

Federated Learning for Privacy-Preserving Systems

Anjana Bala¹, Kanchan Jyoti²

¹Assistant Professor, Dept. of CSE, SBSSU Gurdaspur, Punjab, India

²Assistant Professor, Dept. of CSE, SBSSU Gurdaspur, Punjab, India

Abstract - Federated Learning (FL) has emerged as a revolutionary paradigm in distributed machine learning that enables multiple clients to collaboratively train models without sharing raw data. This approach is particularly significant in privacy-sensitive domains such as healthcare, finance, and IoT systems. Traditional centralized learning methods expose sensitive data to risks such as breaches, misuse, and regulatory violations. FL addresses these challenges by ensuring data locality and enabling decentralized model training. However, despite its inherent privacy-preserving capabilities, FL is vulnerable to various attacks, including inference and poisoning attacks. This paper provides a comprehensive overview of FL, its architecture, privacy-preserving techniques, challenges, and future directions. A detailed literature survey is presented to highlight recent advancements and research gaps. The study concludes by proposing potential improvements for robust and secure FL systems.

Key Words: Federated Learning, Privacy Preservation, Differential Privacy, Secure Aggregation, Distributed Machine Learning

1. INTRODUCTION

The rapid growth of machine learning applications has significantly increased the demand for large-scale data. However, centralized data collection raises serious privacy concerns, especially in domains involving sensitive information such as healthcare and finance. Federated Learning (FL) addresses these concerns by enabling decentralized training where data remains on local devices and only model updates are shared (Xie, Koyejo and Gupta, 2019).

FL operates on a collaborative framework where multiple clients train local models and send updates to a central server for aggregation (Albshaier, Almarri and Albuali, 2025). This ensures that raw data never leaves the device, thereby preserving user privacy. The paradigm is particularly useful in distributed environments such as edge computing and IoT systems (Wang *et al.*, 2020).

Despite its advantages, FL does not guarantee complete privacy. Model updates may still leak sensitive information through inference attacks. Therefore, integrating additional privacy-preserving mechanisms such as differential privacy and secure aggregation is essential (Agrahari, Dinker and Singh, 2026).




Figure 1: Federated learning with privacy learning

2. LITERATURE SURVEY

Federated Learning has gained substantial attention in recent years due to its potential to enable privacy-preserving machine learning. Numerous studies have explored its architecture, challenges, and privacy-enhancing techniques. One of the foundational works in FL highlights its ability to perform distributed learning without sharing raw data, making it suitable for privacy-sensitive applications. (Long *et al.*, 2020) This paradigm is particularly beneficial in healthcare systems, where patient data confidentiality is critical. (White and Legg, 2024) Studies have demonstrated that FL can train models across multiple institutions while preserving data privacy and maintaining model performance.

(Xu *et al.*, 2020) conducted a comprehensive survey focusing on privacy preservation in FL systems, emphasizing compliance with regulations such as GDPR. (Odeh *et al.*, 2025) The study identified that although FL reduces the need for centralized data storage, it is still susceptible to privacy attacks such as model inversion and membership inference. The authors highlighted the need for integrating additional privacy-preserving techniques to achieve regulatory compliance.

Another significant contribution is the work on privacy-preserving aggregation protocols, which are essential for secure model updates. These protocols ensure that individual updates cannot be inferred from aggregated results. (Li *et al.*, 2020) reviewed various aggregation techniques, including secure multi-party computation (SMC) and homomorphic encryption, and analyzed their trade-offs

in terms of computational complexity and security. Recent research has focused on combining FL with differential privacy (DP) to enhance privacy guarantees. (Kairouz *et al.*, 2021) Differential privacy introduces noise into model updates, ensuring that individual data points cannot be identified. However, this approach often results in a trade-off between privacy and model accuracy. Studies indicate that lower privacy budgets improve privacy but degrade performance.

In addition to DP, homomorphic encryption has been widely explored as a privacy-preserving technique in FL. This method allows computations to be performed on encrypted data, ensuring that sensitive information remains protected during processing. (Liu *et al.*, 2021) demonstrated that combining homomorphic encryption with secure multi-party computation can provide strong privacy guarantees while maintaining efficiency.

