

Comparative Analysis of Federated Learning Architectures, Algorithms, and Privacy-Preserving Techniques for Healthcare Applications

Tejas H V, Sachin K, Kishore Kumar K, Monica K P
Department of CS&E, K.V.G College of Engineering, Sullia
Email: tejasvh23@gmail.com

Department of CS&E, K.V.G College of Engineering, Sullia
Email: sachinkudekallu@gmail.com

Department of CS&E(AI&ML), K.V.G College of Engineering, Sullia
Email: kishorkajjodi@gmail.com

Department of CS&E, K.V.G College of Engineering, Sullia
Email: monicakp86@gmail.com

Abstract:

The deployment of artificial intelligence in healthcare presents a critical tension between leveraging datasets for medical advancement and adhering to stringent privacy regulations such as HIPAA and GDPR. Federated Learning (FL) offers a distributed architecture that enables institutions to collaboratively train models without exchanging sensitive, raw patient data. This paper presents a comparative analysis of FL frameworks for privacy-preserving predictive analytics, evaluating core architectures—Horizontal FL, Vertical FL, and Federated Transfer Learning—alongside the performance of algorithms like FedAvg, FedSGD, and FedProx in navigating the challenges of non-independent and identically distributed (non-IID) clinical data. Furthermore, the study quantifies the effectiveness of privacy-preserving techniques, including Differential Privacy (DP), Homomorphic Encryption (HE), and Secure Multi-Party Computation (SMPC), against threat vectors such as model inversion, backdoor, and data poisoning attacks. Our findings indicate that while advanced algorithms like FedProx significantly enhance stability in heterogeneous environments, hybrid privacy mechanisms offer the strongest defense against adversarial threats. However, these integrated frameworks introduce inherent trade-offs between privacy guarantees, computational overhead, and model utility. Ultimately, no single FL algorithm or privacy technique is universally optimal; successful clinical deployment requires context-aware frameworks that carefully balance architectural design, data modality, and regulatory compliance to achieve reliable predictive analytics.

Keywords - Federated Learning (FL), Privacy-Preserving Machine Learning, Healthcare Artificial Intelligence, Predictive Models, Data Heterogeneity (Non-IID), Differential Privacy (DP), Homomorphic Encryption (HE), Regulatory Compliance (HIPAA, GDPR).

I. INTRODUCTION

The rapid evolution of artificial intelligence (AI) and machine learning (ML) in healthcare has opened unprecedented avenues for advanced clinical diagnostics, personalized treatment planning, and robust predictive analytics. Modern medical AI relies heavily on access to massive, high-quality, and diverse datasets to train deep learning models effectively. However, centralizing medical data faces severe roadblocks due to the sensitive nature of patient records and the strict enforcement of regulatory compliance frameworks such as the Health Insurance

Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in the European Union. These regulations prioritize patient data confidentiality, data minimization, and restriction of cross-border or inter-institutional data transfers, creating a critical tension between data utilization for medical AI advancement and institutional privacy mandates.

To bridge this chasm, Federated Learning (FL) has emerged as a transformative decentralized machine learning paradigm. Unlike traditional centralized learning where local data must be transferred to a single repository, FL allows multiple healthcare institutions to

collaboratively train a global model while keeping sensitive patient records entirely localized at their source. By restricting data movement to only local model updates (such as gradients or weights), FL naturally minimizes the logistical, legal, and security risks inherent to centralized data sharing.

II. CORE FEDERATED LEARNING ARCHITECTURES

Federated Learning (FL) is a distributed machine learning framework that enables multiple parties to collaboratively train a model without directly sharing their raw, localized data. In the healthcare domain, this decentralized paradigm addresses the critical tension between leveraging large-scale datasets for medical AI advancement and adhering to stringent privacy regulations like HIPAA and GDPR. By keeping sensitive patient data decentralized at its source, FL provides a foundational architecture for privacy-preserving collaboration across hospitals, clinics, and research institutions.

A. Horizontal Federated Learning (HFL)

This is the most common architecture, applicable when different institutions hold data with the same feature space but different patient samples. This setup is typical in multi-hospital collaborations where each site holds similar types of patient records, clinical charts, or medical images. A systematic review highlights its dominance, showing that HFL was utilized in 511 out of 577 studies.

B Vertical Federated Learning (VFL)

VFL is deployed when different entities hold different types of information (features) about the same or an overlapping set of patients. For example, a regional hospital might possess clinical history records, while a separate laboratory holds genomic sequencing data for the same patient cohort. This architecture, used in 23 studies, allows for building richer predictive models by combining diverse data types without centralizing them.

