



Federated Learning in Healthcare: A Privacy-Preserving Approach to Medical AI

Kothari Lakshmi kasturi¹, Jajam Naresh²

Amrita Vishwa Vidyapeetham, Amaravathi, Andhra Pradesh, India

Vignan's foundation for science, Technology and Research, Guntur, Andhra Pradesh, India

kothari.kasturi@gmail.com

Abstract

Artificial intelligence (AI) is reshaping healthcare by advancing predictive analytics, medical imaging, drug discovery, and personalized treatment. Yet, progress is constrained by the need for large, diverse datasets and the challenges of privacy, regulation, and institutional data silos. Federated learning (FL) has emerged as a privacy-preserving solution that trains models locally and shares parameters instead of raw data, enabling collaboration without compromising confidentiality. FL has been applied in cancer diagnosis, COVID-19 research, genomics, and electronic health records, showing strong potential for medical innovation. However, technical issues—such as data heterogeneity, communication overhead, model convergence, and security risks—along with economic and organizational barriers like cost, skills, and interoperability, slow adoption. Despite these challenges, FL is both a technological breakthrough and a strategic enabler for collaborative healthcare research. Hybrid approaches integrating FL with homomorphic encryption, differential privacy, and blockchain governance, combined with supportive policies and international cooperation, will be vital to realizing its full potential in healthcare.

Keywords: Federated learning, Healthcare AI, Privacy-preserving machine learning, Medical data, Collaborative intelligence

1. Introduction

Healthcare sector generates vast amounts of data in the form of electronic health records, scans, laboratory tests, and already worn devices. Data mining potential based on AI is promising as it could fundamentally change the nature of diagnosis, treatment, and drug development, but concentration of medical sensitive data raises privacy and control concerns that limit the extent of data storage at large scales (Brisimi et al., 2018; Xu et al., 2021).

In this case, the answer is federated learning: an algorithm that is decentralized with the parameters being transferred to the central node and models are trained locally so that they can work together without raw data being transferred (Kairouz et al., 2021). This enables hospitals, clinics, and research organizations in the medical sector to develop models together and archive patient data locally (Sheller et al., 2020).

The present paper explains the potential of the FL to solve the dilemma between data-driven innovation and privacy protection in healthcare, its technical basis, applications, benefits and limitations thereof, and outlines the opportunities of applying it sustainably to clinical practice (Kairouz et al., 2021; Xu et al., 2021).

2. Background of the Study

Older machine learning models rely on centralized datasets, which in healthcare is not preferred due to privacy concerns and regulations as well as institutional siloing (Brisimi et al., 2018; Xu et al., 2021). FL is a decentralized algorithm for training, which relies on transporting the models to the local scenarios and acquiring the updates there and upholding the privacy and ensuring the nonhomogenous environment learning feasible (Kairouz et al., 2021). Initial applications have demonstrated FL potential in various clinical settings such as COVID-19 CT scans and cancer imaging, but also in wearable analytics-applications with data privacy concerns that are of special relevance

3. Justification

This research will be important due to five reasons. First, in terms of privacy, a solution must be created that guarantees that patient confidentiality will be preserved in the creation of a model (Brisimi et al., 2018). Second, the regulatory balance is skewed to the strategies that would reduce the data flows, the exposure risk (Kairouz et al., 2021). Third, data diversity and the presence of infrequent cases can be achieved in an institutions-collaborating setting without sharing the raw data (Xu et al., 2021). Fourth, AI creation needs robust and representative datasets that may be structured with the support of FL (Kairouz et al., 2021). Finally, privacy as an ethical concern and problem has posed a wakeup call to the unavoidable implementation of AI into clinical practice (Truex et al., 2019).

4. Objectives of the Study

This paper aims to:

1. Apply the notion of federated learning in medical practice.
2. Get to grips with its applications in medical AI.
3. Identify technical, economic and organizational barriers.
4. Nurse review case studies of FL.
5. Prepare future adoption and research recommendations.

