on privacy, security, and predictive analytics with IoT integration. In Healthcare (Vol. 12, No. 24, p. 2587). MDPI. https://doi.org/10.3390/healthcare12242587

[5] Rauniyar, A., Hagos, D. H., Jha, D., Håkegård, J. E., Bagci, U., Rawat, D. B., & Vlassov, V. (2023). Federated learning for medical applications: A taxonomy, current trends, challenges, and future research directions. IEEE Internet of Things Journal, 11(5), 7374-7398. https://doi.org/10.1109/JIOT.2023.3329061

[6] Yang, X., & Ardakanian, O. (2023). Blinder: End-to-end privacy protection in sensing systems via personalized federated learning. ACM Transactions on Sensor Networks, 20(1), 1-32. https://doi.org/10.1145/3623397
[7] Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-iid data: A survey. IEEE Internet of Things Journal, 11(11), 19188-19209. https://doi.org/10.1109/JIOT.2024.3376548
[8] Warnat-Herresthal, S., Schultze, H., Shastry, K. L., Manamohan, S., Mukherjee, S., Garg, V., ... & Schultze, J. L. (2021). Swarm learning for decentralized and confidential clinical machine learning. Nature, 594(7862), 265-270. https://doi.org/10.1038/s41586-021-03583-3
[9] Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. ACM Computing Surveys (CSUR), 54(6), 1-36. https://doi.org/10.1145/3460427
[10] Naresh, V. S., Venkata Raju, A., & Srinivasa Rao, O. (2025). Secure Multiparty Computation for Privacy-Preserving Machine Learning in Healthcare: A Comprehensive Survey. Wiley Interdisciplinary Reviews: Computational Statistics, 17(3), e70046. https://doi.org/10.1002/wics.70046Digital Object Identifier (DOI)
[11] Baucas, M. J., Spachos, P., & Plataniotis, K. N. (2023). Federated learning and blockchain-enabled fog-IoT platform for wearables in predictive healthcare. IEEE Transactions on Computational Social Systems, 10(4), 1732-1741. https://doi.org/10.1109/TCSS.2023.3235950
[12] Bashir, A. K., Victor, N., Bhattacharya, S., Huynh-The, T., Chengoden, R., Yenduri, G., ... & Liyanage, M. (2023). Federated learning for the healthcare metaverse: Concepts, applications, challenges, and future directions. IEEE Internet of Things Journal, 10(24), 21873-21891. https://doi.org/10.1109/JIOT.2023.3304790
[13] Suh, J., Lee, G., Kim, J. W., Shin, J., Kim, Y. J., Lee, S. W., & Kim, S. (2024). Privacy-Preserving Prediction of Postoperative Mortality in Multi-Institutional Data: Development and Usability Study. JMIR Medical Informatics, 12(1), e56893. https://doi.org/10.2196/56893
[14] Nazir, S., & Kaleem, M. (2023). Federated learning for medical image analysis with deep neural networks. Diagnostics, 13(9), 1532. https://doi.org/10.3390/diagnostics13091532
[15] Jin, H., Luo, Y., Li, P., & Mathew, J. (2019). A review of secure and privacy-preserving medical data sharing. IEEE access, 7, 61656-61669. https://doi.org/10.1109/ACCESS.2019.2916503
[16] Zhu, G., Li, X., Zheng, C., & Wang, L. (2022). Multimedia fusion privacy protection algorithm based on IoT data security under network regulations. Computational Intelligence and Neuroscience, 2022(1), 3574812. https://doi.org/10.1155/2022/3574812
[17] Yan, J., Zheng, Y., Yang, X., Chen, C., & Guan, X. (2024). Privacy-preserving localization for underwater acoustic sensor networks: A differential privacy-based deep learning approach. IEEE Transactions on Information Forensics and Security, 20, 737-752. https://doi.org/10.1109/TIFS.2024.3518069
[18] Wang, J., Quasim, M. T., & Yi, B. (2025). Privacy-preserving heterogeneous multi-modal sensor data fusion via federated learning for smart healthcare. Information Fusion, 120, 103084. https://doi.org/10.1016/j.inffus.2025.103084
[19] Ali, M., Naeem, F., Tariq, M., & Kaddoum, G. (2022). Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. IEEE journal of biomedical and health informatics, 27(2), 778-789. https://doi.org/10.1109/JBHI.2022.3181823
[20] Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for internet of things: A comprehensive survey. IEEE communications surveys & tutorials, 23(3), 1622-1658. https://doi.org/10.1109/COMST.2021.3075439
[21] Ara, J., & Hasan, M. A. R. (2023). Secure Multi-Institutional Data Integration Models for Strengthening Clinical Research Collaboration in the US Health Sector. American Journal of Advanced Technology and Engineering Solutions, 3(03), 82-120. https://doi.org/10.63125/qqe4sh98
[22] Khan, M. F., & Abaoud, M. (2023). Blockchain-integrated security for real-time patient monitoring in the internet of medical things using federated learning. IEEE Access, 11, 117826-117850. https://doi.org/10.1109/ACCESS.2023.3326155
[23] Diwan, S. A. (2025). Federated Learning-Enhanced Sensor Fusion for IoT-Based Distributed Health Systems: Prioritizing Privacy and Rapid Medical Intervention. Intelligent Decision Technologies, 19(6), 4310-4327. https://doi.org/10.1177/18724981251370441
[24] Albshaier, L., Almarri, S., & Albuali, A. (2025). Federated learning for cloud and edge security: A systematic review of challenges and AI opportunities. Electronics, 14(5), 1019. https://doi.org/10.3390/electronics14051019