C. Federated Transfer Learning (FTL)

FTL applies when datasets differ significantly in both their sample and feature spaces. This is common in multi-institution research where knowledge from a data-rich source domain (e.g., a large, well-funded research hospital) can be transferred to improve models in a data-scarce target domain (e.g., a smaller rural clinic), even with completely different data distributions. This approach was leveraged in 24 studies.

III. PRIVACY THREATS AND PRESERVATION TECHNIQUES

While FL inherently keeps raw data local, the continuous exchange of intermediate model updates (such as weights or gradients) can still leak sensitive information, necessitating additional privacy-preserving techniques to secure the framework.

A. Differential Privacy (DP)

DP adds carefully calibrated noise (e.g., Gaussian or Laplace) to model updates or outputs to prevent the inference of individual data points, providing a rigorous mathematical privacy guarantee. A prominent DP technique is the Private Aggregation of Teacher Ensembles (PATE), which uses an ensemble of "teacher" models trained on disjoint data to generate noisy labels for a "student" model.

B. Homomorphic Encryption (HE)

A HE allows computations, such as model aggregation, to be performed directly on encrypted data without needing to decrypt it first. The CKKS scheme is a prominent HE method used in healthcare FL to encrypt local model weights before they are transmitted to the server.

C. Secure Multi-Party Computation (SMPC)

SMPC enables multiple parties to jointly compute a function over their private inputs without revealing those inputs to each other. Within an FL system, SMPC is primarily used to secure the model aggregation process itself on the central server.

IV. COMPARATIVE ANALYSIS OF FL ALGORITHMS

The performance of FL algorithms varies significantly based on data distribution and system constraints. A key challenge in healthcare is the prevalence of non-independent and identically distributed (non-IID) data across institutions, which can severely degrade model performance. For instance, studies using medical imaging datasets have shown significant accuracy drops in non-IID settings compared to IID settings when using common algorithms like Federated Averaging (FedAvg)

TABLE I
COMPARATIVE ANALYSIS ON ALGORITHMS

Algorithm	Key Strength	Key Limitation
FedAvg	Simple, communication-efficient via aggregated model updates.	Performance degrades significantly under extreme non-IID conditions.
FedProx	Enhanced stability and convergence in heterogeneous environments due to a proximal regularizer.	Requires careful hyperparameter tuning
FedSGD	Provides frequent gradient updates.	Low communication efficiency and poor handling of non-IID data.
Hierarchical FL	Efficient for large healthcare systems with	Introduces complexity in

nested structures (e.g., departments, hospitals).	synchronization and trust management.
---	---------------------------------------

V. SUITABILITY FOR DIFFERENT HEALTHCARE DATA TYPES

Healthcare applications involve diverse data modalities, each presenting unique challenges for FL.

A. Medical Imaging Data

Applications in radiology, pathology, and ophthalmology involve high-dimensional data. FedProx demonstrates superior robustness in tasks like retinal optical coherence tomography (OCT) image classification under heterogeneous conditions. Hierarchical FL approaches are also well-suited to the structured networks of hospitals and imaging centers.

B. Electronic Health Records (EHR)

Structured EHR data is used for prediction tasks like hospital length of stay or disease risk stratification. The choice of algorithm here may depend on the primary goal: statistical FL methods are preferable for understanding associations, while FedProx is effective for building predictive models from non-IID EHR data across institutions.

C. Genomic and Multi-Modal Data

Integrating sensitive data types like genomics with imaging or EHRs presents unprecedented challenges for FL. Approaches that combine strong cryptographic privacy guarantees, such as Homomorphic Encryption (HE), with FL algorithms may be necessary, though they incur higher computational and communication overhead.

Fig. 4 Dendrogram for hierarchical clustering of patient features (10,000 patients)

VI. SYSTEM ARCHITECTURE

While the core distributed architecture of Federated Learning (FL) inherently prevents the direct exposure of raw clinical data, transmitting intermediate model parameters or gradients still presents clear information leakage vectors. Deployed models can be vulnerable to sophisticated adversarial strategies designed to extract sensitive patient information or manipulate model behavior. Securing healthcare FL networks against these threats requires augmenting the base optimization framework with additional mathematical and cryptographic privacy-preserving techniques.

The defence capabilities of standalone and hybrid privacy configurations have been rigorously analysed against severe clinical attack vectors using key security metrics.

