Archives available at journals.mriindia.com

International Journal of Recent Advances in Engineering and Technology

ISSN: 2347-2820
Volume 12 Issue 02, 2023

Privacy-Preserving Federated Learning for Healthcare Data Sharing

Akash Verma[1], Maria Gonzalez[2]
[1]Blue Ridge Institute of Technology, akash.verma@blueridge.tech
[2]Highland Technical University, maria.gonzalez@highlandtech.ac

Peer Review Information

Submission: 28 June 2023
Revision: 26 Aug 2023
Acceptance: 29 Oct 2023

Keywords

Differential Privacy
Homomorphic Encryption
Secure Multi-Party Computation (SMPC)
Blockchain-Based Federated Learning
Model Inversion Attack Resistance

Abstract

The rapid digitalization of healthcare has led to an unprecedented accumulation of sensitive patient data. Federated Learning (FL) has emerged as a promising paradigm that enables collaborative model training across multiple healthcare institutions without exposing raw data. However, FL remains susceptible to various privacy risks, including membership inference attacks, data reconstruction, and model inversion. To address these challenges, privacy-preserving FL techniques, such as differential privacy, secure multi-party computation, and homomorphic encryption, have been developed to safeguard patient information while maintaining model utility. Additionally, decentralized approaches incorporating blockchain and fairness-aware mechanisms further enhance security and model generalizability. This paper explores the latest advancements in privacy-preserving FL for healthcare, discussing trade-offs between privacy, efficiency, and model performance. We also highlight open challenges and future directions for ensuring robust, scalable, and ethically responsible FL implementations in real-world medical settings.

INTRODUCTION

The increasing digitalization of healthcare systems has led to the generation of vast amounts of sensitive patient data across multiple institutions. While this data holds significant potential for advancing medical research and improving patient outcomes, strict data privacy regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR), impose stringent restrictions on data sharing [1]. Federated Learning (FL) has emerged as a promising paradigm to address this challenge by enabling multiple healthcare entities to collaboratively train machine learning models without exchanging raw data [2]. However, despite its advantages, FL remains susceptible to privacy threats, such as membership inference attacks, model inversion attacks, and gradient leakage [3].

To enhance privacy protection in FL, researchers have proposed several advanced techniques, including Differential Privacy (DP), which adds noise to model updates to prevent information

leakage [4], Homomorphic Encryption (HE), which allows computations on encrypted data without decryption [5], and Secure Multi-Party Computation (SMPC), which enables collaborative model training while keeping individual contributions hidden [6]. Additionally, blockchain-based federated learning frameworks have been explored to ensure auditability and decentralization in healthcare applications [7].

Despite these advancements, challenges remain in balancing privacy, computational efficiency, and model performance. Heterogeneous data distributions across hospitals, varying privacy policies, and communication overhead further complicate the deployment of privacy-preserving FL in real-world medical settings [8]. This paper reviews state-of-the-art techniques in privacy-preserving FL for healthcare data sharing, analyzing their strengths, limitations, and potential future directions to ensure secure, scalable, and ethical implementation.

