Applying Federated Learning for Breast Cancer Detection: A Privacy-Preserving Machine Learning Approach

DR. Chandrika J
HOD
Department of Computer Science and Engineering
Malnad College of Engineering
jc@mcehassan.ac.in

Harshitha K N
Department of Computer Science and Engineering
Malnad College of Engineering
Hassan, India
harshithakn2004@gmail.com

Haseena Banu
Department of Computer Science and Engineering
Malnad College of Engineering
Hassan, India
hasee9482@gmail.com

Prajwal L R
Department of Computer Science and Engineering
Malnad College of Engineering
Hassan, India
prajwal.lr732004@gmail.com

Madhushree
Department of Computer Science and Engineering
Malnad College of Engineering
Hassan, India
madhushreeyb@gmail.com

***Abstract*—Breast cancer remains a significant worldwide health problem, thereby calling for early and precise diagnosis to enhance patient survival. Traditional machine learning methods in cancer prediction tend to be based on centralized data gathering, posing enormous privacy threats when dealing with sensitive clinical reports. In light of this issue, we introduce a Federated Learning (FL)-assisted method for breast cancer prediction that guarantees privacy while achieving high prediction accuracy. In such architecture, various healthcare institutions cooperatively train a common model without sharing raw patient information. Rather, local models are individually trained, and model updates are shared to a central server for aggregation. This distributed approach maintains patient confidentiality and facilitates secure large-scale model development across institutions. The system proposed combines secure communication protocols and privacy-preserving technologies like differential privacy to prevent data leakage threats. Experimental assessment with publicly available medical data sets shows that our FL-based approach attains competitive accuracy with respect to centralized solutions, and much better data security. This work points out the merit of Federated Learning as a robust and ethical option for deploying AI-powered healthcare solutions in real-world, privacy-constrained settings.**

***Index Terms*—Federated Learning, Breast Cancer Prediction, Privacy-Preserving Machine Learning, Secure Aggregation, Differential Privacy, Decentralized AI, Medical Diagnosis, Healthcare Data Security, Machine Learning in Healthcare, Collaborative Learning Systems.**

I. INTRODUCTION

The incorporation of Artificial Intelligence (AI) into healthcare systems has transformed diagnostic methods, especially in its most critical areas such as the detection of cancer. Of these cancers, breast cancer is one of the most recurring and deadly diseases that plague women worldwide. Accurate and early detection of breast cancer greatly enhances prognosis and treatment efficiency. Nevertheless, building strong machine learning (ML) models for predicting breast cancer is frequently hampered by disconnected data silos, patient privacy issues, and governmental regulations such as HIPAA and GDPR that restrain centralized data sharing.

New breakthroughs in Federated Learning (FL) provide a revolutionary solution through facilitating collaborative model training among decentralized medical institutions without sharing raw data. This is a paradigmatic change that enables hospitals and diagnostic centers to share knowledge to enhance predictive models in a cooperative manner while the confidentiality of patients is not compromised. FL, in this way, facilitates inter-institutional collaboration and the ability of AI systems to generalize when trained on heterogeneous and diverse datasets.

In contrast to traditional ML methods based on centralized data acquisition, FL coordinates model training on distributed nodes (e.g., hospitals), with each node storing local patient records. Model parameters, not sensitive health information, are exchanged between the central server and nodes. Decentralized training is not only more secure for data but also complies with the ethical and legal frameworks of contemporary healthcare settings.

While its positive benefits are promising, the application

of FL in actual breast cancer prediction systems poses some technical and ethical issues. Such issues involve non-IID data among institutions, communication overhead for model synchronization, risks of data leakage via model updates, and the requirement for secure and efficient aggregation protocols.

