MULTIDISCIPLINARY RESEARCH IN COMPUTING INFORMATION SYSTEMS

VOL 05 ISSUE 06 2025

P-ISSN: 3080-7182
E-ISSN: 3080-7190

https://mrcis.org

FEDERATED LEARNING FOR PRIVACY-PRESERVING HEALTHCARE DATA ANALYTICS

Farah Javed [1], Imran Hussain [2]

Abstract. *The rising ubiquity of electronic health records, wearable sensors and hospital imaging systems presents vast opportunities for data-driven healthcare analytics. However, concerns over patient privacy, regulatory compliance and data siloing hinder the sharing of raw medical data across institutions. Federated Learning (FL) offers a collaborative machine learning paradigm in which local models are trained at individual healthcare sites and only model updates (not raw patient data) are exchanged. This article examines the architecture of FL in healthcare settings, explores privacy-preserving mechanisms (secure aggregation, differential privacy, homomorphic encryption), analyses domain-specific challenges such as non-IID medical data and communication overhead, and discusses empirical trends and deployment pathways. Two conceptual charts visualise the trade-offs between number of sites and model accuracy, and between privacy strength and model error. We conclude by outlining key recommendations for translational deployment in multi-site healthcare analytics.*

Keywords: *Federated Learning, Privacy-Preserving Analytics, Healthcare Data Collaboration, Secure Aggregation.*

INTRODUCTION

In modern healthcare ecosystems, patient data is distributed across hospitals, clinics, and personal devices, creating both analytic potential and privacy concerns. Traditional machine-learning approaches often require pooling raw data centrally, which is restricted by regulators and institutions. Federated Learning (FL) enables multiple healthcare sites to collaboratively train a global model while each retains local control of sensitive patient records. This paradigm aligns strongly with the imperative for data minimisation and privacy by design in medical analytics. In healthcare applications, FL enables cross-institutional predictive modelling (e.g., diagnosis, outcome prediction) without direct data sharing. Yet the deployment of FL in healthcare

[1] *Department of Computer Science, University of Karachi, Pakistan.*

[2] *School of Electrical Engineering and Computer Science, National University of Sciences & Technology (NUST), Islamabad, Pakistan.*

encounters challenges: data heterogeneity, non-IID distributions, resource limitations at sites, communication latency, and risks of information leakage through model updates. This article reviews the foundational architecture of FL for healthcare, explores the privacy-enhancing techniques required, discusses domain-specific obstacles and outlines practical steps toward implementation.

1. Architecture and Workflow of Federated Learning in Healthcare

Federated Learning (FL) in healthcare enables multiple institutions (e.g., hospitals, research centres, or clinics) to collaboratively train machine learning models without sharing raw patient data. The FL architecture and workflow depend on the type of data partitioning model (horizontal, vertical, or hybrid) and the system's architecture (client-server vs. peer-to-peer).

Horizontal FL (Data-Parallel FL)

In horizontal FL, data from multiple institutions are similar in structure (i.e., they contain the same features, but different samples of patient data). For example, different hospitals may hold data on various patients, but the variables for each patient (e.g., age, gender, diagnosis) are the same. In this model, each institution trains its local model using its data, and only the model weights (not the data) are exchanged. Horizontal FL is ideal for federated systems where the data has the same feature space but is distributed across institutions.

Vertical FL (Feature-Parallel FL)

In vertical FL, data from different institutions may share the same set of patients, but the features (variables) differ. For instance, a hospital might have clinical data about patients (e.g., EHR), while a diagnostic lab may have additional features like lab results or genetic data for the same set of patients. In this case, the institutions collaborate by splitting the model training across the features each possesses, exchanging only the model parameters. Vertical FL is effective when different data modalities (e.g., imaging, clinical records) are collected across different sites but with shared patients.

Hybrid FL (Combination of Horizontal and Vertical)

Hybrid FL is a combination of both horizontal and vertical approaches. It is used when institutions share both common features and patient sets but may also possess unique features. For example, a clinic may have patient data on specific diseases that other clinics do not have but can still collaborate on shared features like demographic information.