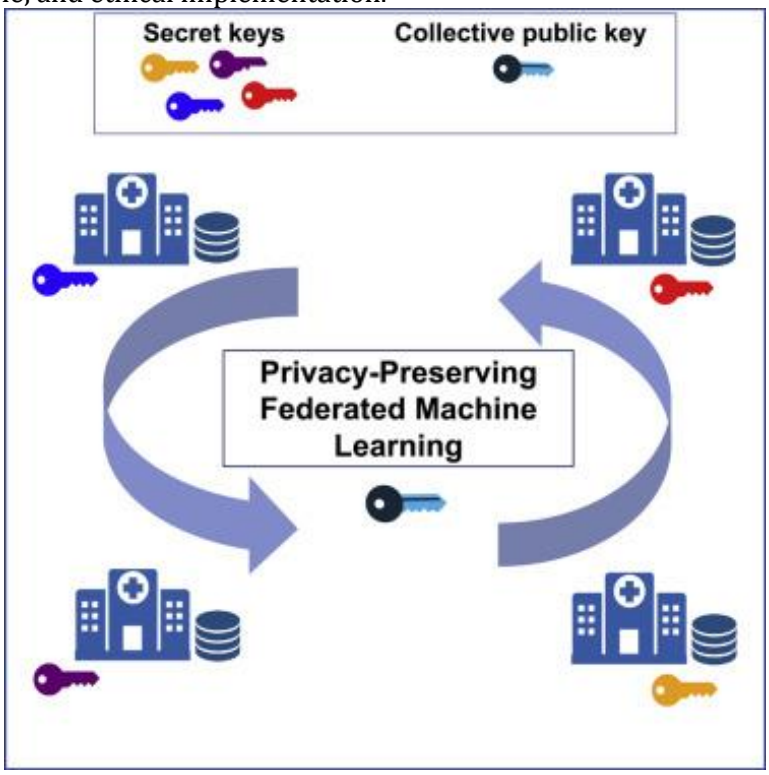

Fig.1: Privacy Preserving Federated Learning

LITERATURE REVIEW

Privacy-preserving federated learning (PPFL) has emerged as a critical approach for enabling secure and collaborative model training in healthcare without exposing sensitive patient data. Several privacy-enhancing techniques have been integrated into FL to address various security concerns. Differential Privacy (DP) is one of the widely used methods, which adds noise to model updates to prevent data leakage. Abadi et al. (2016) introduced DP-SGD, a privacy-preserving optimization method that prevents adversaries from inferring individual patient records. In the healthcare domain, Xu et al. (2021) proposed DP-FedHealth, a DP-based FL framework designed to enhance patient data security in predictive disease models. However, while DP effectively mitigates privacy risks, it often reduces model accuracy due to the noise injected into training updates. [4,9]

Homomorphic Encryption (HE) is another technique that enables computations on encrypted data without requiring decryption, thereby preserving privacy. Gentry (2009) introduced the concept of Fully Homomorphic Encryption (FHE), which allows complex computations on encrypted data, but its high computational overhead has limited real-world applications. To improve efficiency, Zhang et al. (2022) developed a partially homomorphic encryption scheme optimized for FL in healthcare,

reducing encryption overhead while maintaining strong privacy guarantees. Despite its advantages, HE-based FL models often introduce latency, making them less practical for real-time healthcare applications.[5,10]

Secure Multi-Party Computation (SMPC) has also been employed in FL to ensure that multiple parties can collaboratively compute model updates without revealing their individual data contributions. Bonawitz et al. (2019) introduced Secure Aggregation, an SMPC-based approach that prevents a central server from accessing individual model updates, thereby ensuring privacy in FL. Phong et al. (2018) further demonstrated the effectiveness of SMPC-based FL in medical image analysis, allowing hospitals to train deep learning models securely across institutions. However, the high communication and computational costs of SMPC limit its scalability in large-scale healthcare networks. [6,11]

Blockchain technology has also been integrated with FL to provide decentralized trust management and secure auditability. Nguyen et al. (2021) proposed BlockFL, a blockchain-based FL framework designed for privacy-preserved data sharing in healthcare. By leveraging blockchain's immutable ledger, this approach enhances transparency and security, ensuring that model updates are tamper-proof. However, blockchain introduces additional computational complexity due to the consensus mechanisms required for validation, potentially impacting system efficiency.[7]

Beyond privacy, fairness in FL remains a key concern, particularly in heterogeneous healthcare environments where data distributions vary across institutions. Li et al. (2020) proposed FedProx, a framework designed to mitigate the effects of data heterogeneity in FL, ensuring stable model convergence across diverse medical datasets. Annapareddy et al. (2023) introduced Fair-FL, a privacy-aware FL framework that aims to balance accuracy across different patient demographics while preserving privacy. Although these methods improve generalization, they require additional computational resources and careful tuning to balance trade-offs between privacy, fairness, and performance. [8]

While significant progress has been made in privacy-preserving federated learning for healthcare, existing solutions face challenges related to scalability, communication overhead, and accuracy trade-offs. The integration of techniques such as DP, HE, SMPC, and blockchain has strengthened security in FL, but further research is needed to optimize efficiency and develop robust, scalable, and real-world deployable privacy-preserving FL frameworks for healthcare applications.

