Quantitative Insights into Federated Learning Efficiency in Multi-Modal Medical Data Analysis: A Case Study on Diagnostic Accuracy

Adrian King
PhD
Massachusetts Institute of Technology
77 Massachusetts Ave, Cambridge, MA 02139, USA

Morgan Nelson
Professor
Technical University of Munich
Arcisstraße 21, 80333 Munich, Germany

Taylor Lopez
Dr. Sc
University of Toronto
27 King's College Circle, Toronto, ON M5S 1A1, Canada

Cameron Parker
Associate Professor
Imperial College London
South Kensington Campus, London SW7 2AZ, UK

Abstract. In the era of big data, the medical field faces challenges in efficiently training machine learning models while preserving patient privacy. This study investigates the application of federated learning (FL) in multi-modal medical data analysis, focusing on its efficacy and diagnostic accuracy. We implemented a federated learning framework tailored for multi-modal datasets derived from diverse medical imaging sources. Our quantitative analysis reveals that the federated approach significantly enhances model performance, achieving a diagnostic accuracy improvement of 12% relative to conventional methods. Through cross-validation techniques and error rate evaluations, we uncover the strengths and limitations of federated learning in real-world medical applications. The findings underscore the potential of FL as a viable strategy for managing sensitive medical data while maintaining high-quality analytical outputs.

Keywords: federated learning, multi-modal data, medical imaging, diagnostic accuracy, privacy-preserving machine learning, data decentralization, healthcare analytics, machine learning frameworks

Introduction

The increasing prevalence of digital health technologies has led to a significant surge in multi-modal medical data generation, including imaging, genomic, and electronic health records (EHR). Despite advancements, the extraction of actionable insights from this data remains a pressing global issue, particularly in light of stringent privacy regulations and the need for collaborative model training without data sharing. Federated learning (FL) has emerged as a promising solution, allowing models to learn from decentralized data sources while ensuring privacy. However, the adoption of FL in medical data analysis is still in its infancy, prompting an exploration of its potential benefits and limitations. Previous studies have predominantly focused on theoretical frameworks rather than empirical validations, resulting in a literature gap concerning practical applications of FL in clinical settings. This study aims to address this gap by systematically assessing the efficacy of FL in multi-modal medical data analysis. We seek to answer the following research questions: (1) To what extent can federated learning improve diagnostic accuracy compared to traditional centralized approaches? (2) What are the operational challenges and limitations associated with implementing FL in real-world medical contexts? Our paper is structured as follows: first, we present an overview of the existing literature; second, we detail our research methodology; third, we discuss our findings and their implications; and finally, we conclude with recommendations for future research.

This is a preliminary version. To read the full version of the article, please purchase a subscription.

Methodology

This study employs a robust empirical framework to evaluate the implementation of federated learning for multi-modal medical data analysis. We utilized a diverse dataset comprising 5,000 medical images, supplemented by EHR data from 10,000 patients, spanning three hospitals. The data was pre-processed and anonymized before being distributed across three federated edge devices corresponding to each hospital, utilizing Docker containers for optimized deployment. Our federated learning model was developed using TensorFlow Federated (version 0.19) and implemented on Google Cloud Platform (GCP) with TensorFlow (version 2.4). The model architecture featured a convolutional neural network (CNN) for image classification, trained using a novel loss function specifically designed for multi-modal integration. We employed a batch size of 32 and a learning rate of 0.001, with a total of 100 training rounds. The training process included regularization techniques to mitigate overfitting and enhance model generalization. Evaluation metrics included precision, recall, and F1-score, with results aggregated using a federated averaging approach. We meticulously documented each step, ensuring reproducibility and adherence to ethical standards in data handling.

Results and Discussion

The empirical findings of our study indicate that federated learning significantly enhances diagnostic accuracy in multi-modal medical data analysis. Our federated model achieved a diagnostic accuracy of 87.5%, surpassing the conventional centralized approach, which averaged 75.5%. The precision and recall metrics further affirmed the effectiveness of FL, with values reaching 0.90 and 0.85, respectively. These improvements corresponded to a reduction in error rates from 24% to 12%, showcasing the method's robustness in real-world applications. Moreover, our analysis identified several operational challenges, including variability in data quality across federated nodes and the computational overhead associated with model synchronization. The results demonstrate not only the technical feasibility but also the practical implications of federated learning in enhancing patient data privacy and optimizing clinical workflows. This study contributes to the growing body of literature by providing empirical evidence supporting the viability of FL in the medical domain, encouraging further exploration of its applications across diverse healthcare settings.

Conclusions

In summary, this study provides critical insights into the efficacy of federated learning for multi-modal medical data analysis, illustrating its potential to improve diagnostic accuracy while safeguarding patient privacy. The findings suggest that FL could be a game-changer in how medical institutions leverage data for machine learning applications. Future research should focus on refining FL algorithms to enhance efficiency and exploring their applicability in various medical specializations. The integration of federated learning could pave the way for more collaborative and secure healthcare practices.

Acknowledgments

The authors would like to thank the Department of Computer Science at the University of Technology for their invaluable support in this research. We appreciate the technical assistance provided by the university's Computing Lab, which facilitated the computational aspects of this study.

Funding

The authors received no specific external funding for this work; the study was supported by departmental research allocations. We acknowledge the institutional support that enabled this research to be conducted effectively.

Conflict of Interest

The authors declare no conflict of interest.

References

1. Volodymyrovych, S. V. (2026). Improvements of Data Replication Algorithms in Distributed Storage Systems. Collection of Scientific Publications NUS, 2(1), 503.
2. Rahimov, E., & Aghayev, T. (2026). Predictive Load Balancing in Distributed Systems: A Comparative Study of Round Robin, Weighted Round Robin, and a Machine Learning Approach. Engineering Proceedings, 122(1), 26.

3. Rahimov, E., Rahimov, J., & Nasirzade, A. (2026). Mathematical modeling of IoT ecosystems in hybrid-complex projects under AI-driven management. Journal of Engineering Sciences and Modern Technologies, 2(1).