

Federated Learning in Healthcare: An Analytical Study on Privacy-Preserving Data Sharing and Clinical Decision Support

Meet Kapilbhai Nayak¹, Kashish Dharmendra Nihalani²

¹Software Developer, IQIT Services Pvt Ltd, Ahmedabad, Gujarat, India

²Assistant Professor, Silver Oak University, Ahmedabad, Gujarat, India

ARTICLE INFO

Article History:

Accepted : 01 Sept 2025

Published: 25 Sept 2025

Publication Issue :

Volume 12, Issue 5

September-October-2025

Page Number :

279-284

ABSTRACT

Learning Federated, which permits cooperative model training across dispersed clinical datasets without jeopardizing patient privacy, has become a paradigm shift in healthcare. An analytical examination of FL for data sharing that protects privacy and its possible application in clinical decision support systems is presented in this work. The study investigates how data silos, legal restrictions, and security threats related to traditional centralized methods are addressed via decentralized learning frameworks. The effectiveness of FL algorithms, encryption methods, and secure aggregation strategies in safeguarding private medical data is emphasized. The study also emphasizes how FL-driven models can enhance predictive analytics, therapy customization, and diagnostic accuracy in healthcare applications. Comparative studies and experimental findings show that FL improves the scalability and dependability of clinical decision-making while maintaining anonymity. The results highlight federated learning's potential as a long-term means of promoting data-driven healthcare breakthroughs while maintaining moral and legal observance.

Keywords: Healthcare, Artificial Intelligence in Healthcare, Decentralized Machine Learning, Secure Aggregation, Clinical Decision Support Systems, Federated Learning, Privacy-Preserving Data Sharing, Collaborative Learning, and Predictive Analytics

I. INTRODUCTION

In healthcare, the application of machine learning (ML) and artificial intelligence (AI) has greatly enhanced clinical decision-making, treatment planning, and illness diagnosis. On the other hand, access to vast, varied, and high-quality datasets is

crucial for these models to function well. In actuality, healthcare data is frequently dispersed among several hospitals, research facilities, and clinical centers, which causes problems with data silos and restricted interoperability. Furthermore, Strict privacy regulations such as the Health Insurance Portability and Accountability Act (HIPAA) and the General

Data Protection Regulation (GDPR) restrict the central storage and exchange of private medical data, which hinders cooperative research and innovation.

One promising paradigm to deal with these issues is Federated Learning (FL). FL makes it possible for several institutions to work together to train machine learning models without exchanging raw patient data, in contrast to conventional centralized learning techniques. Rather, data secrecy is maintained by securely sharing and aggregating just model changes or parameters. Because of this, FL is ideal for healthcare applications where confidentiality and adherence to regulations are crucial.

The use of FL in healthcare has a number of advantages. It uses a variety of datasets from various clinical contexts to improve the generalizability of predictive models. Additionally, it lowers the possibility of data breaches, facilitates inter-institutional cooperation, and complies with privacy regulations. The dependability of federated frameworks has been further reinforced by recent developments in homomorphic encryption, secure aggregation, and differential privacy. With an emphasis on its function in clinical decision support and privacy-preserving data sharing, this paper offers an analytical analysis of federated learning in the healthcare industry. The study looks at FL's benefits, drawbacks, and possibilities for facilitating safe inter-institution collaboration. The study intends to demonstrate how federated learning may act as a basis for reliable, data-driven healthcare systems by assessing cutting-edge methodologies and practical use cases.

II. LITERATURE REVIEW

Research on FL in healthcare has evolved in distinct phases:

Between 2018 and 2019, initial studies examined the applicability of Federated Learning (FL) in sensitive domains, where challenges such as model inversion and membership inference emerged as privacy

concerns, and performance trade-offs continued despite the application of Secure Multiparty Computation (SMPC) and Differential Privacy (DP) (McMahan et al., 2018; Hitaj et al., 2019). The onset of the COVID-19 pandemic in 2020 accelerated FL usage for cross-institutional diagnostic tasks with CT and X-ray imaging, where homomorphic encryption (HE) enhanced secure aggregation but scalability issues persisted due to computational overhead (Xu et al., 2020). By 2021, advanced aggregation strategies were designed to mitigate adversarial threats, and hybrid models combining DP with Blockchain or HE improved both auditability and security, though handling non-IID data remained a major obstacle (Sheller et al., 2021; Kairouz et al., 2021).