Client-Server vs. Peer-to-Peer Architectures

FL can adopt either a client-server or peer-to-peer architecture, depending on the healthcare setup:

Client-Server Architecture: In a client-server model, the central server coordinates the training process. Each client (hospital or healthcare institution) trains a local model and sends model updates to the central server. The server aggregates these updates using an aggregation protocol (e.g., Federated Averaging - FedAvg) and updates the global model. This architecture is beneficial for healthcare systems that rely on centralised coordination.

Peer-to-Peer Architecture: Peer-to-peer FL allows direct communication and model aggregation between participating institutions. In this setup, there is no central server; instead, each participant collaborates with others to improve the model. This decentralized approach can be beneficial in distributed healthcare networks with several autonomous units.

Aggregation Protocols

The most widely used aggregation protocol is Federated Averaging (FedAvg), where each site computes model updates and shares them with the central server. The server then averages the updates (weights or gradients) to form a global model. While FedAvg is simple and efficient, it may not work well with heterogeneous or non-IID (independent and identically distributed) data across healthcare institutions. Alternative methods such as secure aggregation or differential privacy can be employed to further protect patient privacy during aggregation.

Site Onboarding and Participation

When onboarding sites for FL, healthcare institutions must ensure that data-sharing agreements and regulatory compliance (e.g., HIPAA or GDPR) are met. Furthermore, each site must have a compatible infrastructure to support FL training (e.g., computational resources, secure data storage). Participation can vary across sites: some may be more involved with data sharing and model updates, while others may only be limited to occasional participation based on their available resources or privacy concerns.

2. Privacy Enhancement & Security Mechanisms

Privacy preservation and data security are critical in Federated Learning, particularly when dealing with sensitive healthcare data. The following privacy and security mechanisms help ensure that data privacy is maintained while enabling model training in a collaborative environment.

Secure Aggregation

Secure aggregation ensures that the central server cannot access the individual model updates sent by each institution. Instead of directly summing the model updates, a cryptographic protocol is used to aggregate the data securely. This ensures that even if the server is compromised, it cannot access the data updates from individual institutions. One widely used approach is Homomorphic Encryption (HE), which allows computations to be performed on encrypted data without decrypting it. After computation, the data is only decrypted at the intended endpoint, ensuring that privacy is preserved throughout the process.

Differential Privacy (DP)

Differential privacy is a technique used to ensure that individual data points are not exposed when the model is trained. In healthcare, this involves adding noise to the data before training, ensuring that the model learns generalized patterns and not specific information about any one patient. Differential privacy can be incorporated into FL by adding noise to the model updates before they are sent back to the server for aggregation. This ensures that the model does not reveal any private information about individual patients during the training process.

Homomorphic Encryption (HE)

Homomorphic encryption is another key privacy-enhancing technology that allows computations to be done on encrypted data without revealing any private information. When used in federated learning, HE ensures that sensitive patient data remains encrypted throughout the training process, and only encrypted model updates are exchanged. The central server can aggregate the encrypted updates to form a global model, which is then decrypted only after it is aggregated, ensuring that sensitive patient data is never exposed.

Trusted Execution Environments (TEEs)

Trusted Execution Environments are hardware-based security measures used to protect data during the model training process. TEEs create a secure, isolated environment within a processor to execute computations in a manner that prevents unauthorized access. In the context of FL, TEEs can be used to securely train the models on a local site without exposing the underlying data, even in the case of malicious attacks.

Adversarial Threats: Inference Attacks and Poisoning

Federated Learning in healthcare is also susceptible to adversarial attacks, such as inference attacks, where attackers try to extract private information from the model by observing updates. Model poisoning attacks can also occur, where an attacker manipulates local model updates to degrade the performance of the global model. Techniques like robust aggregation methods, anomaly detection, and encryption-based methods can be used to mitigate such threats. Researchers are actively working on methods to make FL more resilient to adversarial attacks in healthcare settings.

3. Healthcare-Specific Challenges and Data Considerations

Federated Learning (FL) in healthcare presents several unique challenges due to the heterogeneity of medical data, differences across sites, and regulatory considerations. These challenges must be addressed to ensure effective model training and deployment in a privacy-preserving manner.

