

# Shard-Unlearn: A Sharded Elastic SGD Privacy-Preserving Federated Unlearning Framework for 5G-Assisted Healthcare

Sudip Chatterjee<sup>1</sup>, Samyak Jain<sup>1</sup>, Pronaya Bhattacharya<sup>1</sup>, Sandip Roy<sup>2</sup>, Soumya Banerjee<sup>2</sup>, Pratip Rana<sup>3</sup>, Sachin Shetty<sup>2</sup>

<sup>1</sup>Department of CSE, Amity School of Engineering and Technology, Amity University, Kolkata, India

<sup>2</sup>Center for Secure & Intelligent Critical Systems, Old Dominion University, Virginia, USA

<sup>3</sup>Department of Computer Science, Old Dominion University, Virginia, USA

{schatterjee1, pbhattacharya}@kol.amity.edu, samyakjain86602@gmail.com, {sroy, slbanerj, prana, sshetty}@odu.edu

## Abstract

Smart healthcare systems are generating unprecedented volumes of sensitive data, making robust privacy preservation a critical requirement. Traditional machine unlearning (MU) techniques aim to excise specific data points and their statistical influence from trained machine learning (ML) models. Thus, they suffer from limited computational efficiency, poor scalability, and suboptimal model convergence when applied to large-scale, big-data (BD) healthcare environments. These limitations become even more significant in 5G-assisted settings, where real-time connectivity and rapid data processing are essential. To address these challenges, we introduce the concept of data sharding which partitions healthcare datasets into manageable segments. In the paper, we introduce *Shard-Unlearn* framework, that implements federated unlearning (FU) process to the shards that contain sensitive data. This reduces the overall computational overhead and optimizes model convergence over 5G networks. In the framework, we present the elastic stochastic gradient descent (SGD) optimization which effectively removes the targeted data and associated statistical perturbations from the local models. The framework is tested over the *ADMISSIONS* benchmark dataset, which is divided into 10 shards. The framework is compared on computational efficiency, model robustness, and privacy preservation metrics. Statistical findings reveal a 47.14% improvement in unlearning impact (as measured by recall) while striking a balanced trade-off between performance and data security. These results underscore the viability of the framework as a scalable and privacy-preserving solution for modern 5G-assisted healthcare systems.

## Introduction

The advent of big data (BD) assisted healthcare has resulted in exponential data surge, fueled by the widespread adoption of healthcare sensors, and patient electronic health records (EHRs). This leads to unprecedented challenges in managing, processing, and securing sensitive information (Cao and Yang 2015), (Banerjee et al. 2024). As data volumes continue to rise, the necessity for real-time data transmission and analysis becomes critical for time-sensitive applications such as intensive care unit (ICU) monitoring and emergency medical response (Roy et al. 2016), (Mishra et al. 2023),

(Bhattacharya et al. 2020).

To address this, fifth generation (5G) connectivity guarantees seamless, uninterrupted data exchange between remote monitoring devices and decentralized systems, which is vital for instantaneous clinical decision-making (Thakur, Kharbas, and Manashree 2024). As healthcare grows more networked and data-driven, preserving privacy is vital, with identifiers and records vulnerable to inference attacks. Although privacy techniques like data anonymization (Vovk, Pihou, and Ross 2021), differential privacy (Javed et al. 2023), and  $t$ -closeness (Verma et al. 2022) have been proposed, but these solutions suffer from diminishing utility problem, addition of statistical noise to data (outside the dataset boundary) and struggling with high computational overhead and data sparsity issues respectively. To add to the complexity, the legal and ethical guidelines embodied in regulations like the general data protection regulation (GDPR), and the health insurance portability and accountability act (HIPAA) demand not only robust privacy measures but also the ability to selectively remove individual data points upon request (Chatterjee and Roy 2018).

To address the privacy and regulation guidelines, machine unlearning (MU) emerged as a promising solution. Unlike traditional data deletion that requires complete model retraining, MU seeks to selectively eliminate specific data points and their statistical influence from a trained model without the excessive computational burden of full-scale retraining (Li et al. 2024). Federated unlearning (FU) extends this concept into distributed environments, ensuring that privacy-sensitive operations are performed locally without exposing raw data.

Several studies have explored FU mechanisms to enhance data privacy without compromising model performance. For instance, (Ge 2024) applied FU in medical image analysis, demonstrating improved scalability and efficacy in removing sensitive information. Another approach in (Chen et al. 2024) iteratively adjusts the original model to entirely erase traces of the forgotten data. Authors in (Zhou et al. 2024) introduced a simplified FU strategy (SFU) capable of eliminating targeted data impacts while maintaining overall model integrity.

