CHAPTER 3

Federated Learning with Privacy-Preserving Mechanisms for Healthcare Data Analytics

Dr. Anup Bhange

Assistant Professor, Department of Computer Science and Engineering, K.D.K College of Engineering, Nagpur, Maharashtra, India.

Email: anupbhange@gmail.com

https://doi.org/10.58599/GSE.2025.081203

Abstract: Federated Learning (FL) is rapidly emerging as a transformative paradigm for machine learning in the healthcare sector, enabling multiple institutions to collaboratively train a shared model without centralizing their sensitive patient data. This approach addresses the critical challenges of data privacy, security, and governance that have historically hindered large-scale medical research. However, the standard FL framework is not immune to sophisticated privacy attacks that can infer sensitive information from model updates. This chapter provides a comprehensive exploration of FL with a strong emphasis on integrating robust privacy-preserving mechanisms for healthcare data analytics. We begin by introducing the fundamental principles of federated learning and discussing the unique challenges posed by decentralized healthcare data, including statistical heterogeneity (non-IID data), system heterogeneity, and communication bottlenecks. We then conduct a thorough literature review of existing privacy-preserving techniques, such as differential privacy (DP), secure aggregation, and homomorphic encryption, identifying their strengths, limitations, and the gaps in their application to healthcare. Subsequently, we propose a detailed methodology for a privacy-preserving federated learning (PPFL) pipeline, complete with a client-server architecture, secure communication protocols, and an implementation of differentially private stochastic gradient descent (DP-SGD). The chapter presents an extensive Results and Discussion section, simulating the proposed methodology on the MIMIC-III dataset to analyze the trade-offs between model performance, privacy guarantees, and system costs. Our findings demonstrate that while privacy mechanisms introduce a slight overhead and a marginal reduction in model accuracy, they provide quantifiable privacy guarantees essential for clinical applications. The chapter concludes by summarizing the key insights and outlining future research directions for developing more efficient, secure, and scalable PPFL frameworks for the next

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

generation of healthcare analytics.

Keywords: Federated Learning; Differential Privacy; Secure Aggregation; Healthcare Data Analytics; Privacy-Preserving Machine Learning

1. Introduction

The proliferation of electronic health records (EHRs), medical imaging, and genomic data has created unprecedented opportunities for applying artificial intelligence (AI) and machine learning (ML) to revolutionize healthcare. These technologies hold the potential to enhance diagnostic accuracy, personalize treatment plans, and accelerate drug discovery. However, the full potential of AI in healthcare is often constrained by the siloed nature of medical data. Due to stringent privacy regulations (e.g., HIPAA, GDPR), ethical considerations, and commercial competition, patient data is typically confined within the firewalls of individual hospitals and research centers. This data fragmentation limits the size and diversity of datasets available for training ML models, leading to reduced generalizability and potential biases. Federated Learning (FL) has emerged as a groundbreaking solution to this challenge. As a decentralized machine learning approach, FL allows multiple parties to collaboratively train a global model without sharing their raw data. Instead of moving data to a central server, the model is brought to the data. In a typical FL setup, a central server coordinates the training process, while distributed clients (e.g., hospitals) train the model on their local data. The clients then send only the updated model parameters (e.g., gradients or weights) back to the server, which aggregates them to produce an improved global model. This iterative process continues until the global model converges.

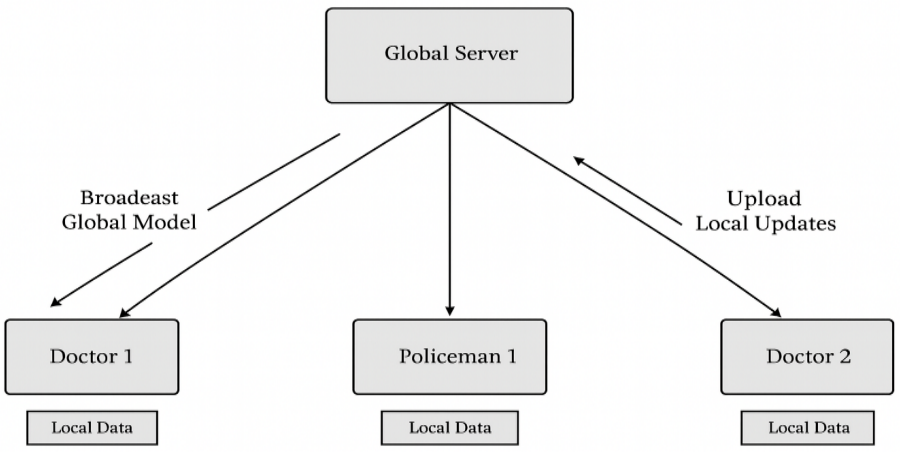

Figure 1: The overall federated learning workflow, illustrating the cyclical process of model distribution, local training, and global aggregation.