Medical Data Heterogeneity (EHRs, Imaging, Genomics)

One of the most significant challenges in healthcare FL is the heterogeneity of data across different institutions. Healthcare data varies widely in format, structure, and content. For example, Electronic Health Records (EHRs) typically contain patient demographics, medical history, lab results, and diagnosis codes. Imaging data (such as X-rays, CT scans, MRIs) are typically unstructured and require specialized algorithms for analysis. Genomic data involves complex sequences of DNA information and is often high-dimensional and sparse. These differences in data types present a challenge for federated learning models, which may require pre-processing to standardize and handle this diverse data.

These data types may not always align, leading to non-IID (non-identical and independently distributed) data, meaning the distribution of data across institutions is not uniform. This can make training a federated model difficult as local models trained on different data distributions may not generalize well when aggregated, potentially leading to biases or reduced model accuracy.

Label Imbalance

Label imbalance is another prevalent issue in healthcare data, especially when dealing with rare diseases or conditions. For example, cancer datasets may have a significantly higher number of healthy patients compared to those diagnosed with cancer, creating a skewed class distribution. When training models across different sites with unbalanced data, federated models might prioritize majority classes and overlook minority class data, leading to poor performance in predicting rare conditions. Strategies like class-weight adjustment, oversampling, and data augmentation are typically used to mitigate label imbalance in federated learning systems.

Non-IID Data Across Sites

Non-IID data across sites presents one of the biggest challenges in federated learning. In healthcare, data collected at different hospitals or clinics may have significant variations in terms of patient populations, data collection methods, and feature distributions. For example, different hospitals may use varying diagnostic criteria, resulting in variations in the clinical data. Additionally, when aggregating models from heterogeneous datasets, the variance in data between sites could cause the global model to perform poorly on data from new, unseen sites. Solutions such as personalized federated learning or domain adaptation techniques are being explored to address this challenge, ensuring that models generalize well across diverse healthcare environments.

Site Dropout

In federated learning, site dropout is a challenge that can occur when participating healthcare institutions or sites drop out of the training process due to technical issues, lack of resources, or regulatory concerns. This can impact the model's training process, especially if critical sites are missing or if the dropout occurs at various points during training. Model robustness can be compromised if dropout is frequent. Techniques such as dynamic federated learning, where models are trained to adjust to the continuous addition or dropout of sites, are being explored to make federated models more resilient to site dropout.

Communication and Compute Constraints

Federated learning, by its nature, involves frequent communication between participating sites and the central server (or peers in peer-to-peer architecture). Healthcare institutions, especially those in remote or low-resource settings, may face significant challenges due to communication bandwidth limitations and compute constraints (e.g., limited computational power or storage). The communication overhead of constantly transferring model updates can be high, especially in multi-site settings with large datasets. Optimizing the number of communication rounds and reducing the size of updates via techniques like model compression, sparsification, and adaptive learning rates can help mitigate these constraints.

Regulatory Frameworks (HIPAA, GDPR)

Healthcare federated learning systems must comply with stringent regulations, such as HIPAA (Health Insurance Portability and Accountability Act) in the United States and GDPR (General Data Protection Regulation) in the European Union. These regulations mandate the protection of personal health information and ensure that patient data is not exposed or misused. In the context of federated learning, these regulations are enforced by ensuring that raw data is never shared, only aggregated model updates are communicated, and privacy-preserving techniques like differential privacy are implemented. Compliance with these regulations is crucial, and healthcare FL models must be designed to meet these legal requirements to avoid penalties and ensure ethical data use.

4. Empirical Trade-Offs and Performance Considerations

Federated learning in healthcare presents several trade-offs that must be considered when evaluating its feasibility and performance. These include scaling model accuracy with the number of participating sites, the cost of implementing privacy mechanisms, communication rounds, and comparing FL with centralised learning approaches.