However, personalized medicine demands precise data control, which requires an integration of FU with differential privacy (DP). More specifically,  $\epsilon$ -DP quantitatively con-

trols the privacy-utility trade-off in FU setting. For example, smaller  $\epsilon$  values inject more noise for stronger privacy guarantees while still preserving overall model performance. However, despite the key benefits of FU and  $\epsilon$ -DP bound guarantees, the schemes face challenges related to scalability and efficiency when applied to large-scale, BD-assisted healthcare datasets. This motivates the use of smaller networks, or shards (autonomous controlled segments) in the BD setup to effectively manage the incoming flux of unlearning requests.

## Motivation and Contribution

Conventional FU methods operate over monolithic datasets impose excessive computational overhead and introduce delays. Motivated by these limitations, the paper presents the *Shard-Unlearn* framework that partitioning the model data into independent, autonomous segments, termed a shards. The framework localize the unlearning process to only those shards that contain the sensitive data. This reduces the computational overhead significantly, improves convergence rates, and minimizes the disruptive impact on the overall model.

Further, in combination of sharding, the framework employs an elastic stochastic gradient descent (SGD) optimization that dynamically adapts to the inherent fluctuations introduced by data removal and privacy-preserving noise. Alternative techniques such as gradient inversion or attention-based unlearning are proposed, but they often require iterative, resource-intensive adjustments. This flexibility is critical in environments where model parameters span millions of entries and the cost of full retraining is prohibitively high, often estimated in the tens of thousands of dollars for terabyte-scale datasets. The contributions of this article are as follows.

- We introduce a three layered system model that integrates data sharding with FU in BD-assisted healthcare without the need for full model retraining.
- Based on the system model, the framework implements the elastic SGD for patient data removal from trained data based on sharded node's locations.
- Extensive simulations on the real world *ADMISSIONS* dataset is conducted in terms of metrics like computational efficiency, model robustness, and retraining times.

The rest of the article is constructed as follows. Section discusses the system model and problem formulation. Section ?? presents the elastic SGD optimization based on multi-objective problem. Section presents performance analysis of the proposed model. Finally, Section offers the concluding remarks of the article.

### *Shard-Unlearn*: System Model and Problem Formulation

The section presents a three layered system model. Figure 1 presents the details. the three-layer model.

#### The layered design

The details of the three layers and the connected components are described below.

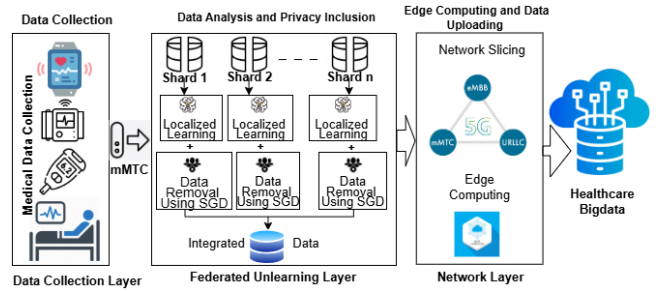


Figure 1: The layered architecture

**Data Collection Layer** : This layer deals with the set of medical Internet of Things (MIoT) and other sensor devices (e.g. accelerometer, gyroscope, pressure sensors, and others) associated with patients  $E_p$ . In  $E_p$ , consider  $n$  patients,  $E_p = \{E_1, E_2, \dots, E_n\}$ . The association of a patient and equipment is represented by the mapping as follows.

$$M = \{E_i, E_{mj}, E_{st}\} \quad (1)$$

where  $E_i$  is a unique patient identification number,  $E_{mj}$  is the MIoT  $j^{th}$  associated with the patient  $i^{th}$  and  $E_{st}$  represents  $t^{th}$  sensor collected data from the passenger. The data generated from  $j^{th}$  IoT/sensor is a time series data stream as follows.

$$D_i(t) = \{d_{j1}, d_{j2}, \dots, d_{jn}\} \quad (2)$$

where  $d_{jk}$  represents  $k$ -th data sample from  $j$ -th IoT or sensor.

Once the data collection is completed, we consider data is transferred to heterogeneous sensor nodes via the 5G massive machine type communication (mMTC) network. The collected data is transferred using constrained application protocol (CoAP), in JSON format for uniformity and lightweight transfer over nodes.