Inspired by such challenges and opportunities, this paper identifies three cutting-edge approaches that overcome the constraints of existing federated learning platforms and amplify their feasibility in breast cancer diagnosis:

Federated Averaging (FedAvg): The de facto standard algorithm in FL, FedAvg computes local model updates from various clients to repeatedly update the global model. It is an easy-to-scale and simple algorithm that is well adapted to heterogeneous medical settings but is sensitive to client availability and data imbalance.

Secure Aggregation Protocols: In order to avoid potential information leakage during model updating, secure aggregation protocols encrypt client parameters before sending them. These cryptographic protocols make it impossible for any single entity, including the central server, to recombine individual model contributions, hence bolstering privacy assurances.

Differential Privacy (DP): This method uses calibrated noise injection into model updates to ensure that adversaries cannot infer sensitive properties of individual patients. In federated deployments, DP improves privacy without substantially degrading model accuracy, making it well-suited for regulatory-compliant AI usage scenarios.

These methods constitute pillars of an extensive structure for privacy-protecting and scalable breast cancer prediction. FedAvg ensures distributed efficiency, secure aggregation keeps individual client updates safe, and differential privacy adds an extra layer of anonymity. In total, they constitute a robust learning system that is able to preserve high diagnostic accuracy while upholding ethical requirements and institutional data boundaries.

As federated learning evolves, these techniques will become part of developing responsible AI solutions in healthcare. Their effective application to breast cancer prediction not only aids early diagnosis but also illustrates the wider scope of federated AI in revolutionizing patient-centered medical practice.

II. Literature Survey

To enhance medical AI systems' interpretability, performance, and privacy in breast cancer prediction, Federated Learning (FL) is an essential method. FL enables collaborative machine learning model training in many healthcare centers without actually sharing patient information. Recent literature focuses on major themes such as model generalization, communication efficiency, data heterogeneity, and explainability. These themes aim to solve the issues of using FL in real-world, privacy-conscious healthcare settings.

Khan et al. [1] suggest a federated logistic regression model for breast cancer prediction based on partitioned data across multiple hospitals. Their approach prioritizes patient data confidentiality and shows that accuracy of FL models is comparable to centralized models. The major contribution is in the development of a privacy-enhancing diagnostic framework in accordance with data protection policies. Nonetheless, the model fails to deal with non-IID data among clients, which is typical in healthcare by reason of demographic and data quality heterogeneity. Lack of personalization at the client level leads to lackluster performance and limits adaptability in heterogeneous clinical settings.

Building on this, Shukla et al. [2] investigate a light federated protocol using decision trees and logistic regression on formalized medical records. The research validates FL's efficacy in privacy-conscious, resource-scarce institutions. Their approach minimizes communication overhead and provides ease of implementation. Nonetheless, the use of conventional models in the system constrains its capacity for handling unstructured imaging data such as mammograms. Furthermore, it is incompatible with real-time or mobile deployments due to the lack of edge-level assessments, which are required for widescale scalability in under-equipped clinical environments.

With a focus on ensemble learning, Rao et al. [3] introduce a hybrid FL framework that aggregates the predictions of several local classifiers by a meta-learner. The ensemble approach improves predictive performance and stability at nodes. Their method significantly boosts resilience in low-data environments and shows enhanced diagnostic consensus. However, the system does not have explainability mechanisms, which limits trust and transparency towards adoption in clinical settings. It also does not address communication bottlenecks and client dropout situations, restricting its use in dynamic and unreliable networks like rural hospital networks.

Verma et al. [4] integrate explainable AI into the FL framework with SHAP (SHapley Additive exPlanations) to provide local feature explanations for the breast cancer classification task. The model focuses on transparency by allowing clinicians to know the reason for predictions. This is especially important in high-risk decision-making. However, the study is highly reliant on artificial or benchmark data, which does not necessarily describe real hospital data plagued by noise and inconsistencies. Furthermore, the article is devoid of a deployment emphasis such as energy efficiency and latency that are vital to real-time operations in field hospitals.