A. Model Inversion Attacks

These attacks aim to reverse-engineer and reconstruct original training data from exposed model gradients. Defense efficacy is measured using the Mean Squared Error (MSE) between the original clinical images or data and the adversary's reconstructed output, where a higher MSE signifies stronger privacy protection. A hybrid framework combining FL with PATE and SMPC (FL_PATE_SMPC) achieves exceptional resistance, yielding an MSE of \$19.267\$ compared to the highly vulnerable base model's MSE of \$0.9676\$.

B Backdoor Attacks

This threat vector involves a malicious client embedding hidden malicious triggers inside the local updates to hijack the global model's behavior. Success is quantified via the Attack Success Rate (ASR). An integrated model combining all three core techniques (FL_PATE_CKKS_SMPC) establishes the most robust defense, restricting the ASR to just \$0.0920\$, whereas a vanilla FL baseline suffers a high ASR of \$0.6800\$.

C. Federated Transfer Learning (FTL)

These attacks focus on corrupting training inputs to degrade performance. For *untargeted poisoning*, the combination of homomorphic encryption and secure multi-party computation (FL_CKKS_SMPC) delivers optimal resilience with a minimal ASR of \$0.0010\$. For *targeted poisoning*, both standalone encryption (FL_CKKS) and the hybrid FLCKKS_SMPC scheme serve as top performers, matching a low ASR of \$0.0020\$.

D. Man-in-the-Middle (MITM) Attacks

These attacks involve intercepting and tampering with model updates during transmission over network channels. When subjected to MITM tampering, a base FL model undergoes a massive \$48.45\%\$ drop in predictive accuracy. Conversely, the comprehensive \$FL_PATE_CKKS_SMPC\$ framework maintains strict operational stability, experiencing a negligible accuracy degradation of only \$1.68\%\$.

TABLE II
KEY ATTACK TYPES AND THEIR MITIGATION

Attack Type	Most Effective Model Configuration	Key Performance Metric	Improvement Over Base FL Baseline
Model Inversion	FL_PATE_SMPC	MSE: 19.267	~20times X
Backdoor	FL_PATE_CKKS_SMPC	ASR: 0.0920	~7x
Untargeted Poisoning	FL_CKKS_SMPC	ASR: 0.0010	~400x
Targeted Poisoning	FL_CKKS	ASR: 0.0020	~490x

g			
MITM	FL_PATE_CKKS_S MPC	Accuracy Degradation: 1.68%	~29x

VII. FUTURE HORIZONS: OPEN RESEARCH DIRECTIONS

At its core, Federated Learning (FL) works like a decentralized team project where everyone collaborates without sharing their private notes. Instead of gathering sensitive patient data from various hospitals into one central database—which raises massive security and privacy concerns—the master AI model travels directly to each individual institution. Each hospital uses its own local computational resources to train this visiting model on its unique set of patient records, clinical charts, or medical images. Once this local training is complete, the hospital sends only the isolated model improvements, such as adjusted weights and gradients, back to a central server.

To fully transition healthcare federated learning (FL) frameworks from theoretical models into ubiquitous, production-grade clinical tools, several critical research paths must be pursued. First, the development of Explainable Federated Learning (XFL) is paramount. Future work must design techniques capable of computing interpretable, post-hoc clinical explanations at individual local sites, which can then be safely aggregated into a coherent global explanation without breaching privacy or degrading under differential privacy noise. Second, automating privacy budget optimization represents a major step forward. Rather than relying on rigid, manually configured parameters that lead to suboptimal utility trade-offs, adaptive algorithms should dynamically throttle noise injection across iterations—allocating less noise during early training phases when gradients are large, and utilizing a higher privacy cost during late-stage clinical fine-tuning

Third, the medical AI field requires standardized, multi-dimensional evaluation platforms. Future benchmarking frameworks must simultaneously evaluate downstream predictive performance across diverse non-IID clinical populations, measure mathematical resilience against reconstruction attacks, and monitor computational efficiency in bandwidth-constrained edge environments. Finally, exploring technical synergies with complementary emerging technologies will accelerate real-world adoption. Integrating blockchain architecture can provide immutable audit trails for model updates and automate cross-institutional compliance via smart contracts. When paired with advanced communication-efficient mechanisms—such as compressed error feedback, partial-layer aggregation, and Federated Meta-Learning—these hybrid

configurations will dramatically lower the barriers to scaling decentralized predictive analytics globally.