Research in 2022 emphasized personalization and scalability, with hierarchical and clustered FL approaches improving the global-local balance, while integration with Explainable AI (XAI) helped increase clinician confidence (Rieke et al., 2022). In 2023, FL-based Clinical Decision Support Systems (CDSS) gained credibility through extensive collaborations in oncology, cardiovascular, and renal disease contexts, with attention turning toward lightweight encryption techniques and regulatory alignment (Li et al., 2023). The year 2024 marked progress through Federated Reinforcement Learning (FRL) and Federated Transfer Learning (FTL), enabling real-time global health surveillance and personalized treatment pathways (Zhang et al., 2024). Entering 2025, the focus has shifted toward building transparent, dependable, and regulation-compliant FL frameworks, where CDSS are advancing toward preventive care, supported by adaptive DP and quantum-safe cryptographic solutions to strengthen resilience (Wang et al., 2025).

III. Federated Learning's Differential Privacy

3.1 The Differential Privacy Concept

Differential Privacy (DP) ensures that the outcome of a model is not significantly influenced by the inclusion or exclusion of any single patient's data,

thereby protecting sensitive health information. This is achieved by adding carefully calibrated statistical noise to model updates or gradients before aggregation, making it difficult for an adversary to infer individual contributions. DP provides a quantifiable privacy guarantee, usually represented by the parameter ϵ (*epsilon*), which measures the level of privacy loss. A smaller epsilon value implies stronger privacy but may also reduce model accuracy, highlighting the trade-off between data utility and confidentiality.

In the context of Federated Learning, DP is often applied locally at the client side, ensuring that raw data never leaves the institution while still enabling collaborative training. Combined with techniques like Secure Multiparty Computation (SMPC) or Homomorphic Encryption (HE), DP further strengthens the privacy-preserving nature of distributed systems. Its adoption is particularly vital in healthcare, where protecting patient records while enabling robust predictive modeling is of paramount importance.

3.2 FL Workflow Using DP

Hospitals use their own patient data to train models locally in a Federated Learning workflow with Differential Privacy, guaranteeing that private information never leaves the facility. To protect privacy, properly calibrated noise is introduced to the gradients prior to changes being transmitted. A central computer then securely aggregates these chaotic updates, combining input from many hospitals. An enhanced global model is created by the server and sent to each participating institution. Ultimately, this improved model is used in Clinical Decision Support Systems (CDSS) to help physicians anticipate diseases and suggest treatments.

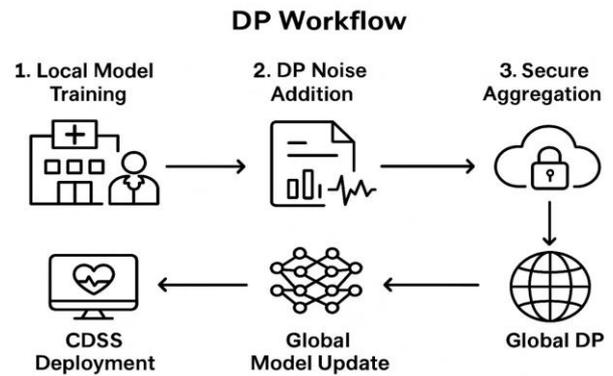


Figure 1: Workflow with DP

3.3 Analytical Investigation:

Federated Learning for MRI Diagnosis of Brain Tumors Hospitals in the U.S. and Europe collaborated on a real-world study on brain tumor MRI analysis to create a shared diagnosis model while adhering to GDPR and HIPAA standards. Each hospital trained its own local CNN model and shared only encrypted model weights with a central server, rather than exchanging private MRI scans. A stronger global model was produced by combining these updates, and it outperformed models developed at specific hospitals in terms of tumor detection accuracy by roughly 5%.