**5G Channel Setup** : For packets transmission, let  $P_t$  denotes the transmitting power,  $G_t$  and  $G_r$  denotes the transmitter and receiver antenna gains, respectively.  $\lambda$  the signal wavelength, and  $d$  the distance between transmitter and receiver. The received power  $P_r$  is given as follows.

$$P_r = \frac{P_t G_t G_r \lambda^2}{(4\pi d)^\alpha} \quad (3)$$

where  $\alpha$  is the path loss exponent. The signal-to-noise ratio (SNR) is then defined as follows.

$$\text{SNR} = \frac{P_r}{N_0 B} \quad (4)$$

with  $N_0$  representing the noise spectral density and  $B$  the channel bandwidth. The channel capacity is determined via the Shannon formula as follows.

$$C = B \log_2 (1 + \text{SNR}) \quad (5)$$

Over the channel, heterogeneous service sets are defined to cater to diverse QoS requirements, which is presented in subsection .

**Federated Unlearning Layer** In our framework, the process of data sharding serves as the foundation for localized FU design. Given the aggregated dataset  $D_{agg}$  collected from various sensors, we partition the data into  $r$  independent shards. Let  $\mathcal{M}$  denote the trained ML model, parameterized by  $\theta$ , obtained by minimizing the loss function over the dataset  $D$  as follows.

$$\theta^* = \arg \min_{\theta} \sum_{x_i \in D} \mathcal{L}(x_i, \theta). \quad (6)$$

When a data removal request is issued—for instance, to delete sensitive attributes like a patient identifier, the affected shard(s) apply unlearning locally to recalibrate their models.

Specifically, to remove the influence of sensitive data in shard  $j$ , we identify the subset of data to be forgotten, denoted by  $D_{rem,j}$ . The gradient contribution of this forgotten data is computed as follows.

$$\Delta\theta_j = \eta \sum_{x_i \in D_{rem,j}} \nabla \mathcal{L}(x_i, \theta_j), \quad (7)$$

where  $\theta_j$  represents the current local model parameters in shard  $j$  and  $\eta$  is the learning rate. The updated model parameter, after unlearning, is then given as follows

$$\theta'_j = \theta_j - \Delta\theta_j. \quad (8)$$

To further ensure robust privacy, our system integrates  $\epsilon$ -DP. For a given privacy budget  $\epsilon$ , noise proportional to the sensitivity  $\Delta L$  of the loss function is injected into the updates to satisfy the  $(\epsilon, \delta)$ -DP guarantee. Specifically, noise  $\nu$  is sampled from a Gaussian distribution as follows.

$$\nu \sim \mathcal{N}(0, \sigma^2), \quad \text{with } \sigma \propto \frac{\Delta L}{\epsilon}. \quad (9)$$

This mechanism bounds the privacy loss while retaining model utility.

Further, the framework employs an elastic SGD algorithm for localized learning and unlearning. Unlike traditional SGD with a fixed learning rate, elastic SGD dynamically adjusts  $\eta$  in response to gradient variations and injected DP noise, ensuring faster convergence and a well-balanced privacy-utility trade-off.

**Network Layer (5G service sets)** At this layer, via the network slicing technology, the 5G spectrum into virtual sub-networks for various healthcare services, resulting in optimal bandwidth utilization and minimized interference. Each slice  $L_k$  is dynamically assigned based on the service type:

$$L = \{L_{uRLLC}, L_{eMBB}, L_{mMTC}\} \quad (10)$$

where  $L_{uRLLC}$  ensures ultra-low latency for critical healthcare tasks,  $L_{eMBB}$  deals with high throughput applications and  $L_{mMTC}$  supports massive IoT connectivity. To reduce latency, edge computing is integrated at the network edge, where data is processed locally before being sent to the cloud. The edge computing latency is determined as follows.

$$T_{edge} = \frac{D_{proc}}{C_{edge}} + \frac{D_{trans}}{R_{edge}} \quad (11)$$

where  $D_{proc}$  and  $D_{trans}$  are the processing and transmission data sizes,  $R_{edge}$  is the local transmission rate,  $C_{edge}$  is the edge computation capacity.