VIII. CONCLUSION

In conclusion, the deployment of federated learning in healthcare offers a transformative, decentralized paradigm that successfully balances the advancement of collaborative medical AI with the stringent privacy mandates of HIPAA and GDPR. This comparative analysis demonstrates that while advanced algorithms like FedProx effectively mitigate the performance degradation caused by non-IID clinical data, the base architecture must be augmented with hybrid privacy-preserving techniques—such as Differential Privacy, Homomorphic Encryption, and Secure Multi-Party Computation—to defend against sophisticated adversarial threats like model inversion and data poisoning. However, implementing these robust cryptographic defenses introduces a critical "privacy-performance-practicality" triangle, demanding careful calibration between mathematical privacy guarantees, computational execution times, and downstream model accuracy. Ultimately, because no single algorithm or privacy mechanism is universally optimal, the future of clinical federated learning lies in the maturation of context-aware, integrated, and hybrid frameworks that optimize specific data modalities and resource constraints to realize secure, auditable, and scalable predictive analytics. Moving forward, the successful translation of these frameworks from theory to production will rely on fostering institutional trust and establishing standardized governance models alongside technical advancements. As healthcare networks expand, defining transparent rules for data contribution, collaborative ownership, and multi-tier network accountability will be just as critical as the cryptographic protocols securing them. By unifying robust data regularization algorithms like FedProx, verifiable blockchain audit trails, and multi-layered hybrid privacy defenses, the medical community can confidently move past isolated data silos. Ultimately, these integrated federated learning systems will lay the groundwork for a secure, globally interconnected digital health ecosystem, driving clinical innovation and improving patient outcomes while keeping patient data entirely confidential and protected at its source.

REFERENCES

[1] Zhen Ling Teo, Liyuan Jin, Nan Liu, Siqi Li, Di Miao, Xiaoman Zhang, Wei Yan Ng, Ting Fang Tan, Deborah Meixuan Lee, Kai Jie Chua, John Heng, Yong Liu, Rick Siow Mong Goh, Daniel Shu Wei Ting, Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture, Cell Reports Medicine, Volume 5, Issue 2, 2024,101419,

ISSN 2666-3791,

<https://doi.org/10.1016/j.xcrm.2024.101419>.

(<https://www.sciencedirect.com/science/article/pii/S2666379124000429>)

- [2] Shalabi, E.; Khedr, W.; Rushdy, E.; Salah, A. A Comparative Study of Privacy-Preserving Techniques in Federated Learning: A Performance and Security Analysis. *Information* **2025**, *16*, 244. <https://doi.org/10.3390/info16030244>.
- [3] Babar M, Qureshi B, Koubaa A (2024) Investigating the impact of data heterogeneity on the performance of federated learning algorithm using medical imaging. *PLoS ONE* **19**(5): e0302539. <https://doi.org/10.1371/journal.pone.0302539>.
- [4] Dai, J. (2024). Comparative analysis of federated learning algorithms under non-IID data. *Applied and Computational Engineering*, 86, 91-100.
- [5] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *J. Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [6] Li S, Miao D, Wu Q, Hong C, D'Agostino D, Li X, Ning Y, Shang Y, Wang Z, Liu M, Fu H, Ong MEH, Haddadi H, Liu N. Federated Learning in Healthcare: A Benchmark Comparison of Engineering and Statistical Approaches for Structured Data Analysis. *Health Data Sci.* 2024 Dec 4;4:0196. doi: 10.34133/hds.0196. PMID: 39635226; PMCID: PMC11615161.
- [7] Singh, J.P.; Aqsa, A.; Ghani, I.; Sonani, R.; Govindarajan, V. Privacy-Aware Hierarchical Federated Learning in Healthcare: Integrating Differential Privacy and Secure Multi-Party Computation. *Future Internet* **2025**, *17*, 345. <https://doi.org/10.3390/fi17080345>
- [8] Bamidele Samuel Adelus, Damilola Osamika, MariaTheresa Chinyeaka Kelvin-Agwu, Ashiata Yetunde Mustapha, Adelaide Yeboah Forkuo, & Nura Ikhalea. (2025). A Federated Interoperability Framework for Seamless Health Data Exchange Using FHIR Standards Across Multi-Hospital Systems. *Engineering And Technology Journal*, *10*(5), 4672–4695. <https://doi.org/10.47191/etj/v10i05.03>.
- [9] N. Ramesh, M. Anwar and K. K, "Privacy-Preserving Federated Learning in Healthcare: Integrating Homomorphic Encryption, Smart Contracts, and Adaptive Strategies," *2025 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI)*, Chennai, India, 2025, pp. 1-6, doi: 10.1109/ICDSAAI65575.2025.11011810.