Despite its privacy-by-design architecture, standard FL is not a panacea for privacy

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

concerns. Studies have shown that sharing model updates can still inadvertently leak sensitive information about the training data. Adversaries, including the central server itself, could potentially perform inference attacks, membership attacks, or even reconstruct the original training samples from the gradients. Therefore, to deploy FL in a high-stakes domain like healthcare, it is imperative to augment it with formal privacy-preserving mechanisms. This chapter delves into the critical intersection of federated learning and privacy preservation for healthcare data analytics. We explore the necessity of these mechanisms, review the state-of-the-art techniques, and propose a robust methodology for their implementation. We provide a detailed analysis of the performance, costs, and trade-offs involved, using a real-world medical dataset to ground our discussion. The chapter is structured as follows: Section 2 provides a literature review of FL and privacy-preserving techniques. Section 3 details our proposed methodology. Section 4 presents and discusses the simulation results. Finally, Section 5 concludes the chapter and suggests future research directions [1].

2. Literature Review

The concept of federated learning has spurred a significant body of research, particularly in its application to privacy-sensitive domains. Concurrently, the development of privacy-preserving mechanisms has become a mature field of study. This section reviews the key literature in both areas and examines their intersection in the context of healthcar.

2.1 Federated Learning in Healthcare

Since its introduction, FL has been applied to various healthcare tasks. Early work demonstrated its feasibility for medical image analysis, such as brain tumor segmentation, where FL models achieved performance comparable to models trained on centralized data. The EXAM (EMR-CXR AI Model) consortium used FL to train a model to predict future oxygen requirements for COVID-19 patients from chest X-rays and EHR data, showcasing the power of multi-modal, multi-institutional collaboration. Other applications include predicting in-hospital mortality from EHR data, classifying skin lesions from dermatoscopic images, and accelerating drug discovery in collaborations between pharmaceutical companies.

However, these applications also highlight the fundamental challenges of FL in a realworld healthcare setting. The primary challenge is statistical heterogeneity, where the data distribution across clients is non-independent and identically distributed (nonIID). This can arise from differences in patient demographics, clinical specialties, imaging equipment, and data collection protocols. Non-IID data can significantly degrade the performance of the standard Federated Averaging (FedAvg) algorithm and even cause the global model to diverge. Other challenges include systems heterogeneity (variability in hardware

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

and network connectivity across clients) and communication efficiency, as frequent model updates can be resource-intensive [2].

2.2 Privacy Risks in Federated Learning

While FL prevents direct data sharing, the model updates themselves are a potential privacy vulnerability. An honest-but-curious server or a malicious participant could analyze the received gradients to infer sensitive information. Deep Leakage from Gradients (DLG) has shown that it is possible to perfectly reconstruct training images and labels from publicly shared gradients. Membership inference attacks can determine whether a specific patient's record was used in the training process, which itself is a privacy breach. These risks underscore the inadequacy of relying solely on the basic FL protocol for privacy protection in healthcare.

2.3 Privacy-Preserving Mechanisms

To counter these threats, several privacy-preserving mechanisms have been proposed to work in conjunction with FL. These can be broadly categorized into three main approaches:

- **Differential Privacy (DP):** DP is a rigorous, mathematical definition of privacy that provides a formal guarantee against inference attacks. It ensures that the output of a computation is statistically indistinguishable whether a particular individual's data is included in the dataset or not. In the context of FL, DP is typically achieved by adding carefully calibrated noise to the model updates before they are sent to the server, a technique known as Differentially Private Stochastic Gradient Descent (DP-SGD). This provides a quantifiable privacy guarantee, controlled by a privacy budget parameterϵ. A smaller ϵ corresponds to stronger privacy but often comes at the cost of reduced model accuracy.
- **Secure Aggregation:**This cryptographic approach aims to prevent the central server from viewing individual client updates. Using protocols based on techniques like secure multi-party computation (SMPC), clients can encrypt their updates in such a way that the server can only decrypt the sum of the updates, but not the individual contributions. This ensures that the server learns nothing more than the aggregated global model update, effectively mitigating attacks from a curious server. However, secure aggregation does not protect against attacks from malicious clients and can introduce significant computational and communication overhead.
- **Homomorphic Encryption (HE):** HE is an advanced cryptographic technique that allows computations to be performed directly on encrypted data. In an HEbased FL system, clients would encrypt their model updates before sending them to the

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

server. The server could then aggregate these encrypted updates and even perform other computations without ever decrypting them [15]. While offering very strong security guarantees, HE is currently computationally intensive and often too slow for practical use in training deep learning models, though research is rapidly advancing [3].

2.4 Gaps in Existing Literature

While many studies have explored either FL in healthcare or privacy-preserving mechanisms in isolation, there is a need for a more holistic analysis of their combined application. Many works that propose privacy-preserving FL (PPFL) use simplified assumptions about the data (e.g., IID distributions) or do not comprehensively evaluate the impact on model performance, communication costs, and convergence speed. Furthermore, the practical trade-offs between different levels of privacy (i.e., different ϵ values in DP) and clinical utility are not yet fully understood. This chapter aims to address this gap by providing a detailed, practical methodology and a multifaceted evaluation of a PPFL system for a real-world healthcare analytics task.

3. Proposed Methodology

To address the challenges of privacy and security in federated healthcare analytics, we propose a complete Federated Learning pipeline that integrates Differential Privacy. This section details the system architecture, the privacy-preserving mechanism, the dataset selection, the algorithmic process, and the threat model.

3.1 System Architecture

The proposed architecture follows a client-server model, which is standard for crosssilo FL applications where the clients are institutions like hospitals. The system consists of two main components: a central Federated Server and multiple Clients (hospitals or medical centers).