## Problem Formulation

In this subsection, a joint optimization is presented to minimize three key metrics: the cumulative unlearning loss across all data shards, the aggregate latency from edge processing and network transmission, and the overall privacy leakage as measured under an  $\epsilon$ -DP guarantee.

$$W = \{W_j\}_{j=1}^r \quad (12)$$

denote the set of local model parameters across the  $r$  shards. We define the corresponding objectives as follows:

$$\begin{aligned} f_1(W) &= \sum_{j=1}^r \mathcal{L}_j(W_j) \\ f_2(W) &= \sum_{j=1}^r \left( \frac{D_{proc,j}}{C_{edge}} + \frac{D_{trans,j}}{R_{edge}} \right) \\ f_3(W, \epsilon) &= \sum_{j=1}^r \phi(W_j, \epsilon) \end{aligned} \quad (13)$$

Thus, the multi-objective optimization problem is formulated as:

$$\min_{W, \epsilon, \eta} \{f_1(W), f_2(W), f_3(W, \epsilon)\} \quad (14)$$

subject to the following constraints:

$$\begin{aligned} C_1 : \text{SNR} &= \frac{P_t G_t G_r \lambda^2}{(4\pi d)^\alpha N_0 B} \geq \gamma \\ C_2 : B \log_2(1 + \text{SNR}) &\geq R_{\min} \\ C_3 : \epsilon &\leq \epsilon_{\max}, \quad W_j \in \mathcal{W}, \quad \forall j = 1, \dots, r \\ C_4 : \sum_{j=1}^r \left( \frac{D_{proc,j}}{C_{edge}} + \frac{D_{trans,j}}{R_{edge}} \right) &\leq T_{\max} \end{aligned} \quad (15)$$

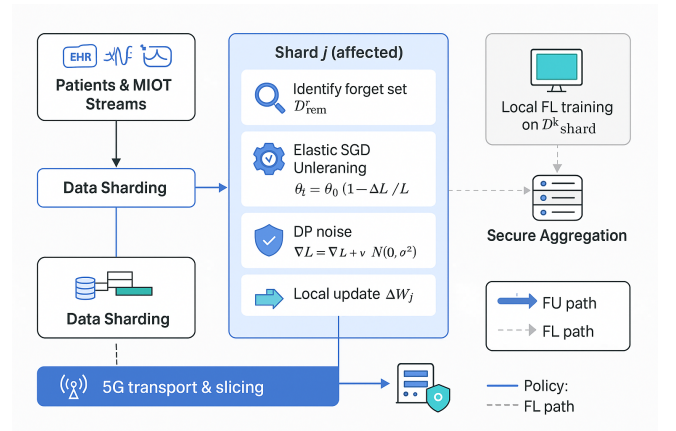


Figure 2: The FU process

## Performance Evaluation

This section comprises all evaluation and analysis of our proposed ShardedFU algorithm. The evaluation focuses on the preferability and effectiveness of the FU

mechanism in healthcare applications. To carry out the evaluation, we utilize the *ADMISSIONS* dataset from *mimic\_iii\_clinical\_dataset\_demo* (kag 2020). This dataset contains 19 columns related to patients. Patient identity is represented by the column *'hadm\_id'*. Choosing *'hospital\_expire\_flag'* column as the target variable of the *ADMISSIONS.csv* dataset, which reflects patient mortality (1 if deceased, 0 if survived), enables privacy preservation in the *Shard-Unlearn* framework for 5G-powered healthcare. As a binary target, it reduces the leakage of individual medical records or personal information, maintaining confidentiality in federated learning in which data stays local to hospitals.

Targeting unlearning on this flag (e.g., deleting deceased patient records) reduces the likelihood of leakage of sensitive attributes such as diagnosis or admission time. Combined with *ShardedFU's* differential privacy, adding noise to avoid re-identification, this choice becomes HIPAA compliant and maintains the framework's privacy-preserving objectives for secure health data processing. The comparison is made based on three primary aspects: MU, FL variants, and other unlearning models. The detailed results and analyses are provided below.

### Preferability of Federated Unlearning

The proposed *ShardedFU* model integrates DP, injecting noise into data to obscure individual identities while preserving privacy during FL and FU. *ShardedFU* model is compared with the traditional MU mechanism implemented on the same dataset.