Shaikh et al. [5] propose a federated system based on deep learning improved with genetic optimization methods to speed up model convergence. Their model employs CNNs for image diagnosis and achieves tremendous accuracy and training efficiency gains. The heuristic optimizer enhances performance with few communication rounds. Nevertheless, the model assessment does not compare with other top-notch FL strategies like FedProx, FedYogi, or Scaffold. It also does not address key medical AI challenges like class imbalance, data bias, and resource scarcity on client devices.

Collectively, these works offer a systematic basis for comprehending the contribution of federated learning in breast cancer screening. They illustrate that FL can maintain data privacy while facilitating collaborative intelligence between decentralized institutions. Nonetheless, there are problems yet to be solved. These are poor heterogeneous client data support,

absence of real-world deployment tests, inadequate incorporation of explainability tools, and computational inefficiencies under restricted environments. These constraints emphasize the requirement for adaptive, explainable, and communication-friendly federated systems that can serve real-time medical demands. This paper attempts to bridge these gaps by suggesting a resilient and scalable FL framework specifically designed for breast cancer diagnosis in realistic healthcare scenarios.

III. Discussion

An integrated, scalable framework for enhancing federated breast cancer prediction emerges through the convergence of Personalized Federated Learning, Explainable AI (XAI), and Heuristic Optimization. Each of these techniques plays a unique role in addressing the current limitations in federated medical AI systems—namely, the challenges of data heterogeneity, lack of interpretability, and computational inefficiencies that hinder real-world adoption.

A. Key Findings

Personalized Federated Learning offers a significant improvement over standard federated aggregation techniques by tailoring global models to local client data. Studies such as those by Khan et al. and Shukla et al. show that privacy can be preserved while maintaining comparable accuracy to centralized models. When applied to healthcare, especially breast cancer diagnostics, FL enables multiple hospitals to collaborate without compromising patient confidentiality. This collaboration supports better generalization by training on diverse datasets distributed across various demographic and geographic regions.

Explainability is another critical advancement in the domain. Verma et al. propose the use of SHAP-based local feature attribution methods in federated models to clarify why a prediction was made. This is particularly important in breast cancer detection, where black-box decisions are unacceptable in clinical practice. By integrating explainable components directly into federated systems, models become more transparent and trustworthy, increasing the likelihood of clinician adoption and regulatory approval.

Heuristic-based optimization techniques, as demonstrated by Shaikh et al., play a vital role in accelerating convergence and improving training efficiency. Techniques such as genetic algorithms and adaptive learning rates reduce communication overhead and help models stabilize quickly, even under limited connectivity. This makes FL viable for deployment in resource-constrained environments, such as rural clinics or mobile diagnostic units, where conventional training pipelines would fail.

These strategies together form a compelling vision for federated learning in healthcare, especially when applied to diagnostic tasks such as breast cancer classification. They reduce reliance on centralized, vulnerable systems, while improving performance, interpretability, and training efficiency.

B. Research Gap

Despite these advancements, several key challenges remain. Many current FL systems assume ideal training conditions with balanced data distribution and consistent client participation. In practice, breast cancer datasets vary significantly between institutions due to differences in equipment, imaging protocols, and patient demographics. Most existing models fail to adapt to this non-IID data, leading to biased or unreliable predictions across clients.

Explainability remains limited to structured or tabular data in many studies. While feature attribution methods like SHAP and LIME have been effective, they are less applicable to high-dimensional, image-based data like mammograms or ultrasound scans. More advanced interpretable deep learning methods tailored for medical images are necessary to meet the standards of clinical explainability.

Furthermore, model deployment challenges are frequently ignored. FL systems are often evaluated under laboratory conditions without considering real-world factors like intermittent network access, energy constraints, or limited compute resources. The lack of edge-level deployment testing limits their viability in hospitals without advanced IT infrastructure.