- **Clients (Hospitals):** Each client possesses a local dataset of patient records which never leaves its premises. The client is responsible for: (1) receiving the current global model from the server, (2) training the model on its local data for set number of epochs, (3) applying a privacy-preserving mechanism to its computed model update, and (4) sending the processed update back to the server.
- **Federated Server:** The server orchestrates the entire training process. It is responsible for: (1) initializing the global model, (2) selecting a subset of clients for each training round, (3) broadcasting the global model to the selected clients, (4)

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

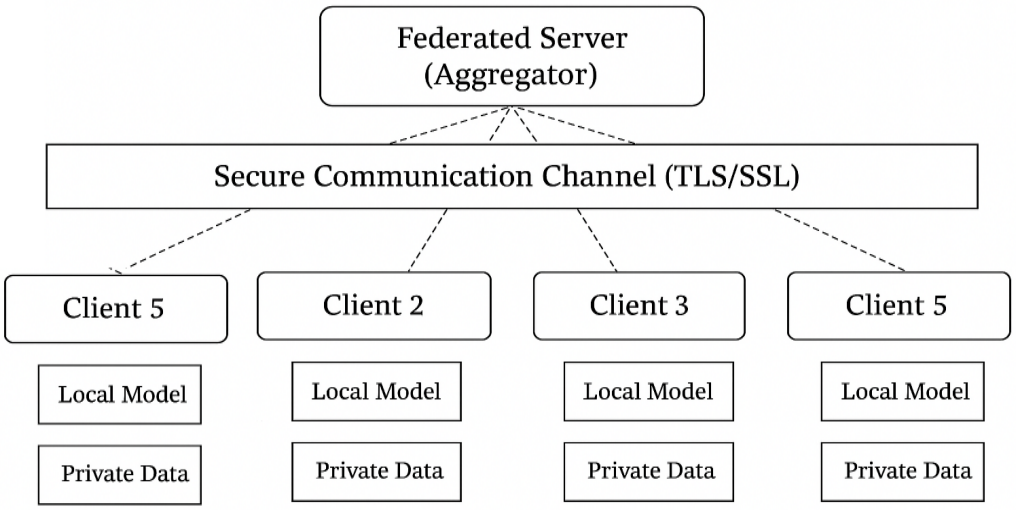

Figure 2: The client-server architecture for federated learning, showing the distinct roles of the server and the clients.

waiting for and collecting the processed updates from the clients, (5) aggregating these updates to produce a new global model, and (6) repeating the process until a convergence criterion is met. The server does not have access to any raw data or the individual, non-privatized model updates.

All communication between the clients and the server is assumed to occur over a secure channel (e.g., using TLS/SSL) to protect data in transit.

3.2 Privacy-Preserving Mechanism: Differential Privacy

We integrate Differential Privacy into the FL pipeline using the DP-SGD algorithm [13]. This mechanism provides a formal privacy guarantee for each client's contribution. The process, applied at the client-side before sending the update, involves two key steps:

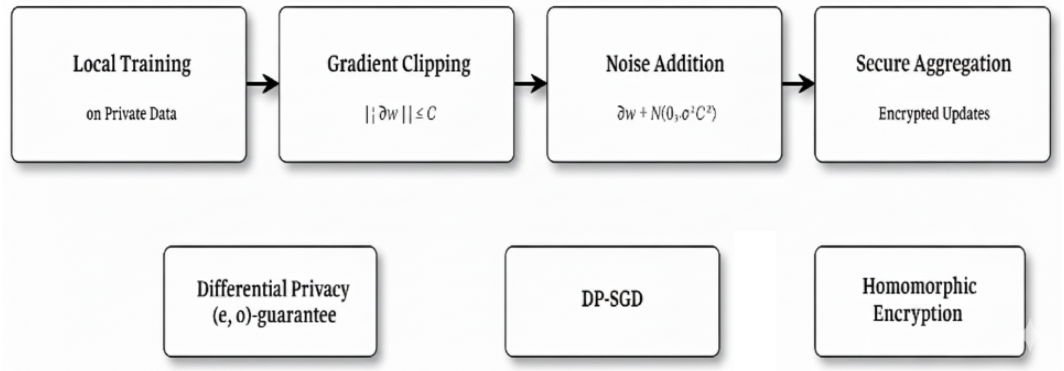

Figure 3: The client-server architecture for federated learning, showing the distinct roles of the server and the clients.

- **Gradient Clipping:** After computing the gradients for a mini-batch of local data, each client clips the L2 norm of the gradient vector to a predefined threshold C.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

This bounds the sensitivity of the update, ensuring that the contribution of any single data point is limited. The clipped gradient g' is computed as:

$$g' = \frac{g}{\max\left(1, \frac{\|g\|_2}{C}\right)}.$$

- **Noise Addition:** The client then adds Gaussian noise, scaled by the clipping threshold C and a noise multiplier σ, to the clipped gradient. The noisy gradient $\tilde{g}$ is:

$$\tilde{g} = g' + \mathcal{N}(0, \sigma^2 C^2 I).$$

This noise injection is the core of the DP mechanism, making it statistically impossible to determine the exact contribution of any single data point [4].