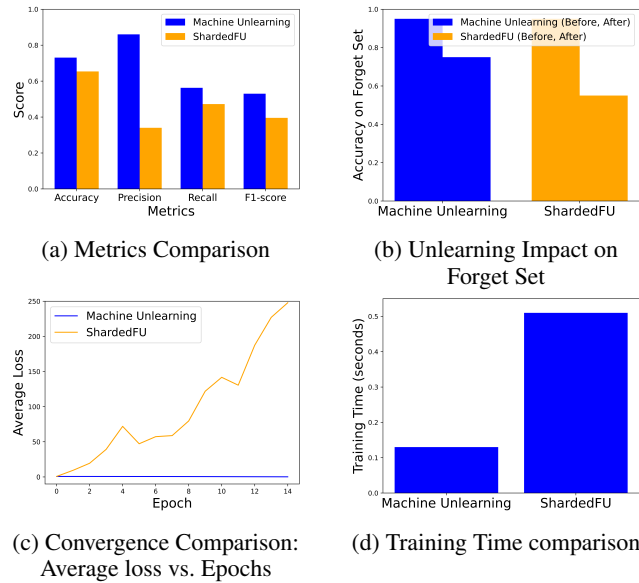


Figure 3: ShardedFU vs. Traditional Machine Unlearning

Figure 3a establishes that the *ShardedFU* model surpasses MU in test metrics. The lower accuracy of the *ShardedFU* denotes the reduced reliance on forget set, which further reduces the risk of data leakage or privacy attacks. The federated approach used in *ShardedFU* ensures privacy as there is no sharing of raw data which is not the case for traditional

MU. The drop in the *ShardedFU's* accuracy after unlearning as shown in Figure 3b which indicates stronger unlearning effect because of the federated approach and integrated DP.

The convergence curve is shown in Figure 3c, which shows the average loss for *ShardedFU* and MU over epochs. MU shows stable loss whereas *ShardedFU* exhibits dynamic loss, reflecting the DP noise and enhanced privacy. The gap in the training time of the models, is depicted in Figure 3d. The benefit is due to MU which involves straightforward retraining while the proposed approach involves federated training, DP noise application, and unlearning steps for multiple clients.

### Federated Unlearning Efficacy

This subsection examines FU's efficacy by comparing it against two FL variants: FL without unlearning -FL (no deletion), and FL with retraining FL (retraining) methodologies. The retraining variant FL (retraining) shows improvement in terms of data security by mitigating the influence of the forget set. The balanced metrics of *ShardedFU*, is shown in Figure 4a, which supports the proposition of using the unlearning mechanism with federated approach in healthcare systems. Referring to Figure 4b, the volatile nature of *ShardedFU* due to DP noise addition and unlearning, showcases effective data removal which is critical for security in healthcare. The other two approaches exhibit stable and low loss but lack unlearning, making them unsuitable for data removal.

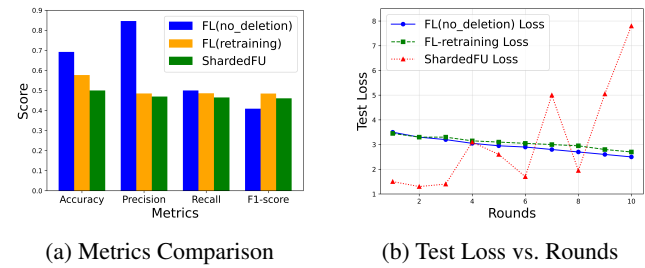


Figure 4: ShardedFU vs. Federated Learning Variants




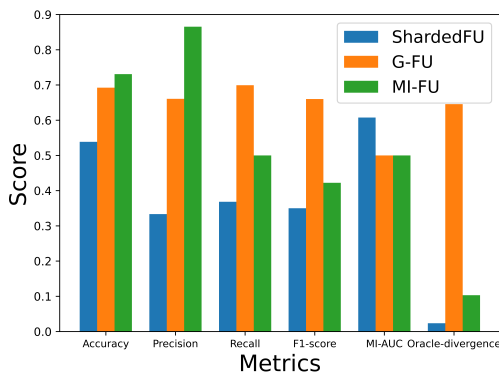
Figure 5: Confusion Matrix for ShardedFU and Federated Learning variants

Referring to Figure 5, the confusion matrix for the two approaches is provided to visualize the use of DP to effectively preserve data. The privacy budget  $\epsilon$  is the measure of noise

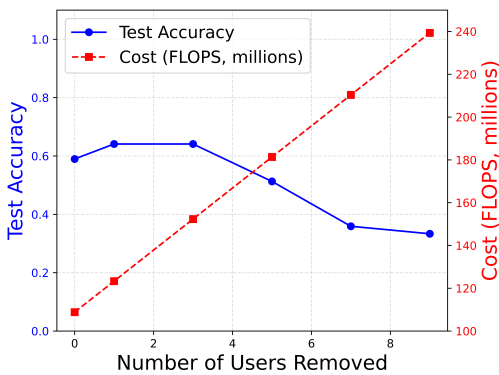
added to the data set. FL (retraining) approach does not show much misclassifications since there is no unlearning effect or integrated privacy preserving technique, but for ShardedFU approach the misclassifications can be seen, demonstrating the unlearning effect and also the influence of DP noise.