Model Accuracy vs. Number of Participating Sites

As the number of sites participating in a federated learning system increases, the global model's accuracy may improve due to exposure to a more diverse set of data. However, there are diminishing returns as the number of sites increases, especially when sites have heterogeneous or

non-IID data. Furthermore, as more sites participate, the model convergence time increases due to the need for more communication rounds and aggregating model updates. Studies have shown that model accuracy improves up to a certain point, after which the gains plateau or even degrade due to increased data diversity. Therefore, it is crucial to determine an optimal number of sites that balances model accuracy with computational and communication costs.

Cost of Privacy Mechanisms

Privacy-preserving techniques such as secure aggregation, differential privacy, and homomorphic encryption come with a performance cost. While these mechanisms protect patient data, they can add computational overhead, increasing the time required for training and reducing the model's overall accuracy. For instance, differential privacy may add noise to model updates to protect individual data, but this can degrade the quality of the learned model. Similarly, homomorphic encryption requires encrypting and decrypting model updates, which can significantly slow down the federated learning process. The trade-off between privacy protection and model performance must be carefully considered based on the healthcare setting's privacy requirements and resources.

Communication Rounds and Latency

The number of communication rounds in federated learning directly affects the model's convergence time. In healthcare applications, communication overhead can be particularly costly, as large healthcare institutions may have limited bandwidth. The frequency of model updates and the number of communication rounds required to achieve an optimal model need to be optimized. Edge computing or local model updates can help alleviate this issue by reducing the need for frequent communication between the central server and participating sites. A balance must be struck between reducing communication costs and ensuring the model learns effectively from the distributed data.

Benchmarking Federated Learning vs. Centralized Learning

Federated learning must be benchmarked against traditional centralized learning approaches to assess its effectiveness in healthcare settings. Centralized learning aggregates all data into a central repository, which often leads to better model performance because of the availability of large, homogeneous datasets. However, this approach risks exposing sensitive patient data, which violates privacy regulations. In contrast, federated learning preserves data privacy, but its performance may be slightly lower due to the challenges discussed (non-IID data, privacy mechanisms, etc.). Comparing the two approaches helps determine when federated learning is the optimal solution and when centralized models might perform better.

Cost-Benefit of FL in Healthcare

The cost-benefit analysis of federated learning in healthcare includes factors like data privacy, infrastructure costs, model performance, and the trade-offs between maintaining data sovereignty and sharing model updates. The initial setup of federated learning (e.g., infrastructure for secure communication, storage for model updates) can be expensive. However, in multi-site healthcare environments, the benefits of collaborating without sharing raw data and improving models for disease detection or treatment predictions justify the investment. The cost-benefit ratio can be optimized through efficient algorithms, data compression techniques, and cost-effective infrastructure management.

5. Deployment Roadmap and Future Research Directions

Deploying Federated Learning (FL) in real-world multi-site healthcare settings presents significant challenges that require a structured roadmap and ongoing research to overcome barriers to adoption. The following recommendations provide a framework for the successful deployment of FL in healthcare:

Creation of Federated Healthcare Datasets

One of the primary challenges in implementing FL in healthcare is the lack of open, diverse, and standardized datasets that can be used for training federated models. Most healthcare data is siloed within individual institutions, making it difficult to create collaborative training sets. To address this, initiatives should focus on the creation of federated healthcare datasets that:

- Represent diverse patient populations, including different demographics, medical histories, and conditions.
- Include data from various healthcare modalities (EHRs, imaging, genomics, wearable data) to create comprehensive models.
- Ensure that datasets are non-IID (non-identically distributed) across sites to simulate real-world conditions and improve model generalization.
- Be shared in a way that complies with privacy regulations such as HIPAA and GDPR, using data anonymization and privacy-preserving techniques.

Creating federated datasets may require partnerships between hospitals, universities, and research organizations to build standardised data-sharing protocols and formats that can be used across institutions while maintaining patient confidentiality.

Standardized Metrics and Auditing

To ensure the effectiveness and reliability of FL models in healthcare, standardized metrics for evaluating performance must be established. These metrics should cover various aspects of healthcare-specific FL models, such as:

- Accuracy, recall, precision, and other classification metrics specific to disease prediction or medical imaging tasks.
- Data privacy metrics, such as how much patient information can be inferred from the model updates (quantifying privacy leakage).
- Model robustness, including how well the model performs under various attack scenarios, like poisoning or inference attacks.