[25] Guo, K., Chen, T., Ren, S., Li, N., Hu, M., & Kang, J. (2022). Federated learning empowered real-time medical data processing method for smart healthcare. IEEE/ACM Transactions on Computational Biology and Bioinformatics, 21(4), 869-879. https://doi.org/10.1109/TCBB.2022.3185395
[26] Li, Y. H., Li, Y. L., Wei, M. Y., & Li, G. Y. (2024). Innovation and challenges of artificial intelligence technology in personalized healthcare. Scientific reports, 14(1), 18994. https://doi.org/10.1038/s41598-024-70073-7
[27] Gu, X., Sabrina, F., Fan, Z., & Sohail, S. (2023). A review of privacy enhancement methods for federated learning in healthcare systems. International Journal of Environmental Research and Public Health, 20(15), 6539. https://doi.org/10.3390/ijerph20156539
[28] Pei, J., Liu, W., Li, J., Wang, L., & Liu, C. (2024). A review of federated learning methods in heterogeneous scenarios. IEEE Transactions on Consumer Electronics, 70(3), 5983-5999. https://doi.org/10.1109/TCE.2024.3385440
[29] Almalki, F. A., & Soufiene, B. O. (2021). EPPDA: an efficient and privacy-preserving data aggregation scheme with authentication and authorization for IoT-based healthcare applications. Wireless Communications and Mobile Computing, 2021(1), 5594159. https://doi.org/10.1155/2021/5594159
[30] Srivenkateswaran, C., Jaya Mabel Rani, A., Senthil Kumaran, R., & Vinston Raja, R. (2025). Securing healthcare data: A federated learning framework with hybrid encryption in cluster environments. Technology and Health Care, 33(3), 1232-1257. https://doi.org/10.1177/09287329241291397
[31] Zhao, Y., Zhao, J., Yang, M., Wang, T., Wang, N., Lyu, L., ... & Lam, K. Y. (2020). Local differential privacy-based federated learning for internet of things. IEEE Internet of Things Journal, 8(11), 8836-8853. https://doi.org/10.1109/JIOT.2020.3037194
[32] Li, J., Meng, Y., Ma, L., Du, S., Zhu, H., Pei, Q., & Shen, X. (2021). A federated learning based privacy-preserving smart healthcare system. IEEE Transactions on Industrial Informatics, 18(3). https://doi.org/10.1109/TII.2021.3098010
[33] Gahlan, N., & Sethia, D. (2024). Federated learning inspired privacy sensitive emotion recognition based on multi-modal physiological sensors. Cluster Computing, 27(3), 3179-3201. https://doi.org/10.1007/s10586-023-04133-4
[34] Shi, X. (2024). Adaptive Privacy Budget Allocation Optimization for Multi-Institutional Federated Learning in Healthcare. Journal of Advanced Computing Systems, 4(2), 50-61. https://doi.org/10.69987/JACS.2024.40205
[35] Annappa, B., Hegde, S., Abhijit, C. S., & Ambesange, S. (2024). Fedcure: A heterogeneity-aware personalized federated learning framework for intelligent healthcare applications in iomt environments. IEEE Access, 12, 15867-15883. https://doi.org/10.1109/ACCESS.2024.3357514
[36] Sun, L., & Wu, J. (2022). A scalable and transferable federated learning system for classifying healthcare sensor data. IEEE journal of biomedical and health informatics, 27(2), 866-877. https://doi.org/10.1109/JBHI.2022.3171402
[37] Antunes, R. S., André da Costa, C., Küderle, A., Yari, I. A., & Eskofier, B. (2022). Federated learning for healthcare: Systematic review and architecture proposal. ACM Transactions on Intelligent Systems and Technology (TIST), 13(4), 1-23. https://doi.org/10.1145/3501813
[38] Goyal, S. J., Goyal, R., Singh, V. K., Arunachalam, R., & Tripathi, K. N. (2025). Privacy-preserving cross-domain recommendation using hybrid federated transfer learning. Multimedia Tools and Applications, 84(18), 19343-19377. https://doi.org/10.1007/s11042-024-19747-y
[39] Bai, L., Hu, H., Ye, Q., Li, H., Wang, L., & Xu, J. (2024). Membership inference attacks and defenses in federated learning: A survey. ACM Computing Surveys, 57(4), 1-35. https://doi.org/10.1145/3704633
[40] Pinto, R. P., Silva, B. M., & Inácio, P. R. (2025). Federated learning for anomaly detection on Internet of Medical Things: A survey. Internet of Things, 33, 101677. https://doi.org/10.1016/j.iot.2025.101677