### Comparative Analysis: ShardedFU vs. Other Unlearning Models

ShardedFU compromises accuracy to achieve higher privacy preservation, but still manages to attain comparable performance and outstands among the three, as discussed in this subsection. To carry out this comparison, the alternative unlearning models used are Graph based FU (G-FU) and Model-Inversion based FU (MI-FU) model. G-FU is a FL system that uses graph-structured perturbations and Gaussian noise in order to improve privacy and facilitate unlearning through the reduction of the effect of individual data (He et al. 2021). MI-FU uses gradient inversion in the name of targeted unlearning, locally perturbing the parameters and accumulating the updates with momentum (Li et al. 2022). ShardedFU partitions clients into shards, and adds uniform noise (0.5), and uses FU to successfully remove individual data (Romandini et al. 2024).



(a) Metrics Comparison of ShardedFU with other unlearning models



(b) Trade-off Calculation for ShardedFU

Figure 6: Comparative Analysis of ShardedFU

Table 1 provides the data on the performance of different

unlearning models, also depicted in Figure 6a. The ShardedFU metrics, which are comparatively lower than the other two models, demonstrate strong privacy preservation and also good unlearning impact. The lowest recall of ShardedFU suggests that true positive cases are unlikely to be accurately identified, benefiting the goal of individual data security. To support the claim of individual data security, we have also evaluated the models on privacy leakage/MI-AUC (Membership Inference AUC) and fidelity to the divergence baseline (Oracle-divergence). MI-AUC is used to determine whether a particular data point was included in the training data set or not (Shokri et al. 2017). Oracle divergence, on the other hand, outlines parameter difference between the developed unlearning model and 'ideal' oracle model (Ginart et al. 2019). According to table Table1, with a very lower divergence value and an acceptable MI-AUC score, ShardedFU demonstrates strong fidelity to the retrain baseline, suggesting the best suitability of this model among all.

Figure 6b shows the trade-off of ShardedFU in the unlearning process by removing increasing number of clients. The cost rises linearly for increasing number of clients removed. For testing on 10 clients, the trade-off is found to be justified for 5 (approx.) clients. Beyond the intersecting point, the computational cost for unlearning client data would become disproportional to the loss in accuracy due to unlearning steps and DP noise, reflecting the privacy-utility trade-off.

Data for evaluating the performance of the unlearning models in epsilon and epochs are provided in Table 2 along with Figure 7a and Figure ???. ShardedFU shows low accuracy for initial values of epsilon due to strong DP noise and then stabilizing at higher values of epsilon and epochs, showing privacy-utility balance. In contrast, G-FU model shows initial peak due to graph-based perturbation, then declines because of unlearning step and then fluctuates as the adjustments of graph parameter changes. It also varies widely for different values of epsilon showing instability and sensitivity. MI-FU model shows increase in the accuracy for initial epochs (5 to 10) and then decreases for higher values of epochs (25 to 100) due to weak unlearning.

There are some differences in terms of use and advantage of using this *Shard-Unlearn* over other mechanisms like SISA (Sharding, Isolation, Separation and Aggregation). *Shard-Unlearn* offers clear advantages over SISA. While SISA requires centralized retraining of affected slices—computationally expensive and unsuitable for 5G—*Shard-Unlearn* employs federated, shard-level unlearning with elastic updates. By excluding flagged clients and aggregating only active ones, it reduces computation and communication, accelerates reconvergence, and achieves higher efficiency under strict 5G bandwidth and latency constraints.

### Network Performance

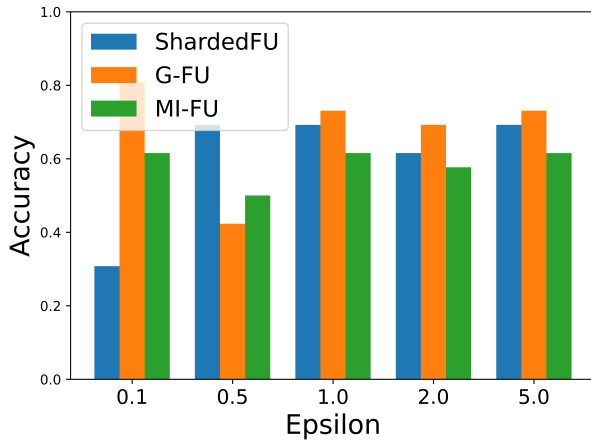
The experiment with the model showed an efficient network performance in the third layer, where 5G service have been used to send data to cloud. The Figure 8 confirms that 5G network slicing (uRLLC & eMBB) is the most efficient solution for and ICU environments. The ability to achieve