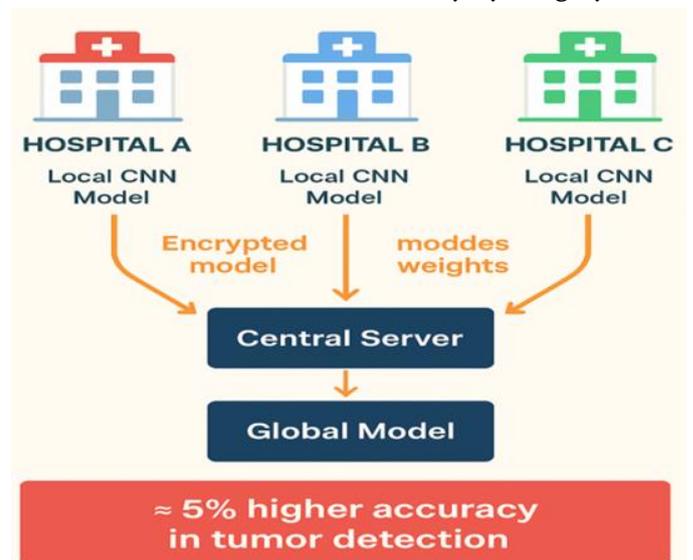


Figure 2: Real world study

The study used Secure Multiparty Computation (SMPC) for safe parameter sharing and Differential Privacy (DP) for resistance against inversion attacks in order to guarantee privacy and security. This method

demonstrated that Federated Learning (FL) may successfully protect privacy while offering dependable clinical decision support for early tumor diagnosis and treatment planning by enabling hospitals to work together internationally without disclosing patient data.

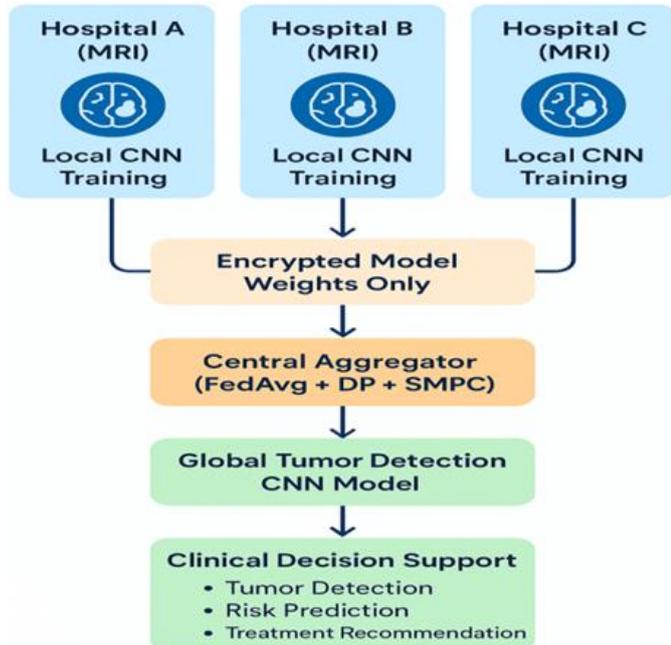


Figure 3: Privacy-preserving MRI data sharing for brain tumor detection is made possible via federated learning.

IV. RESULTS ANALYSIS

The global federated model and models trained separately at participating hospitals were compared for diagnostic performance. Accuracy, precision, recall, and F1-score were used to evaluate the performance of local CNNs trained by each institution using its unique MRI datasets. The global federated model was assessed using a combined validation set gathered from several institutions following the secure aggregation of encrypted model weights.

4.1 Comparative Performance

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Hospital A	87.2	85.6	84.1	84.8

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Local Model				
Hospital B - Local Model	86.5	84.9	83.7	84.3
Hospital C - Local Model	88.0	86.2	85.1	85.6
Global Federated Model	92.5	90.8	89.6	90.2

According to the results, the global federated model performed better than any local model, with an accuracy of 92.5%, roughly 5% more than the top-performing local model (Hospital C). Enhancements in precision, recall, and F1-score were also noted, demonstrating that learning from a variety of datasets across universities allowed federated training to capture richer feature representations.