Alongside standardized metrics, robust auditing frameworks must be developed to verify that federated learning systems are compliant with regulatory requirements and ethical guidelines. Audits should focus on verifying the security of the data aggregation process, ensuring that privacy-preserving techniques are correctly implemented, and monitoring the consistency and fairness of model updates across institutions. Regular audits will ensure that federated learning models can be deployed at scale without compromising patient privacy or regulatory compliance.

Governance Frameworks

A strong governance framework is essential for the smooth operation and long-term sustainability of federated learning in healthcare. These frameworks should define the roles and responsibilities of each participating site, the central server (if any), and data custodians. Key governance components should include:

Data ownership: Establish clear guidelines on who owns the patient data and how the data is shared between institutions for model training.

Ethical guidelines: Ensure that all participating sites adhere to ethical standards for handling sensitive healthcare data and that patients are informed and give consent for their data to be used in federated learning models.

Conflict resolution: Develop mechanisms for resolving disputes between participating sites or for addressing issues such as site dropout, data discrepancies, or model degradation.

Transparency and accountability: Ensure transparency in decision-making processes and regular reporting to stakeholders on the performance and privacy of the FL systems.

A well-structured governance framework will help mitigate the risk of conflicts, ensure regulatory compliance, and maintain trust between healthcare institutions and patients.

Cross-Border Federations

In a globalized healthcare system, multi-country collaboration may be necessary to create more diverse and robust federated learning models. Cross-border federations present both opportunities and challenges. These federations can:

Broaden the diversity of patient data, especially for rare diseases or conditions that may not be prevalent in all regions.

Facilitate collaboration between international healthcare institutions, creating larger and more representative datasets for model training.

Address the global health disparities by pooling data from different countries, leading to more accurate and inclusive models.

Cross-border federations face significant regulatory and legal challenges, such as differences in data protection laws and cross-border data transfer regulations (e.g., GDPR in Europe). Overcoming these challenges requires international cooperation and the establishment of global standards for federated learning systems in healthcare. Regulatory frameworks must be harmonized, and data sovereignty must be respected to ensure compliance across jurisdictions.

Further Research Gaps

While federated learning holds great promise for healthcare data analytics, several research gaps remain that must be addressed to ensure its successful adoption and deployment:

Non-IID Data Handling: Federated learning systems often struggle with non-IID data distributions. Research into new algorithms that can handle such data distributions effectively will be crucial. Methods such as personalized federated learning and domain adaptation could help address this gap.

Privacy-Preserving Algorithms: There is still significant research needed to improve privacy-preserving mechanisms such as differential privacy, secure aggregation, and homomorphic encryption. Enhancing the efficiency of these techniques while maintaining strong privacy guarantees will be essential for healthcare applications.

Federated Learning Optimization: Federated learning systems often face issues such as slow convergence and high communication overhead. Research into more efficient optimization algorithms, data compression techniques, and adaptive learning approaches is necessary to improve the speed and performance of FL models in healthcare.

Model Interpretability: One of the challenges of deploying FL in healthcare is ensuring that the models are interpretable. This is critical for healthcare professionals who need to trust the models' decisions. Research into explainable AI (XAI) for federated learning will be key to addressing this challenge.

Real-World Validation and Deployment: While much of the research on FL in healthcare is theoretical, there is a need for more real-world validation and large-scale deployments. Collaborations with healthcare providers to conduct empirical studies will help demonstrate the effectiveness and practical challenges of federated learning in healthcare.

Ahmad (2025) provides a detailed analysis of eight major Pakistani State-Owned Enterprises (SOEs), including PIA, Pakistan Steel Mills, and Pakistan Railways, over the period 2019–2024. Using both quantitative and qualitative methods such as thematic content analysis, cross-case comparison, and theoretical mapping, the study reveals chronic losses, heavy subsidy dependence, and low operational efficiency. PIA and Pakistan Steel Mills alone consume over 92% of total subsidies, reflecting structural inefficiencies and political interference. Ahmad emphasizes that urgent reforms, including privatization, public-private partnerships, and professionalized governance, are essential to restore public trust, enhance accountability, and promote sustainable public sector management in Pakistan.