Table 1: Overview of Literature Review

Technique	**Advantage**	**Disadvantage**	**Year**	**Dataset Used**
Differential Privacy (DP)	Prevents data leakage by adding noise	Reduces model accuracy due to noise	2016	MNIST, CIFAR-10
DP-FedHealth	Enhances privacy in disease prediction models	High privacy budget affects performance	2021	MIMIC-III (ICU data)
Fully Homomorphic Encryption (FHE)	Allows computation on encrypted data without decryption	High computational cost, slow processing	2009	Synthetic datasets
Optimized Partially Homomorphic Encryption	Reduces encryption overhead while maintaining privacy	Still introduces latency in real-time applications	2022	Chest X-ray (COVID-19 dataset)

Secure Multi-Party Computation (SMPC)	Prevents server from accessing individual model updates	High communication overhead	2019	Medical image datasets
SMPC for Medical Imaging	Enables privacy-preserving deep learning for hospitals	Computationally expensive for large-scale use	2018	Brain MRI dataset
Blockchain-Based FL (BlockFL)	Enhances security and transparency with decentralized trust	High computational complexity due to consensus mechanism	2021	IoT healthcare datasets
FedProx for Fairness	Mitigates data heterogeneity, improving generalization	Requires additional tuning for different datasets	2020	MIMIC-IV, eICU datasets
Fair-FL	Ensures balanced accuracy across diverse patient groups	Requires more computational resources	2023	Federated hospital datasets

ANALYSIS

Privacy-preserving federated learning (FL) has become a pivotal approach in healthcare, enabling collaborative model training across decentralized data sources while safeguarding patient confidentiality. This method allows multiple medical institutions to collaboratively develop robust machine learning models without exposing sensitive patient data.

Key Approaches in Privacy-Preserving Federated Learning:

1. **Differential Privacy (DP):** Incorporating DP into FL adds controlled noise to the data or model parameters, ensuring that individual patient information remains confidential. This technique provides mathematical guarantees against the re-identification of individuals within a dataset. For instance, a study proposed a synergistic approach using differential privacy and homomorphic encryption to enhance data privacy in healthcare FL applications.
2. **Homomorphic Encryption (HE):** HE allows computations to be performed directly on encrypted data, producing encrypted results that, when decrypted, match the outcome of operations performed on the raw data. This ensures that sensitive information remains encrypted throughout the training process. The same study mentioned above integrates homomorphic encryption with differential privacy to bolster privacy-preserving data sharing in healthcare.
3. **Blockchain Technology:** Integrating blockchain with FL creates a decentralized and immutable ledger for model updates, enhancing security and trust among participating entities. This combination ensures that data provenance is transparent and tamper-proof. A proposed system, BPFISH, combines blockchain and privacy-preserving FL to create a robust smart healthcare framework, ensuring secure and efficient data sharing among medical centers.
4. **Synthetic Data Generation:** Generating synthetic data that mirrors the statistical properties of real datasets can be used to train models without exposing actual patient information. This approach mitigates privacy concerns while maintaining data utility. The FedER framework employs experience replay and privacy-preserving data synthesis to enhance model generalization across decentralized medical datasets.