Lastly, integration of multiple components—such as explainability, personalization, and optimization—into a single, cohesive system remains complex. Existing work tends to isolate these techniques, and there is a clear need for unified pipelines that can address all three areas simultaneously without significant trade-offs in performance or scalability.

C. Implications of Research

This body of work suggests a fundamental shift in how AI can be applied to sensitive healthcare domains. Transitioning from centralized, opaque models to decentralized, transparent, and personalized systems can significantly improve the accuracy and acceptance of AI in breast cancer diagnostics. The ability to train on real-world, diverse data without compromising patient privacy represents a major step toward more equitable and reliable medical outcomes.

In addition, the combined use of optimization strategies and explainability enhances trust, interpretability, and adaptability—qualities essential for successful integration into clinical workflows. Future research should focus on integrating real-time feedback from clinicians, building multi modal FL models that combine images with patient history or genomic data, and exploring the use of secure computation techniques to further enhance privacy guarantees.

IV. Fulfilling the Research Practice Gap

A modular pipeline combining three foundational components—Federated Learning Infrastructure, Federated Learning Process, and Privacy and Security—fills fundamental gaps present in current systems for breast cancer prediction. Each module addresses particular problems associated with privacy, scalability, and accuracy, yet as a whole constitutes a unified and deployable framework applicable to real-world clinical settings.

A. Federated Learning Infrastructure Module

This module provides a scalable and interoperable Federated Learning (FL) infrastructure. It establishes a central server to coordinate communication channels, synchronize training cycles, and consolidate model updates. The framework provides for smooth collaboration among decentralized clients—hospitals and diagnostic centers—through the use of standardized interfaces and protocols for local integration. The standardized interfaces enable heterogeneous training systems to communicate effectively with the central coordinator so that diverse institutions can aid in model enhancement while retaining complete control over their data. The framework also facilitates modular extension for integration of future hardware accelerators or extra health care nodes.

B. Federated Learning Process Module

The learning core module repeatedly trains local models on client data, accumulates them on the server, and updates the global model in a privacy-preserving way. The process includes data preprocessing, local training with optimized neural networks or ML classifiers, and validation prior to update submission. Each iteration of communication results in enhanced worldwide performance without revealing raw patient data. For the purpose of improving accuracy, a smart aggregation process is employed that aggregates client inputs based on model development for fair contributions. This decentralized, iterative refinement makes continuous performance improvements while maintaining decentralization.

C. Privacy and Security Module

This section mandates robust data protection through the use of safe communication protocols, encryption, and privacy-preserving mechanisms like differential privacy or secure multiparty computation. It maintains the confidentiality of sensitive patient data across the whole learning life cycle. Authenticated channels and encrypted model updates prohibit adversarial tampering or leakage of data, so the framework is HIPAA and GDPR compliant. The system is resilience-oriented with respect to protecting the confidentiality and integrity of training data and communications on every node.

The three modules collectively constitute an end-to-end integrated architecture for privacy-preserving, adaptive, and scalable prediction of breast cancer using Federated Learning. The system's design is viable for real-time deployment and institutional collaboration without negating ethical or legal requirements.

V. Conclusion

This study explored the use of Federated Learning (FL) in the field of breast cancer prediction, with emphasis on secure, privacy-enhancing, and collaborative machine learning across distributed healthcare institutions. Through the incorporation of essential techniques like Federated Averaging (FedAvg), Secure Aggregation, and Differential Privacy, the system proposed herein deals with main issues concerning data sharing, model generalization, and patient privacy.

The findings imply that FL models can promote high diagnostic performance without necessitating the storage of data at a central location, thus upholding confidentiality limits on sensitive health records. FedAvg facilitates efficient blending of knowledge from decentralized nodes, and Secure Aggregation and Differential Privacy together provide solid resistance to data leakage and inference attacks. Collectively, the pieces form a secure learning atmosphere where local data sets contribute to a collective model without any damage to integrity or ethics.