Model	Accuracy	Precision	Recall	F1-score	MI-AUC	Oracle-divergence
<b>ShardedFU</b>	<b>0.54</b>	<b>0.33</b>	<b>0.37</b>	<b>0.35</b>	<b>0.61</b>	<b>0.02</b>
G-FU (He et al. 2021)	0.69	0.66	0.70	0.66	0.50	0.65
MI-FU (Li et al. 2022)	0.73	0.87	0.50	0.42	0.50	0.10

Table 1: Performance comparison of different unlearning models

Model Type	Epsilon ( $\epsilon$ )					Epochs				
	0.1	0.5	1.0	2.0	5.0	5	10	25	75	100
ShardedFU	30.77%	69.23%	69.23%	61.54%	69.23%	34.62%	30.77%	69.23%	69.23%	69.23%
G-FU (He et al. 2021)	80.77%	42.31%	73.08%	69.23%	73.08%	65.38%	57.69%	69.23%	53.85%	73.08%
MI-FU (Li et al. 2022)	61.54%	50.00%	61.54%	57.69%	61.54%	69.23%	73.08%	69.23%	50.00%	42.31%

Table 2: Model Performance Across Epsilon and Epochs



(a) Accuracy vs. Epsilon

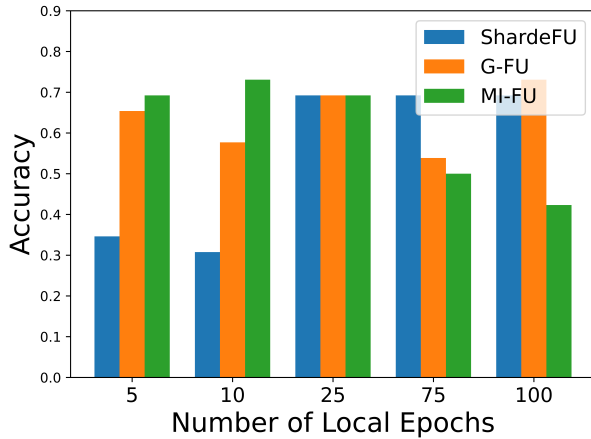


Figure 7: Model Accuracy across Epsilon and Epochs

throughput above 1 Gbps ensures seamless remote monitoring, real-time AI processing, and telemedicine applications. Performance metrics and network configuration for the 5G ICU network based on the *ADMISSIONS* dataset are presented in Table 3.

CDF of Network Throughput in Smart Healthcare Scenarios




Figure 8: 5G Precision Analysis in ICU Network (CDF Plot)

Network Slice	Data Type	Latency (ms)	Data Rate (Mbps)	Reliability (%)
Critical Patient Monitoring (uRLLC)	Vital Signs (Heart Rate, BP, SpO <sub>2</sub> )	≤ 1	>512	99.999
Clinical Decision Support (AI Processing)	Severity Scores, Lab Results	≤ 10	55	99.99
EHR Data Access (eMBB)	Clinical and patient Data	≤ 50	6.2	98.9

Table 3: Performance Metrics and Network Configuration for 5G ICU Network

## Concluding Remarks

In this paper, we proposed a shardedFU framework, *Shard-Unlearn*, designed for 5G-assisted smart healthcare data. We implemented ShardedFU using elastic SGD over a sharded *ADMISSIONS* dataset. ShardedFU maintains an optimal balance among model accuracy, robustness, and data privacy. Our evaluation shows that it surpasses existing methods like G-FU and MI-FU in both unlearning efficiency and preservation of individual privacy. Additionally, performance across  $\epsilon$  and epochs remains stable and balanced.

Future extension of the work would integrate multi-modal healthcare data and federated reinforcement learning tech-

niques to improve model robustness and scalability.