Ahmad (2025) examines human–AI collaboration in professional knowledge work, focusing on productivity, error patterns, and ethical risks. Participants were assigned to human-only, AI-assisted, and optional AI-only task groups performing writing, summarization, decision-support, and problem-solving activities. Findings indicate that AI assistance accelerates task completion by 32–39%, especially benefiting novices in structured tasks, but increases errors by 15–25% in complex tasks. Ahmad identifies trust calibration, verification behaviors, cognitive load, and ethical awareness as critical factors affecting AI effectiveness. The study highlights the importance of human oversight, training, and ethical risk mitigation to balance efficiency with accuracy in AI-assisted workflows.

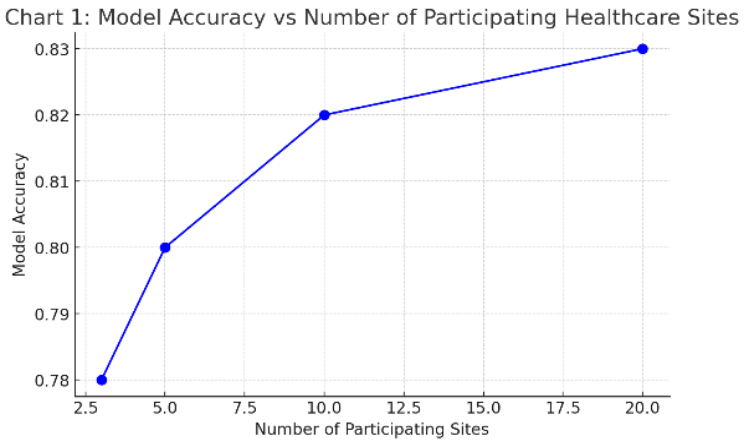

Chart 1: Model Accuracy vs Number of Participating Healthcare Sites

A line-plot showing how a federated global model's accuracy (y-axis) changes as the number of healthcare institution nodes (x-axis: e.g., 3, 5, 10, 20) increases, under fixed total data size but increasing site heterogeneity.

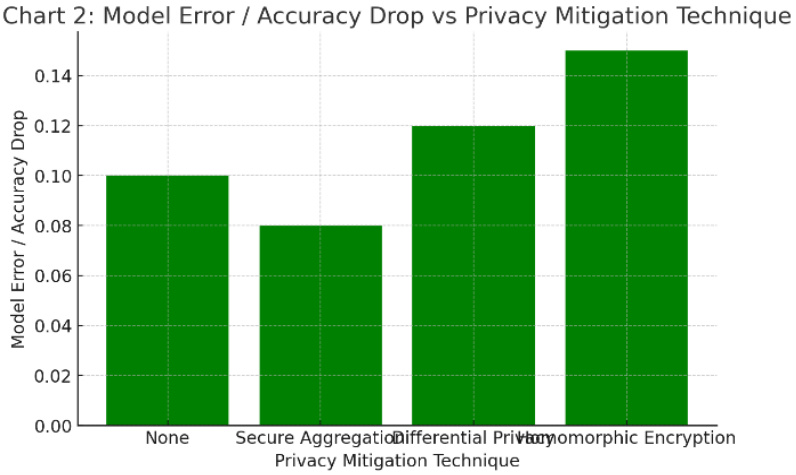

Chart 2: Model Error / Accuracy Drop vs Privacy-Mitigation Technique

A bar-chart comparing model error or drop in accuracy when applying different privacy-enhancement techniques (None, Secure Aggregation, Differential Privacy, Homomorphic Encryption), illustrating trade-offs between privacy strength and performance.