Another emerging area of research is the application of FL in recommender systems. Traditional recommender systems rely heavily on centralized user data, raising privacy concerns. Federated recommender systems address this issue by training models locally and sharing only intermediate parameters. (Toyoda *et al.*, 2020) Studies have shown that this approach improves privacy while maintaining recommendation accuracy.

Tree-based models in FL have also been explored for privacy preservation. A recent survey (Ramanan and Nakayama, 2020) analyzed the integration of decision trees and ensemble methods with FL. The study highlighted the challenges of adapting tree-based algorithms to distributed settings while ensuring privacy through encryption and differential privacy techniques.

Furthermore, FL has been widely applied in intrusion detection systems (IDS) to enhance cybersecurity. Research indicates that while FL improves data privacy by keeping data local, it is still vulnerable to attacks such as data poisoning and adversarial manipulation. The study emphasized the need for robust defense mechanisms to ensure secure FL-based IDS systems (Martinez, Francis and Hafid, 2019).

Fairness and bias in FL systems have also been investigated. Studies reveal that ensuring fairness across diverse clients is challenging due to data heterogeneity. Privacy-preserving techniques may further exacerbate bias, leading to unequal model performance across different user groups. Researchers have highlighted the importance of balancing privacy, fairness, and accuracy in FL systems (Desai, Ozdayi and Kantarcioglu, 2021).

Recent surveys have also explored the role of FL in edge intelligence systems. FL enables distributed learning across edge devices, reducing latency and improving privacy by keeping data local. However, challenges such as limited computational resources, communication overhead, and

system heterogeneity remain significant barriers to large-scale deployment (Korkmaz *et al.*, 2020).

In healthcare, FL has demonstrated promising results in enabling collaborative learning across hospitals without sharing patient data. Studies show that FL can achieve comparable performance to centralized models while preserving privacy. However, challenges such as data heterogeneity and communication efficiency need to be addressed for real-world implementation (Zhang *et al.*, 2020).

Overall, the literature indicates that while FL provides a strong foundation for privacy-preserving systems, it is not inherently secure. Additional techniques such as differential privacy, secure aggregation, and encryption are essential to mitigate risks. Moreover, challenges related to scalability, communication efficiency, and fairness must be addressed to fully realize the potential of FL.

3. PROBLEM DEFINITION

Despite its advantages, Federated Learning faces several challenges:

- **Privacy Leakage:** Model updates can reveal sensitive information.
- **Communication Overhead:** Frequent communication between clients and server.
- **Data Heterogeneity:** Non-IID data affects model performance.
- **Security Threats:** Vulnerability to poisoning and inference attacks.
- **Scalability Issues:** Managing large-scale distributed systems.

4. PROPOSED SOLUTIONS

To effectively address the challenges associated with Federated Learning (FL), particularly in privacy-preserving systems, several advanced techniques have been proposed. These solutions aim to enhance data security, improve model performance, and ensure scalability in distributed environments.