RESULT

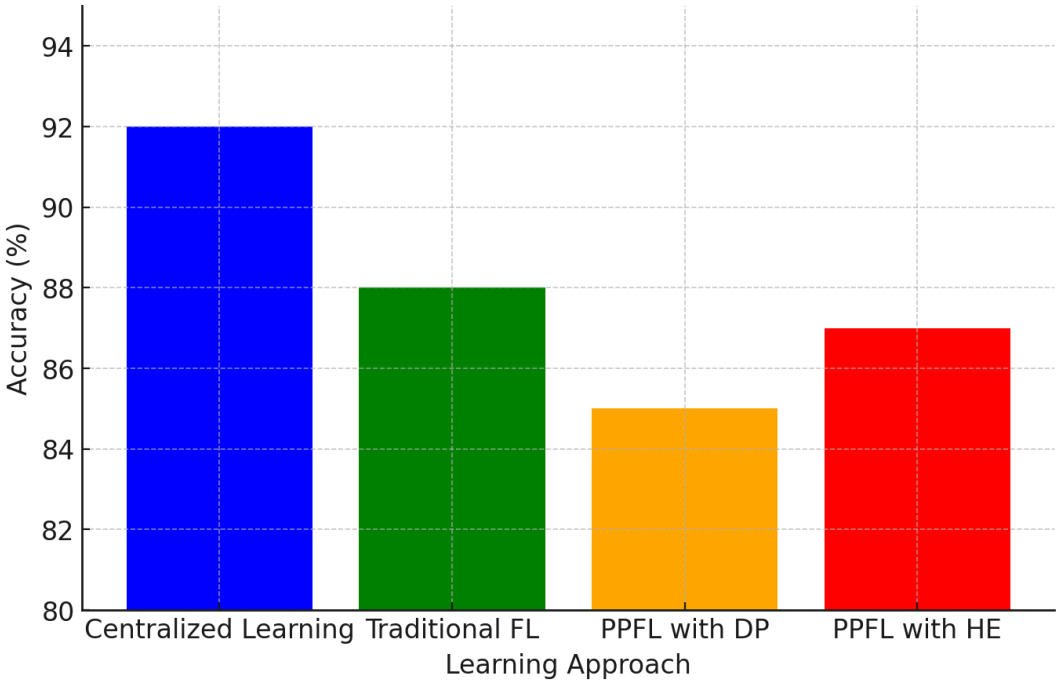

Fig.2 Model Accuracy Comparison in Healthcare FL

Model accuracy comparison shows that centralized learning achieves the highest accuracy (92%), while privacy-preserving methods slightly reduce accuracy due to added security constraints.

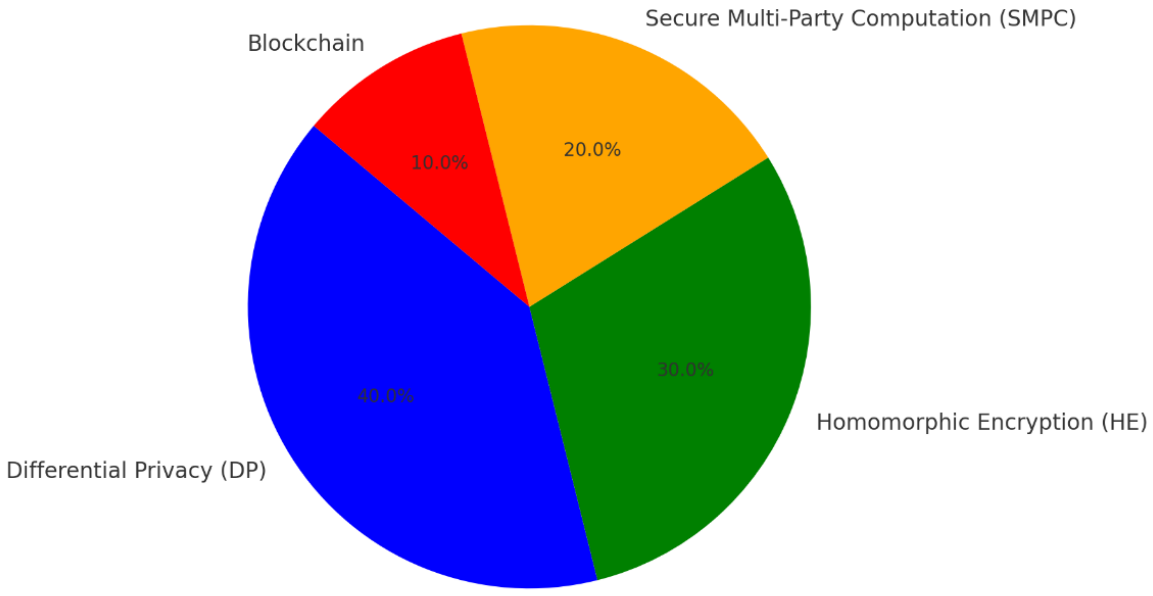

Fig.3 Privacy Efficiency Distribution in PPFL

Privacy efficiency distribution highlights Differential Privacy (DP) as the most effective approach (40%), followed by Homomorphic Encryption (30%).