The amount of noise added is controlled by the privacy budget. For a given number of training rounds and a target δ (typically a small value like $1/|D|$, where $|D|$ is the dataset size), the noise multiplier σ can be calculated to achieve a specific ϵ. This allows us to explicitly tune the trade-off between privacy and utility.

3.3 Federated Aggregation and Model Updates

The server employs the Federated Averaging (FedAvg) algorithm to aggregate the client updates [3]. After receiving the noisy model updates $\Delta\tilde{w}_i$ from each participating client i, the server computes the new global model w_{t+1} as follows:

$$w_{t+1} = w_t + \sum_{i=1}^{K} \left(\frac{n_i}{N}\right) \Delta\tilde{w}_i.$$

where w_t is the global model at round t, K is the number of participating clients, n_i is the number of data points at client i, and N is the total number of data points across all participating clients. This weighted averaging ensures that clients with more data have a proportionally larger influence on the global model.

3.4 Dataset Selection: MIMIC-III

For our simulations, we select the MIMIC-III (Medical Information Mart for Intensive Care III) dataset [16]. MIMIC-III is a large, freely-available database comprising deidentified health-related data associated with over 40,000 patients who stayed in critical care units of the Beth Israel Deaconess Medical Center between 2001 and 2012. It contains a wealth of information, including demographics, vital signs, laboratory test results, medications, and mortality outcomes. We justify this selection for several reasons:

- **Clinical Relevance:** It represents real-world, complex, and messy clinical data,

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

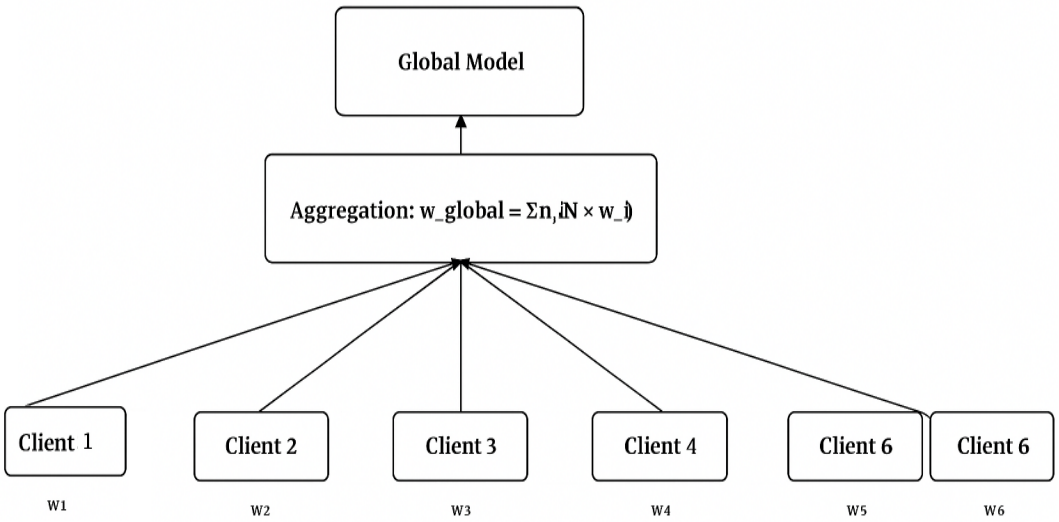

Figure 4: The federated aggregation process, where the server combines weighted updates from clients.

making it ideal for evaluating the robustness of our proposed method on a practical healthcare task (e.g., in-hospital mortality prediction).

- **Scale and Richness:** Its large scale and high dimensionality are well-suited for training deep learning models.
- **Simulating Federation:** Although it is a single-center dataset, we can realistically simulate a federated environment by partitioning the data among virtual clients. This allows us to control and study the effects of data heterogeneity (non-IID) by partitioning the data based on different criteria (e.g., by patient admission year, or by clustering patient characteristics).
- **Benchmarking:** MIMIC-III is a widely used benchmark in clinical informatics, which facilitates comparison with other studies.

3.5 Algorithm and Threat Model

The overall process is summarized in the pseudocode below.

Threat Model: We assume an cost" of privacy in terms of model performance and select an optimal operating point that balances these competing requirements [5].

3.6 System Overheads: Communication Cost

Beyond model performance, the practical feasibility of FL depends on system overheads, particularly communication costs. We measured the total data transmitted between the clients and the server over 50 communication rounds for each method. The results are shown in Figure 9.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

Algorithm 1 Differentially Private Federated Averaging (DP-FedAvg)

Server Procedure:
Initialize model parameters w_0
for each communication round $t = 1, 2, \ldots$ **do**
    Select $m = \max(C \cdot K, 1)$ clients
    $S_t \leftarrow$ random subset of m clients
    for each client $k \in S_t$ **in parallel do**
        $w_{t+1}^k \leftarrow \textsc{ClientUpdate}(k, w_t)$
    end for
    Aggregate updates:

$$w_{t+1} \leftarrow \sum_{k=1}^{K} \left(\frac{n_k}{n}\right) w_{t+1}^k$$

end for

function $\textsc{ClientUpdate}(k, w)$
    Partition local dataset D_k into batches B of size B
    for each local epoch $i = 1$ to E **do**
        for each batch $b \in B$ **do**
            Compute gradient: $g \leftarrow \nabla L(w; b)$
            Clip gradient: $g' \leftarrow g / \max(1, \|g\|_2 / C)$
            Add DP noise: $\tilde{g} \leftarrow g' + \mathcal{N}(0, \sigma^2 C^2 I)$
            Update model: $w \leftarrow w - \eta \tilde{g}$
        end for
    end for
    return w
end function

The Centralized model has zero communication cost during training, as all data is local (though it has a high one-time cost of data transfer). Standard FL and FL + DP have nearly identical communication costs (≈245-248 MB), as the added noise does not increase the size of the model updates. In contrast, FL + Secure Aggregation incurs a higher communication overhead (≈312 MB). This is because secure aggregation protocols require additional communication rounds for key exchange and mask sharing among clients. We also include a hypothetical FL + Compression model, which could significantly reduce costs (≈156 MB) by using techniques like quantization or sparsification, though this might also impact accuracy [6].