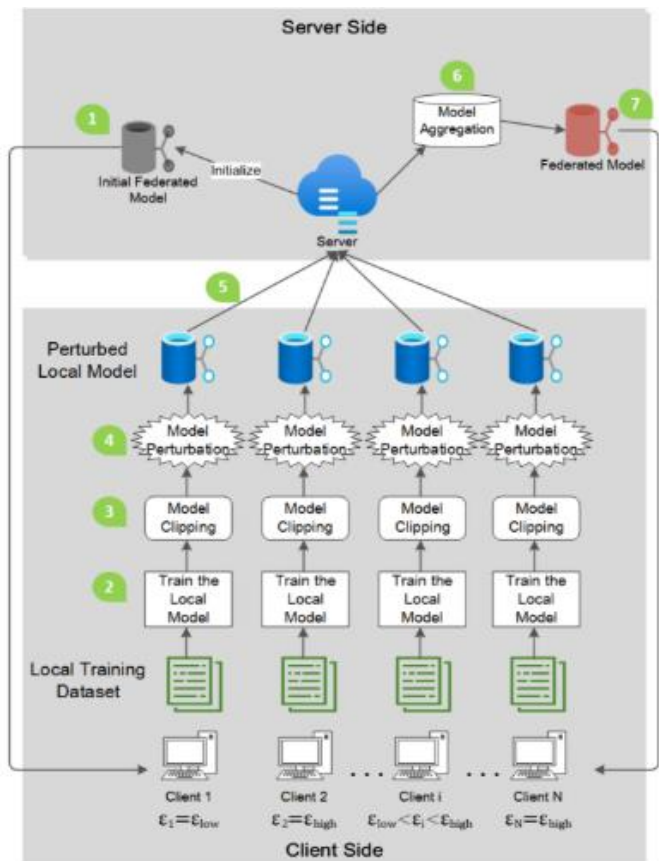


Figure 2: Proposed methodology

4.1 Differential Privacy

Differential Privacy (DP) is one of the most widely adopted techniques for protecting sensitive information in FL systems. It works by introducing carefully calibrated noise into the model updates before they are shared with the central server. This ensures that the contribution of any individual data point cannot be distinguished, thereby preventing attackers from inferring private information (Dwork et al., 2014). In FL, local differential privacy is often applied at the client level, ensuring privacy even before data leaves the device. However, there exists a trade-off between privacy and model accuracy, as excessive noise can degrade performance. Therefore, selecting an optimal privacy budget (ϵ) is crucial for balancing privacy protection and utility.

4.2 Secure Aggregation

Secure aggregation is a cryptographic technique designed to protect individual client updates during the aggregation process. In traditional FL, model updates are sent to a central server, which aggregates them to update the global model. However, these updates may leak sensitive information. Secure aggregation protocols ensure that the server can only access the aggregated result, not individual contributions (Bonawitz et al., 2017). This is typically achieved using encryption schemes where each client masks its update with random values that cancel out during aggregation. As a result, even if the server is compromised, individual data

remains protected. Secure aggregation is essential for maintaining confidentiality in large-scale distributed systems.

4.3 Homomorphic Encryption

Homomorphic Encryption (HE) allows computations to be performed directly on encrypted data without requiring decryption. In the context of FL, this enables the server to aggregate encrypted model updates while preserving data confidentiality. This technique provides a strong level of security, as sensitive information is never exposed during computation (Gentry, 2009). HE is particularly useful in highly sensitive domains such as healthcare and finance. However, it introduces significant computational and communication overhead, making it less practical for resource-constrained environments. Recent advancements in partial and approximate homomorphic encryption aim to reduce this overhead while maintaining acceptable security levels.

4.4 Personalized Federated Learning

One of the major challenges in FL is data heterogeneity, where data distributions vary significantly across clients. Personalized Federated Learning addresses this issue by allowing each client to maintain a customized version of the global model. Instead of enforcing a single global model, personalization techniques adapt the model to local data characteristics, improving performance and user-specific accuracy (Smith et al., 2017). Approaches such as fine-tuning, meta-learning, and multi-task learning are commonly used in this context. Personalized FL not only enhances model accuracy but also improves user satisfaction by tailoring predictions to individual needs.

4.5 Communication Optimization

Communication overhead is a critical bottleneck in FL systems due to frequent exchange of model updates between clients and the server. Communication optimization techniques aim to reduce this overhead while maintaining model performance. Methods such as model compression, quantization, and sparsification are commonly employed to minimize the size of transmitted updates (Konečný et al., 2016). Additionally, strategies like reducing communication rounds and using asynchronous updates can further enhance efficiency. These techniques are particularly important in edge computing environments where bandwidth and resources are limited. Efficient communication mechanisms ensure scalability and faster convergence of FL systems.

5. CONCLUSION

Federated Learning represents a promising approach for privacy-preserving machine learning systems. It eliminates the need for centralized data collection while enabling collaborative model training. However, FL alone is insufficient to guarantee complete privacy. Integrating advanced privacy-preserving techniques such as differential privacy and secure aggregation is essential. Future research

should focus on improving scalability, robustness, and fairness in FL systems. With continued advancements, FL has the potential to revolutionize privacy-preserving AI across various domains.