5. Literature Review

Kairouz et al. (2021) provide an overview of FL which is defined as training information that is not centralized, and enumerates open problems. Sheller et al. (2020) introduce multi-institutional brain tumor segmentation using FL and show similar results to the centralized approaches. Xu et al. (2021) define FL as a healthcare informatics tool that aims to enhance the generalization effect with the help of various, distributed data.

The second issue is that it can be damaged by non-IID information, bottlenecks in communication, and adversarial risks, which can destroy it based on performance and safety (Kairouz et al., 2021; Bagdasaryan et al., 2020). Other benefits of FL in terms of model update leakage are secure aggregation and differential privacy (Truex et al., 2019). According to EHR-based studies, predictive modeling could be possible even without exchanging raw data in order to justify the clinical relevance of FL (Brisimi et al., 2018). Overall, the literature identifies the potential of FL and technical and governance complexity (Kairouz et al., 2021; Xu et al., 2021).

6. Material and Methodology

The current study uses a qualitative review approach to explore how federated learning (FL) can be used in healthcare as a privacy-saving solution to medical AI. The three primary elements are used in the analysis. To begin with, peer-reviewed journal articles, reports by other organisations like the World Health Organization (WHO) and the Organisation Economic Co-operation and Development (OECD), and technical papers delivered at large AI conferences were all included in the analysis of secondary data. Second, several case studies were reviewed to outline the examples of the practical use of FL, such as tumor segmentation using brain imaging, respiratory disease diagnostics with COVID-19, and the modeling of rare diseases. These case studies have been chosen since they illustrate the scope of the applications of FL in clinical and research setting. Third, we performed a comparative study in which we compared the FL with centralized machine learning models, with other methods that are privacy-sensitive. This comparison offered information about the comparative benefits, drawbacks, and feasible issues of applying FL within health care.

7. Results and Discussion

7.1 The use of federated learning in healthcare has been applied in the following ways:


Several fields of health care have used federated learning. Segmentation of brain tumors Medical imaging FL has been applied in the brain tumor segmentation area, is shown in Sheller et al. (2020), as well as lung CT scan analysis to detect COVID-19. FL has also been used with electronic health records (EHRs), in which predictive models are trained on shared patient data without loss of privacy. FL as a concept in the context of genomics can be applied to share rare genetic conditions across institutions. Wearable health monitoring is another space that FL can leverage, allowing specific and privacy-conscious AI solutions that evolve as the needs of various patients alter.

Table 1: Applications, Benefits, and Barriers of Federated Learning in Healthcare

Application	Benefit	Barrier
Medical Imaging (Brain Tumor, COVID-19 CT)	Enables multi-institutional imaging model training without raw data sharing	High communication overhead; data heterogeneity
Electronic Health Records (EHRs)	Predictive modeling with privacy compliance (HIPAA/GDPR)	Interoperability issues; infrastructure cost
Genomics & Rare Diseases	Collaboration on rare genetic conditions across borders	Limited datasets; regulatory fragmentation
Wearable Health Monitoring	Personalized, privacy-preserving real-time monitoring	Technical complexity; lack of awareness among practitioners

7.2 Healthcare FL Advantages

The benefits of FL can be seen in a number of ways. The first advantage is privacy protection since the information about patients is kept locally in hospitals and institutions. FL also can make data more diverse, since modeling in collaboration with multiple centers can better generalize across populations and conditions. Better still, FL assists with regulatory compliance, which aligns with the Health Insurance Portability and Accountability Act (HIPAA) of the U.S. or the General Data Protection Regulation (GDPR) of the EU. Finally, FL encourages collaboration via multi-institutional research facilitation which does not necessitate sharing of data.