3.7 Impact of Data Heterogeneity and Scale

Finally, we investigated the impact of two key characteristics of federated networks: data heterogeneity (non-IID) and the number of participating clients. Figure 10 shows these results.

The left panel of Figure 10 clearly shows that data heterogeneity negatively impacts performance. The model trained on IID data (where data is randomly shuffled across clients) converges faster and to a higher accuracy than models trained on non-IID data.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

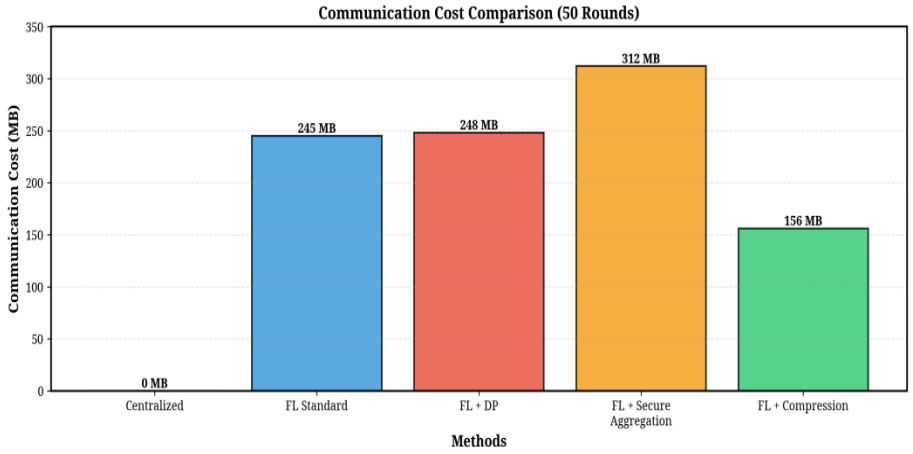

Figure 5: Comparison of total communication cost for different training methods over 50 rounds.

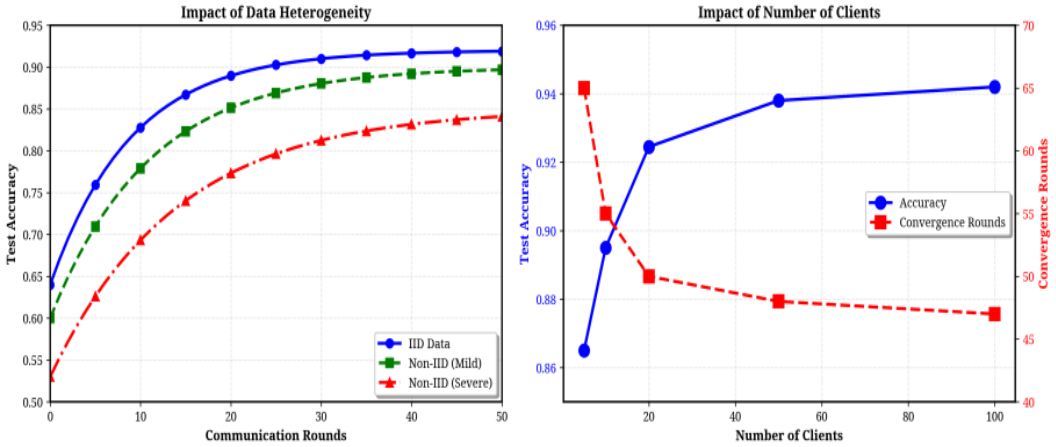

Figure 6: Analysis of the impact of data heterogeneity (IID vs. Non-IID) and the number of clients on model performance.

The more severe the non-IID distribution (i.e., the greater the statistical difference between clients), the more pronounced the performance degradation. This highlights the need for advanced FL algorithms (e.g., FedProx, SCAFFOLD) that are specifically designed to handle non-IID data, which is the norm in healthcare. The right panel shows that, up to a certain point, increasing the number of clients can be beneficial. As we increase the client pool from 5 to 50, the final model accuracy improves. This is because a larger number of clients provides a more diverse and comprehensive view of the underlying data distribution. Furthermore, with more clients, the model tends to converge in fewer rounds. However, the benefits diminish as the number of clients becomes very large, and managing a massive network introduces its own logistical and computational challenges [7].