REFERENCES

1. Agrahari, A.K., Dinker, A.G. and Singh, R.B. (2026) "A review of security threats and privacy issues in federated learning," *International Journal of Data Science and Analytics* 2026 22:1, 22(1), pp. 85-. Available at: <https://doi.org/10.1007/S41060-026-01067-Z>.
2. Albshaiyer, L., Almarri, S. and Albuali, A. (2025) "Federated Learning for Cloud and Edge Security: A Systematic Review of Challenges and AI Opportunities," *Electronics (Switzerland)*, 14(5), p. 1019. Available at: <https://doi.org/10.3390/ELECTRONICS14051019/S1>.
3. Desai, H.B., Ozdayi, M.S. and Kantarcioglu, M. (2021) "BlockFLA: Accountable Federated Learning via Hybrid Blockchain Architecture," *CODASPY 2021 - Proceedings of the 11th ACM Conference on Data and Application Security and Privacy*, pp. 101-112. Available at: <https://doi.org/10.1145/3422337.3447837>.
4. Kairouz, P. et al. (2021) "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, 14(1-2), pp. 1-210. Available at: <https://doi.org/10.1561/22000000083>.
5. Korkmaz, C. et al. (2020) "Chain FL: Decentralized Federated Machine Learning via Blockchain," *2020 2nd International Conference on Blockchain Computing and Applications, BCCA 2020*, pp. 140-146. Available at: <https://doi.org/10.1109/BCCA50787.2020.9274451>.
6. Li, T. et al. (2020) "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Processing Magazine*, 37(3), pp. 50-60. Available at: <https://doi.org/10.1109/MSP.2020.2975749>.
7. Liu, M. et al. (2021) "Federated Learning Meets Natural Language Processing: A Survey." Available at: <http://arxiv.org/abs/2107.12603> (Accessed: July 23, 2024).
8. Long, G. et al. (2020) "Federated Learning for Open Banking," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12500 LNCS, pp. 240-254. Available at: https://doi.org/10.1007/978-3-030-63076-8_17.
9. Martinez, I., Francis, S. and Hafid, A.S. (2019) "Record and reward federated learning contributions with blockchain," *Proceedings - 2019 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, CyberC 2019*, pp. 50-57. Available at: <https://doi.org/10.1109/CYBERC.2019.00018>.
10. Odeh, A. et al. (2025) "Enhancing Data Security and Privacy Using Federated Learning: A Scalable Framework for Distributed Systems," *Communications in Computer and Information Science*, 2652 CCIS, pp. 145-157. Available at: https://doi.org/10.1007/978-3-032-04228-6_11.
11. Ramanan, P. and Nakayama, K. (2020) "BAFFLE : Blockchain Based Aggregator Free Federated Learning," *Proceedings - 2020 IEEE International Conference on Blockchain, Blockchain 2020*, pp. 7281. Available at: <https://doi.org/10.1109/BLOCKCHAIN50366.2020.00017>.
12. Toyoda, K. et al. (2020) "Blockchain-enabled federated learning with mechanism design," *IEEE Access*, 8, pp. 219744-219756. Available at: <https://doi.org/10.1109/access.2020.3043037>.
13. Wang, H. et al. (2020) "FEDERATED LEARNING WITH MATCHED AVERAGING," *8th International Conference on Learning Representations, ICLR 2020 [Preprint]*.
14. White, J. and Legg, P. (2024) "Evaluating Data Distribution Strategies in Federated Learning: A Trade-Off Analysis Between Privacy and Performance for IoT Security," *Lecture Notes in Networks and Systems*, 1032 LNNS, pp. 17-37. Available at: https://doi.org/10.1007/978-981-97-3973-8_2.
15. Xie, C., Koyejo, S. and Gupta, I. (2019) "Asynchronous Federated Optimization." Available at: <http://arxiv.org/abs/1903.03934> (Accessed: July 23, 2024).
16. Xu, J. et al. (2020) "Federated learning for healthcare informatics," *J Healthc Inform Res*, 5(1), pp.119. Available at: <https://doi.org/10.1007/s41666-020-00082-4>.
17. Zhang, Q. et al. (2020) "Demo: A Blockchain Based Protocol for Federated Learning," *Proceedings - International Conference on Network Protocols, ICNP, 2020 October*. Available at: <https://doi.org/10.1109/ICNP49622.2020.9259388>.