**Graph 1: Comparative Benefits and Barriers of FL Applications in Healthcare**

7.3 Technical Barriers

Although this is the case, FL experiences a number of technical challenges. The non-independence and non-identical distribution of data between local datasets (data heterogeneity) makes it harder to train models and converge. Another difficulty is communication overhead because, when the number of exchanged parameters is high, the bandwidth between local models and the central server is large. Moreover, other types of security attacks like model poisoning and inference attacks are also faced by FL and can result in breaches of accuracy and privacy. Lastly, convergence in models between different datasets is challenging to attain and this leads to poor performance relative to centralized models.

7.4 Economic and Organizational Barriers

There are additional economic and organizational concerns that make the implementation of FL difficult. High infrastructure and implementation costs are not easy to embrace by resource-constrained institutions. Lack of interoperability of different hospital IT systems is also the second major problem because the flow of FL across different platforms is not fully rolled out. Moreover, health practitioners and administrators lack the knowledge and experience in utilizing the potential applications and benefits of FL, as a significant number of them are unaware of the tool.




Figure 2: Technical vs Organizational Barriers in Healthcare Federated Learning

7.5 Case Studies

A number of case studies show the potential and the difficulties of FL in healthcare. A U.S. and international initiative called the EXAM Project used FL to create AI models to diagnose COVID-19 without any patient-centered data centralization. The Federated Tumor Segmentation (FeTS) challenge provided a way to test the scalability of FL by uniting more than 30 institutions to build AI models on brain imaging together. Ongoing European projects engage FL in areas such as cancer research and modeling rare diseases once again proving the relevance of the tool and highlighting the existing organizational and regulatory obstacles.

8. Limitations of the Study

Citations to the secondary resources also minimise the possibility of subsequent validation in real healthcare facilities (Xu et al., 2021). Application fields highlighted are imaging and diagnostics, and other fields are relatively underrepresented in this review (Sheller et al., 2020; Xu et al., 2021). In addition, the FL study is carried out at a very high frequency, which is a testament to the risk of obsolescence of some conclusions with the introduction of new algorithms and protocols (Kairouz et al., 2021). Lastly, quantitative benchmarking or attack-resilience metrics that are necessary to measure safety and systemic risks are not presented by us (Bagdasaryan et al., 2020; Truex et al., 2019).

9. Future Scope

The directions to be taken in the future include a composite of FL with privacy and resilience practices based on a differential privacy concept and secure aggregation, as well as, homomorphic encryption (Truex et al., 2019; Kairouz et al., 2021). Multi-site deployments will become easier with the development of interoperability and assessment standards and protocols (Kairouz et al., 2021). The institutionalization of cross-border partnerships would help to achieve modeling inclusiveness and generalization (Xu et al., 2021). Pilots development process should become normalized, should be toolled in MLOps, process-regulated, and should be accompanied by clinical (Xu et al., 2021). Lastly, explainable AI can improve trust and clinical adoption as well as contribute to auditability in the context of FL pipelines (Kairouz et al., 2021).

10. Conclusion

Federated learning is a healthcare AI paradigm shift that allows cooperation without violating privacy. It also covers

the issues of control, data diversification, and building of confidence in cyber health. Nonetheless, it is important to overcome technical, economic, and organizational barriers.

The conclusion to this paper is that federated learning is not a silver bullet, but a building block to privacy-saving AI. The combination of FL and other complementary technologies and policy supportive measures will be needed to unlock the full potential of the technology. Federated learning will empower the new era of healthcare innovation through collaboration, security and transparency through AI.

References

1. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *International Conference on Artificial Intelligence and Statistics*, 2938–2948.
2. Brisimi, T. S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I. C., & Shi, W. (2018). Federated learning of predictive models from federated EHR data. *International Journal of Medical Informatics*, 112, 59–67.
3. Kairouz, P., McMahan, H. B., Avent, B., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2), 1–210.
4. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1), 12598.
5. Truex, S., Liu, L., Gursoy, M. E., Yu, L., & Wei, W. (2019). Demystifying differential privacy in federated learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(5), 1029–1043.
6. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19.