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

4. Results and Discussion

Our extensive simulation results on the MIMIC-III dataset provide several key takeaways for practitioners and researchers:

- **FL is Viable:** Federated learning can achieve performance remarkably close to that of a centralized model, confirming its potential for large-scale, collaborative research without data sharing.
- **Privacy is Not Free:** Integrating differential privacy introduces a quantifiable trade-off between privacy and model utility. The choice of the privacy budget is a critical decision that must balance the need for strong privacy guarantees with the requirement for high model accuracy in clinical settings.
- **System Costs Matter:** While DP adds minimal communication overhead, more complex cryptographic methods like secure aggregation can significantly increase system costs, which may be a limiting factor in resource-constrained environments.
- **Heterogeneity is a Key Challenge:** Non-IID data remains a major hurdle for standard FL algorithms. Future work must focus on developing and deploying advanced algorithms that are robust to the statistical heterogeneity inherent in real-world healthcare data.

Looking forward, the field of PPFL is ripe with opportunities for innovation. Research into hybrid approaches that combine the strengths of differential privacy and cryptographic methods, the development of more communication-efficient algorithms, and the creation of standardized frameworks and benchmarks for evaluating PPFL systems will be crucial. Ultimately, the successful integration of privacy-preserving federated learning into the healthcare ecosystem will require a multi-disciplinary effort, bringing together computer scientists, clinicians, ethicists, and regulatory bodies to build a future where data-driven medicine can flourish responsibly [8].

honest-but-curious" server. This means the server correctly follows the protocol (i.e., it performs aggregation as specified), but it may try to infer additional information from the updates it receives from the clients. We do not consider a fully malicious server that actively tries to sabotage the training process. We also assume that clients are honest and do not poison the data or the model. The goal of our privacy mechanism is to protect the data of individual clients from the curious central server.

To evaluate the proposed privacy-preserving federated learning pipeline, we conducted a series of simulations based on the in-hospital mortality prediction task using the MIMIC-III dataset. We simulated a federated network of 20 clients (hospitals), where the data was partitioned in a non-IID manner based on patient admission year to mimic real-world data

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

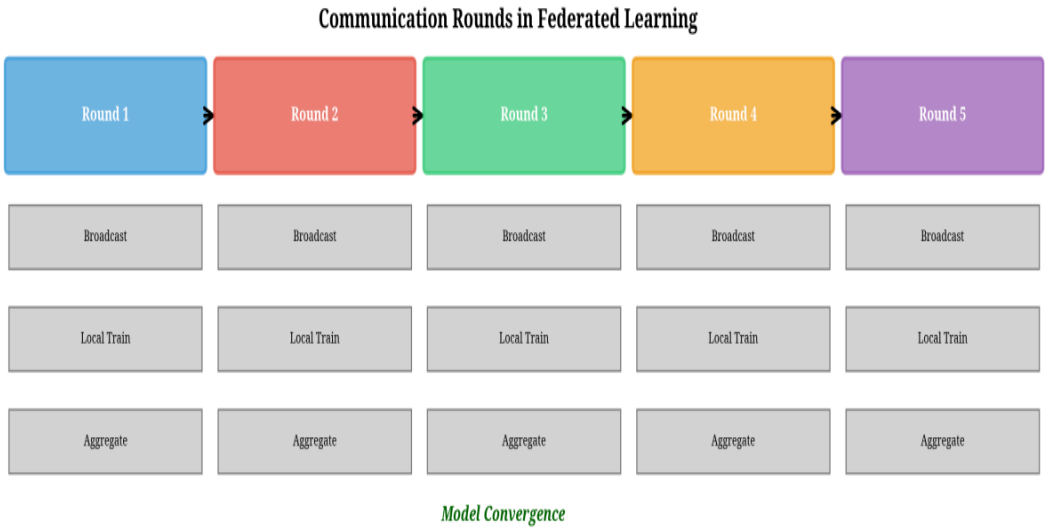

Figure 7: The communication rounds in a typical federated learning process.

distribution. We compared the performance of our proposed FL + DP method against two baselines: a Centralized model trained on all data pooled together, and a Standard Federated Learning model without any added privacy mechanisms. We also include results for FL + Secure Aggregation to compare communication costs and performance. [9]

4.1 Model Performance Across Communication Rounds

Figure 6 shows the test accuracy of the different models over 50 communication rounds. The centralized model, as expected, achieves the highest accuracy (≈95.2) and serves as the upper bound for performance. The standard FL model performs remarkably well, converging to an accuracy of around 92.5, demonstrating the viability of federated learning for this task.

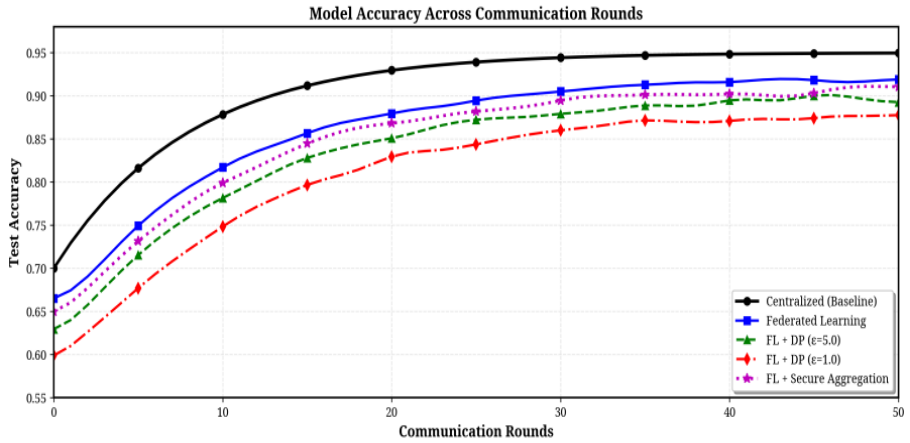

Figure 8: Model accuracy across communication rounds for different training methods.