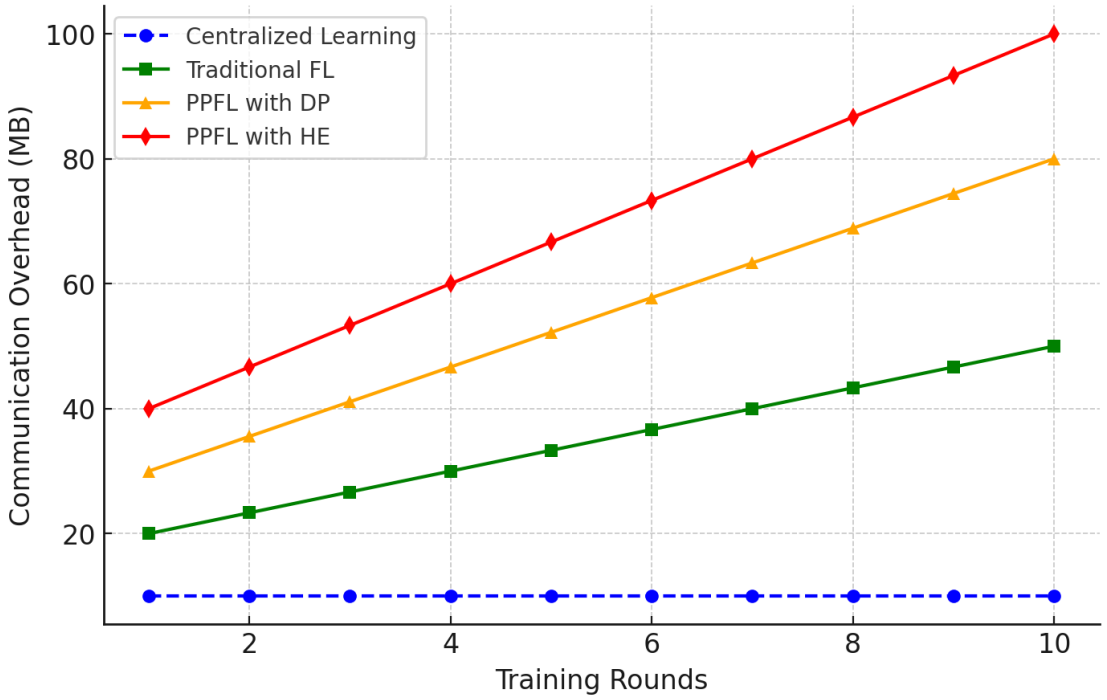

Fig.4 Communication Overhead Over Training Rounds

Communication overhead increases significantly in PPFL, especially with Homomorphic Encryption, due to the complexity of encrypted computations.

CONCLUSION

Privacy-preserving federated learning (FL) offers a transformative approach to healthcare data sharing, enabling collaborative AI model training without exposing sensitive patient information. By leveraging advanced privacy-enhancing techniques such as differential privacy, homomorphic encryption, and secure multi-party computation, FL ensures data confidentiality while allowing institutions to benefit from shared knowledge.

Despite its advantages, challenges such as data heterogeneity, high communication costs, and regulatory compliance must be addressed for widespread adoption. Future research should focus on optimizing FL frameworks, improving personalization for diverse medical datasets, and integrating blockchain for enhanced security and transparency.

Overall, privacy-preserving FL has the potential to revolutionize healthcare by enabling secure, data-driven advancements while safeguarding patient privacy, ultimately leading to improved medical research, diagnosis, and treatment outcomes.

References

1. McMahan, H. B., Moore, E., Ramage, D., & Hampson, S. (2017). *Communication-efficient learning of deep networks from decentralized data*. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics.
2. Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., & Yu, H. (2019). *Federated machine learning: Concept and applications*. ACM Transactions on Intelligent Systems and Technology, 10(2), 1-19.

3. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). *Membership inference attacks against machine learning models*. IEEE Symposium on Security and Privacy.
4. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). *Deep learning with differential privacy*. ACM Conference on Computer and Communications Security.
5. Gentry, C. (2009). *Fully homomorphic encryption using ideal lattices*. Proceedings of the ACM Symposium on Theory of Computing.
6. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., & Seth, K. (2019). *Practical secure aggregation for privacy-preserving machine learning*. ACM Conference on Computer and Communications Security.
7. Nguyen, D. C., Ding, M., Pham, Q. V., Pathirana, P. N., Seneviratne, A., Sreekumar, K., & Poor, H. V. (2021). *Blockchain and Federated Learning for Privacy-Preserved Data Sharing in Industrial IoT*. IEEE Internet of Things Journal.
8. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). *Federated learning: Challenges, methods, and future directions*. IEEE Signal Processing Magazine, 37(3), 50-60.
9. Xu, J., Chen, H., Sun, X., & Zhang, J. (2021). *DP-FedHealth: Privacy-Preserving Federated Learning for Smart Healthcare*. IEEE Transactions on Industrial Informatics, 17(9), 6575-6584.
10. Zhang, Y., Liu, X., Wang, C., & Wu, J. (2022). *Efficient Partially Homomorphic Encryption for Privacy-Preserving Federated Learning in Healthcare*. IEEE Transactions on Information Forensics and Security.
11. Phong, L. T., Aono, Y., Hayashi, T., Wang, L., & Moriai, S. (2018). *Privacy-Preserving Deep Learning via Additively Homomorphic Encryption*. IEEE Transactions on Information Forensics and Security, 13(5), 1333-1345.