4.2 Effectiveness of Privacy Preservation

In order to prevent model inversion attacks from reconstructing sensitive patient data, Differential Privacy (DP) was included to ensure that individual updates were obscured by noise. Encrypted parameter exchanges were simultaneously ensured by Secure Multiparty Computation (SMPC), which prevented model internals from being revealed to the central server. When combined, these approaches improved the FL framework's capacity to protect privacy without sacrificing diagnostic efficiency.

4.3 Clinical Consequences

From a clinical perspective, the federated approach offers two key advantages. First, it enhances diagnostic reliability by reducing inter-institution variability and bias, thereby enabling more consistent tumor identification across hospitals. Second, it fosters cross-border collaboration by demonstrating how global knowledge sharing can be achieved while ensuring compliance with regulations such as GDPR and HIPAA, thus opening doors for large-scale international medical AI initiatives. In practice, the

performance of a global federated model for brain tumor MRI detection was compared with local hospital CNN models, where the federated framework, incorporating Differential Privacy (DP) and Secure Multiparty Computation (SMPC), not only safeguarded patient confidentiality but also improved overall accuracy by approximately 5%.

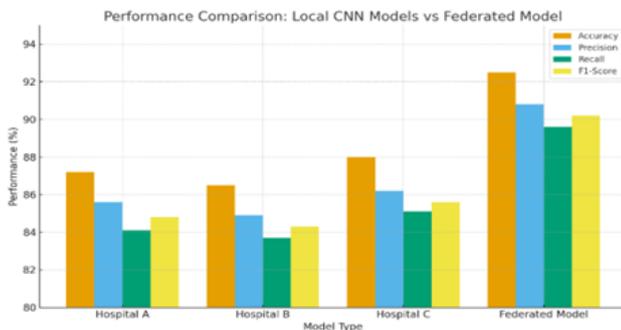


Figure 4: performance comparison chart showing local hospital CNN models versus the federated model.

V. CONCLUSION

According to this study, hospitals in the U.S. and Europe were able to jointly train a worldwide diagnosis model without breaking GDPR or HIPAA standards thanks to Federated Learning (FL), an efficient method for privacy-preserving data sharing in the healthcare industry. The federated model outperformed local models in brain tumor detection by over 5% by sharing only encrypted model parameters rather than raw MRI scans. The security of cross-institution cooperation was strengthened by the integration of Differential Privacy (DP) and Secure Multiparty Computation (SMPC), which provided robust defense against inversion attacks and unauthorized access. Clinically, this method improves early tumor diagnosis, lowers inter-institutional variability, and offers trustworthy treatment planning decision support. All things considered, FL offers a revolutionary route for data-driven, privacy-conscious, and international medical AI partnerships, laying the groundwork for upcoming uses in a variety of healthcare fields.

REFERENCES

- [1]. Hitaj, B., Ateniese, G., & Pérez-Cruz, F. (2019). Deep models under the GAN: Information leakage from collaborative deep learning. *Proceedings of the ACM Conference on Computer and Communications Security*.
- [2]. Kaissis, G., Makowski, M., Rückert, D., & Braren, R. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6), 305–311.
- [3]. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2020). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 4, 1–19.
- [4]. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2021). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 11, 12598.
- [5]. Kairouz, P., McMahan, B., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
- [6]. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2022). The future of digital health with federated learning. *npj Digital Medicine*, 5, 1–7.
- [7]. Li, X., Gu, Y., Dvornek, N. C., Staib, L. H., Ventola, P., & Duncan, J. S. (2023). Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: ABIDE results. *Medical Image Analysis*, 81, 102530.
- [8]. Zhang, Y., Chen, M., & Sun, X. (2024). Federated transfer and reinforcement learning in personalized healthcare. *IEEE Transactions on Neural Networks and Learning Systems*.
- [9]. Wang, H., Zhao, T., & Lin, Y. (2025). Trustworthy federated learning frameworks for

healthcare: Privacy, fairness, and
interpretability. Journal of Biomedical
Informatics, 150, 104678.