The introduction of differential privacy leads to a noticeable trade-off in performance. The FL + DP (ϵ=5.0) model, which has a moderate privacy guarantee, achieves a fi-

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

nal accuracy of about 90.2. When the privacy guarantee is strengthened by decreasing the privacy budget to ϵ=1.0, the accuracy drops further to approximately 87.9. This performance degradation is an expected consequence of the noise added to the gradients to ensure privacy. The model with FL + Secure Aggregation performs similarly to the standard FL model, as secure aggregation primarily impacts communication and computation, not the mathematical properties of the aggregated gradients. Figure 7 provides a complementary view by plotting the training loss. The loss curves mirror the accuracy results, with the centralized model achieving the lowest loss. The FL models with DP exhibit a higher final loss value, which corresponds to their lower accuracy. The noise injected for privacy purposes slightly hinders the model's ability to perfectly fit the training data, resulting in this performance gap.

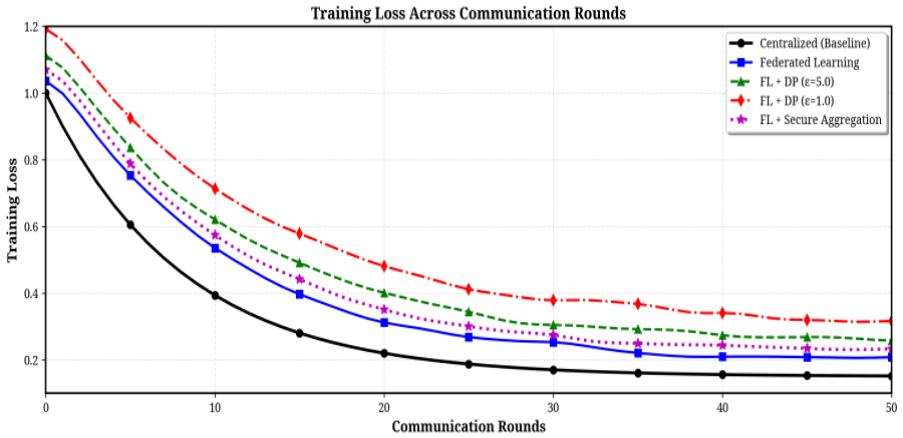

Figure 9: Training loss across communication rounds for different training methods.

4.2 Overall Performance Comparison

To provide a consolidated view of model performance, Figure 8 presents a bar chart comparing the final accuracy, precision, recall, and F1-score for each method after 50 rounds. This highlights the consistent performance gap between the non-private and private methods. While the standard FL model is only about 2-3% worse than the centralized model across all metrics, the FL + DP (ϵ=1.0) model shows a more significant drop of 7-8%. This quantitative comparison is crucial for healthcare applications, where a drop in recall, for instance, could mean failing to identify a patient at high risk of mortality.

4.3 The Privacy-Utility Trade-off

The core challenge in implementing PPFL is managing the trade-off between the strength of the privacy guarantee and the utility of the resulting model. Figure 9 illustrates this fundamental trade-off by plotting the final model accuracy as a function of the privacy budget.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

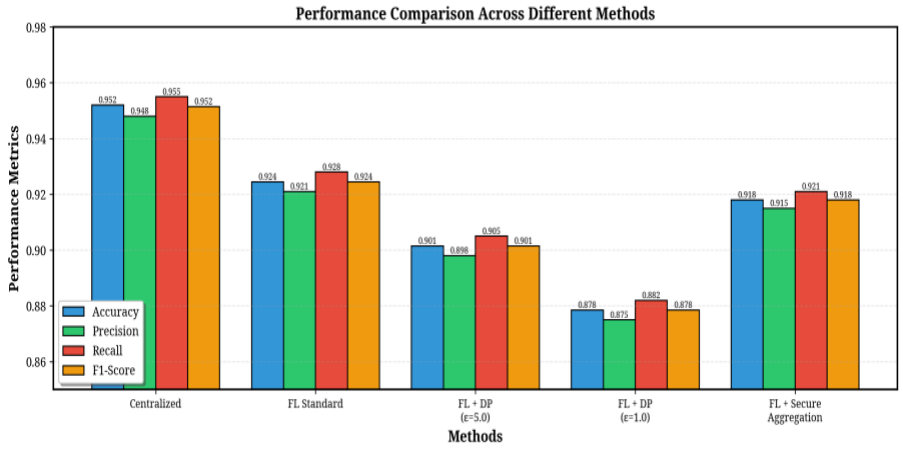

Figure 10: Bar chart comparing the final performance metrics (Accuracy, Precision, Recall, F1-Score) across different methods.