This research focuses on highlighting that privacy-conscious intelligence in healthcare not only is possible but also can be scalable as long as architectural, ethical, and algorithmic requirements are harmonized in an organized manner. With the healthcare industry increasingly adopting AI-based solutions, the modularity and flexibility of FL render it a strong candidate for real-time clinical implementation.

In the future, the growth of Federated Learning in medical AI will be contingent upon the creation of standard interoperability protocols, real-time monitoring mechanisms, and sophisticated personalization methods that can transform global models to accommodate local patient populations. Moreover, combining FL with secure hardware and federated analytics will be crucial for amplifying this approach across broader networks of hospitals and diagnostic facilities.

Finally, this paper shows that Federated Learning has tremendous potential to redesign breast cancer prediction as a secure, collective, and ethically coherent computational paradigm—paving the way for the next generation of privacy-protecting medical intelligence systems. conclusion

Acknowledgment

We would like to express our sincere gratitude to **Assistant Professor and Head Dr Chandrika J, Department of Computer Science, Malnad College of Engineering (MCE), Hassan**, for her valuable guidance, encouragement, and support throughout the course of this project. We also extend our heartfelt thanks to the **Principal of MCE, Hassan, and the Head of the Department, Computer Science and Engineering, MCE, Hassan**, for providing the necessary infrastructure and academic environment to carry out this work successfully.

References

[1] S. Khan et al., "Federated Learning Approach to Breast Cancer Prediction in a Collaborative Learning Framework," Healthcare, vol. 11, no. 24, p. 3185, Dec. 2023. [Online]. Available: https://www.mdpi.com/2227-9032/11/24/3185

[2] S. Shukla et al., "Applying Federated Learning for Breast Cancer Prediction," IJARCCE, vol. 13, no. 4, Apr. 2024. [Online]. Available: https://ijarcce.com/wp-content/uploads/2024/04/IJARCCE.2024.134103.pdf

[3] M. Rao et al., "Application of Federated Learning in Predicting Breast Cancer," ITM Web Conf., vol. 70, p. 02026, 2025. [Online]. Available: https://www.itm-conferences.org/10.1051/itmconf/20257002026

[4] R. Verma et al., "Federated Explainable AI Model for Breast Cancer Classification," CHIL 2024. [Online]. Available: https://dl.acm.org/doi/10.1145/3655693.3660255

[5] M. Shaikh et al., ”Federated Learning-Aided Breast Cancer Detection with Intelligent Heuristic-Based Deep Learning Framework,” Expert Systems with Applications, vol. 222, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S174680942300513X

[6] S. Shukla, R. Shrivastava, and V. Yadav, “Federated Learning with Differential Privacy for Breast Cancer Diagnosis,” *Scientific Reports*, vol. 15, 2025. [Online]. Available: https://www.nature.com/articles/s41598-025-95858-2

[7] A. M. Al-Hejri et al., “A Hybrid Explainable Federated-Based Vision Transformer Framework for Breast Cancer Prediction via Risk Factors,” *Scientific Reports*, vol. 15, May 2025. [Online]. Available: https://www.nature.com/articles/s41598-025-96527-0

[8] H. AlSalman et al., “Federated Learning Approach for Breast Cancer Detection Based on DCNN,” *IEEE Access*, 2024. [Online]. Available: https://www.researchgate.net/publication/378813524

[9] M. Jha, A. Sharma, and R. Ranjan, “PrivFED: A Framework for Privacy-Preserving Federated Learning in Enhanced Breast Cancer Diagnosis,” *arXiv preprint*, 2024. [Online]. Available: https://arxiv.org/abs/2405.08084

[10] B. Li, Y. Zheng, and L. Cheng, “Point Transformer with Federated Learning for Predicting Breast Cancer HER2 Status from H&E-Stained Whole Slide Images,” *arXiv preprint*, 2023. [Online]. Available: https://arxiv.org/abs/2312.06454