## References

2020. MIMIC-III Clinical Dataset Demo. <https://www.kaggle.com/datasets/atamazian/mimic-iii-clinical-dataset-demo>. [Accessed 2025-03-06].
- Banerjee, S.; Roy, S.; Ahamed, S. F.; Quinn, D.; Vucovich, M.; Nandakumar, D.; Choi, K.; Rahman, A.; Bowen, E.; and Shetty, S. 2024. MIA-BAD: An Approach for Enhancing Membership Inference Attack and its Mitigation with Federated Learning. In *2024 International Conference on Computing, Networking and Communications (ICNC)*, 635–640. IEEE Computer Society.
- Bhattacharya, M.; Roy, S.; Mistry, K.; Shum, H. P.; and Chattopadhyay, S. 2020. A privacy-preserving efficient location-sharing scheme for mobile online social network applications. *IEEE Access*, 8: 221330–221351.
- Cao, Y.; and Yang, J. 2015. Towards Making Systems Forget with Machine Unlearning. In *2015 IEEE Symposium on Security and Privacy*, 463–480.
- Chatterjee, S.; and Roy, S. 2018. An efficient dynamic access control scheme for distributed wireless sensor networks. *International Journal of Ad Hoc and Ubiquitous Computing*, 27(1): 1–18.
- Chen, K.; Wang, Y.; Zhao, L.; Jiang, C.; Mai, H.; Wu, Y.; Hong, H.; Shen, Y.; Mo, J.; Huang, L.-L.; Peng, J.; Wang, X.; and Yang, Q. 2024. Private Data Protection With Machine Unlearning for Next-Generation Networks. *IEEE Open Journal of the Communications Society*, 1–1.
- Ge, L. 2024. Erasing memories: implementing client unlearning in medical image analysis. In Qin, C.; and Zhou, H., eds., *International Conference on Image Processing and Artificial Intelligence (ICIPAI 2024)*, volume 13213, 132133V. International Society for Optics and Photonics, SPIE.
- Ginart, A.; Guan, M.; Valiant, G.; and Zou, J. Y. 2019. Making AI Forget You: Data Deletion in Machine Learning. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- He, C.; Balasubramanian, K.; Ceyani, E.; Yang, C.; Xie, H.; Sun, L.; He, L.; Yang, L.; Yu, P. S.; Rong, Y.; et al. 2021. Fedgraphnn: A federated learning system and benchmark for graph neural networks. *arXiv preprint arXiv:2104.07145*.
- Javed, L.; Anjum, A.; Yakubu, B. M.; Iqbal, M.; Moqurrab, S. A.; and Srivastava, G. 2023. ShareChain: Blockchain-enabled model for sharing patient data using federated learning and differential privacy. *Expert Systems*, 40(5): e13131.
- Li, C.; Jiang, H.; Chen, J.; Zhao, Y.; Fu, S.; Jing, F.; and Guo, Y. 2024. An overview of machine unlearning. *High-Confidence Computing*, 100254.
- Li, Z.; Wang, L.; Chen, G.; Zhang, Z.; Shafiq, M.; and Gu, Z. 2022. E2EGI: End-to-end gradient inversion in federated learning. *IEEE Journal of Biomedical and Health Informatics*, 27(2): 756–767.
- Mishra, A. K.; Wazid, M.; Singh, D. P.; Das, A. K.; Roy, S.; and Shetty, S. 2023. ACKS-IA: An access control and key agreement scheme for securing industry 4.0 applications. *IEEE Transactions on Network Science and Engineering*, 11(1): 254–269.
- Romandini, N.; Mora, A.; Mazzocca, C.; Montanari, R.; and Bellavista, P. 2024. Federated unlearning: A survey on methods, design guidelines, and evaluation metrics. *IEEE Transactions on Neural Networks and Learning Systems*.
- Roy, S.; Chatterjee, S.; Chattopadhyay, S.; and Gupta, A. K. 2016. A biometrics-based robust and secure user authentication protocol for e-healthcare service. In *2016 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, 638–644. IEEE.
- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership Inference Attacks Against Machine Learning Models. In *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18.
- Thakur, P. K.; Kharbas, V. K.; and Manashree. 2024. Exploring the Role of Data Analysis in 5G-Enabled Health Care Management Models. In *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 1–6.
- Verma, A.; Bhattacharya, P.; Patel, Y.; Shah, K.; Tanwar, S.; and Khan, B. 2022. Data Localization and Privacy-Preserving Healthcare for Big Data Applications: Architecture and Future Directions. In Singh, P. K.; Kolekar, M. H.; Tanwar, S.; Wierzchoń, S. T.; and Bhatnagar, R. K., eds., *Emerging Technologies for Computing, Communication and Smart Cities*, 233–244. Singapore: Springer Nature Singapore. ISBN 978-981-19-0284-0.
- Vovk, O.; Pihó, G.; and Ross, P. 2021. Anonymization Methods of Structured Health Care Data: A Literature Review. In Attiogbé, C.; and Ben Yahia, S., eds., *Model and Data Engineering*, 175–189. Cham: Springer International Publishing. ISBN 978-3-030-78428-7.
- Zhou, L.; Zhu, Y.; Xue, Q.; Zhang, J.; and Zhang, P. 2024. Streamlined Federated Unlearning: Unite as One to Be Highly Efficient. *arXiv preprint arXiv:2412.00126*.