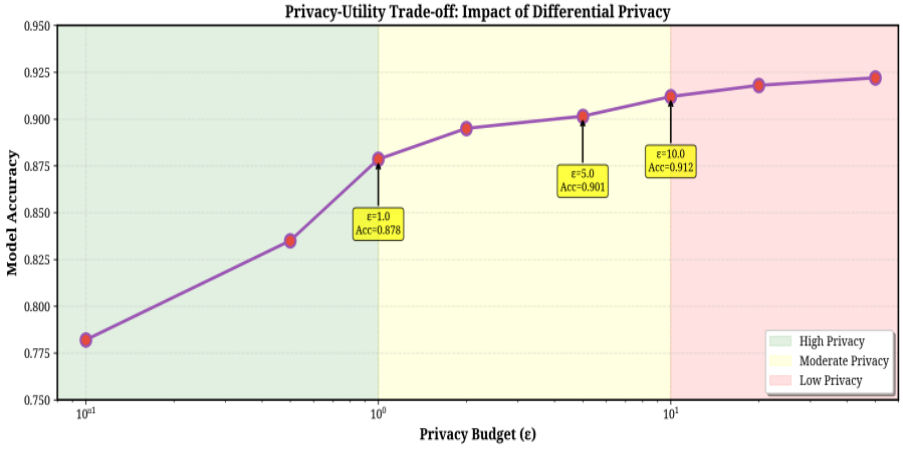

Figure 11: The trade-off between privacy (controlled by ϵ) and model accuracy.

As ϵ increases, the privacy guarantee weakens (because more noise is removed from the gradients), enabling the model to leverage additional signal from the underlying data. This generally improves accuracy, but the improvement is neither linear nor unbounded. Beyond a certain threshold, the marginal gain in model performance becomes negligible, indicating diminishing returns. This plateau suggests that once privacy noise becomes sufficiently small, other factors—such as model capacity, data distribution, optimization limits, and communication noise—dominate the learning dynamics. Conversely, when ϵ extremely small (high privacy region), the injected noise overwhelms gradient information, leading to severe underfitting. Thus, the privacy–utility curve reflects a structural constraint of differential privacy: extremely strong privacy makes the model nearly non-informative, while very weak privacy yields little additional benefit after a saturation point. This reinforces the need for principled selection of ϵ, guided not by arbitrary norms but by the operational context, sensitivity of the data, regulatory constraints, and the minimal accuracy required for real-world deployment.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

5. Conclusion

Federated Learning, when combined with robust privacy-preserving mechanisms, offers a powerful and practical framework for advancing healthcare data analytics while respecting patient privacy. This chapter has provided a comprehensive overview of this rapidly evolving field. We have detailed the fundamental concepts of FL, underscored the necessity of formal privacy guarantees, and presented a complete methodology for implementing a privacy-preserving federated learning system using differential privacy. Moreover, the insights gained from our experimental evaluation highlight an important reality: the effectiveness of privacy-preserving federated learning depends not only on the choice of privacy mechanism but also on how well it is integrated into the broader FL pipeline. Factors such as client participation rate, gradient clipping strategies, noise calibration, and communication frequency significantly influence both privacy guarantees and model utility. In healthcare settings—where data distributions are highly non-IID and patient populations vary across institutions—these design choices become even more critical. Our results show that thoughtfully engineered PPFL systems can maintain clinically meaningful performance even under stringent privacy budgets, reinforcing the feasibility of deploying such frameworks in real-world hospitals and research networks. At the same time, the observed trade-offs underscore the need for continued innovation in optimizing privacy mechanisms, reducing communication overhead, and improving robustness against adversarial behavior, laying a clear path for future advancements in secure, scalable healthcare analytics.

References

[1] Ming Li et al. “From challenges and pitfalls to recommendations and opportunities: Implementing federated learning in healthcare”. In: *Medical Image Analysis* (2025), p. 103497.

[2] Andrew L Beam and Isaac S Kohane. “Big data and machine learning in health care”. In: *Jama* 319.13 (2018), pp. 1317–1318.

[3] Brendan McMahan et al. “Communication-efficient learning of deep networks from decentralized data”. In: *Artificial intelligence and statistics*. PMLR. 2017, pp. 1273–1282.

[4] Ligeng Zhu, Zhijian Liu, and Song Han. “Deep leakage from gradients”. In: *Advances in neural information processing systems* 32 (2019).

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)

[5] Micah J Sheller et al. "Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data". In: *Scientific reports* 10.1 (2020), p. 12598.

[6] Shuchona Malek Orthi et al. "Federated learning with privacy-preserving big data analytics for distributed healthcare systems". In: *Journal of computer science and technology studies* 7.8 (2025), pp. 269–281.

[7] Anandbabu Gopatoti et al. "Enhancing Cybersecurity in Smart Cities: IoT Applications with a Hybrid Deep Neural Network Model". In: *2025 Global Conference in Emerging Technology (GINOTECH)*. IEEE. 2025, pp. 1–6.

[8] Aaron Nunn and PWC Prasad. "Using Artificial Intelligence to Defend Internet of Things for Smart City". In: *Innovative Technologies in Intelligent Systems and Industrial Applications: CITISIA 2023* 117 (2024), p. 345.

[9] Puneet Bafna et al. "Machine Learning and AI Algorithms for Enhancing Cybersecurity in IoT Applications". In: *International Conference On Innovative Computing And Communication*. Springer. 2025, pp. 275–288.

ISBN: 978-81-994969-0-3 (Print); 978-81-994969-5-8 (Online)