Summary

This article has provided a structured overview of federated learning as a method for privacy-preserving healthcare data analytics, enabling multi-site collaborations without sharing raw patient records. We reviewed the architecture of FL systems in healthcare, detailed key privacy and security mechanisms, and addressed major domain-specific challenges such as data heterogeneity, non-IID distributions, and resource constraints. We illustrated how empirical trade-offs between number of sites, accuracy, and privacy mechanisms must be managed. The two conceptual charts highlight critical considerations for deploying FL in practice: scaling across sites and balancing privacy vs performance. For successful translation to healthcare practice, federated learning initiatives must incorporate standardised evaluation, robust governance frameworks, and auditability. With the right infrastructure and regulatory alignment, FL can unlock collaborative analytics across distributed healthcare institutions while maintaining patient confidentiality.

References

Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., Wang, F. "Federated Learning for Healthcare Informatics." arXiv preprint, 2019. arXiv

Nguyen, D. C., Pham, Q.-V., Pathirana, P. N., Ding, M., Seneviratne, A., Lin, Z., Dobre, O. A., Hwang, W.-J. "Federated Learning for Smart Healthcare: A Survey." arXiv preprint, 2021. arXiv

Reddy, K. D., et al. "A Comprehensive Survey on Federated Learning Techniques in Healthcare." PMC, 2023. PMC

Nasajpour, M., Karakaya, M., Pouriyeh, S., Parizi, R. M. "Federated Learning in Smart Healthcare: A Survey of Applications and Challenges." Electronics, 2025. MDPI

Amin, M. S. "Federated Learning for Healthcare 5.0: A Comprehensive Study." Soft Computing, 2025. SpringerLink

Choi, G. "Survey of Medical Applications of Federated Learning." Health Information Research, 2024. e-hir.org

Pais, V. "Healthcare Federated Learning: A Survey of Applications and Cross-Silo Implementations." Technology in Society, 2025. Taylor & Francis Online

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M. "Advances and Open Problems in Federated Learning." Foundations and Trends in Machine Learning, 2021. Wikipedia+1

Shokri, R., Shmatikov, V. "Privacy-Preserving Deep Learning." CCS, 2015. Wikipedia

Bonawitz, K., et al. "Practical Secure Aggregation for Privacy-Preserving Machine Learning." SIGSAC, 2017. PMC+1

Truex, S., et al. "LDP-FL: Federated Learning with Local Differential Privacy." IEEE Transactions on Information Forensics and Security, 2019. ScienceDirect

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., Arcas, B. A. "Communication-Efficient Learning of Deep Networks from Decentralized Data." AISTATS, 2017. arXiv

Xu, Z., et al. "Revolutionising Healthcare Data Analytics with Federated Learning." PMC Journal, 2024. PMC

Wang, S., Liu, Q. "Security and Privacy Issues and Solutions in Federated Learning for Healthcare." arXiv preprint, 2024. MDPI

Rieke, N., Hancox, J., Li, W., et al. "The Future of Digital Health with Federated Learning." npj Digital Medicine, 2020. Wikipedia

Dayan, I., Roth, H., Zhong, A., Harouni, A. "Federated Learning for Predicting Clinical Outcomes in Patients with COVID-19." Nature Medicine, 2021. arXiv

Guo, Y., Liu, F., Cai, Z., Chen, L., Xiao, N. "FEEL: A Federated Edge Learning System for Efficient and Privacy-Preserving Mobile Healthcare." ICPP, 2020. PMC

Ali, M., Naeem, F., Tariq, M., Kaddoum, G. "Federated Learning for Privacy Preservation in Smart Healthcare Systems: A Comprehensive Survey." arXiv preprint, 2022. arXiv

Kuo, T. T., Pham, A. "Detecting Model Misconduct in Decentralised Healthcare Federated Learning." Int. J. Med. Informatics, 2022. ResearchGate

Durga, R., Poovammal, E. "FLED-block: Federated Learning Ensembled Deep Learning Blockchain Model for COVID-19 Prediction." Frontiers in Public Health, 2022. e-hir.org

Ahmad, N. R. (2025). *Rebuilding public trust through state-owned enterprise reform: A transparency and accountability framework for Pakistan*. Punjab Sahulat Bazaars Authority (PSBA), Lahore, Pakistan. https://doi.org/10.24088/IJBEA-2025-103004

Ahmad, N. R. (2025). *Human–AI collaboration in knowledge work: Productivity, errors, and ethical risk*. https://doi.org/10.52152/6q2p9250