

Decentralized Federated Deep Learning for Early Detection of Breast Cancer in Heterogeneous Health Records

Abirami S.P¹, Rajasree S², Poornima S³, Priya S⁴ and Priyadharshini A⁵

¹⁻⁵Coimbatore Institute of Technology, Tamil Nadu, India

Email: abirami.sp@cit.edu.in, rajasree.s@cit.edu.in, poornima.s@cit.edu.in, priya.s@cit.edu.in, priyadharshini@cit.edu.in

Abstract— One of the main causes of death for women worldwide is still breast cancer, which emphasises the need for early and precise diagnosis. Because of institutional silos, traditional centralised machine learning algorithms for cancer diagnosis frequently confront obstacles with data privacy, complying with regulations, and low data variety. This study suggests a decentralised federated learning (FL)-based, productive, secure privacy-based breast cancer detection framework that can address these shortcomings. While maintaining data locality, the recommended method makes use of an assortment of datasets from open-access medical resources such as Kaggle and the UCI Repository. To guarantee compatibility and privacy, preprocessing involves standardising formats, suppressing noise, and anonymising medical information. Local model training is performed amongst several distributed clients employing the Federated Averaging (FedAvg) and Federated Optimisation (FedOpt) algorithms. These locally trained models are securely transferred and aggregated using privacy-preserving approaches like secure aggregation and Byzantine tolerance for faults to create a global model that does not share raw patient data. The system architecture consists of model initialisation, decentralised local training, secure update exchange, and aggregation. Validation datasets are evaluated in order to determine their performance and generalisation. Furthermore, differential privacy and secure multiparty computation are combined to improve data confidentiality and conformance to legal as well as ethical standards. This federated tactics improves model robustness by using heterogeneous data while nonetheless guaranteeing institutional autonomy and regulatory compliance. The experimental results demonstrate the framework's adaptability for real-world applications in medicine. The subsequent phases of this study will look into integrating with multi-modal data sources and edge computing, with the goal of creating a scalable and explainable AI-based diagnosing infrastructure for healthcare.

Index Terms— Federated Learning, Healthcare Analytics, privacy preservation, Federated Optimisation, cancer detection.

I. INTRODUCTION

Breast cancer is the most prevalent malignancy among women worldwide, with an estimated 2.3 million new cases and 685,000 deaths reported in 2020 alone, according to the World Health Organization. Early detection is crucial for improving survival rates—when diagnosed at an early stage, five-year survival rates exceed 90% [1].

Medical imaging modalities such as mammography, magnetic resonance imaging (MRI), and ultrasound remain central to early diagnosis. However, accurate interpretation of these images requires significant expertise, and diagnostic consistency often varies across institutions and regions.

Advances in machine learning (ML), particularly deep learning, have transformed oncology workflows, enabling automated interpretation of medical images. Convolutional Neural Networks (CNNs) and other deep learning architectures have achieved state-of-the-art performance in breast cancer detection and classification tasks [2]. Nonetheless, the success of these models depends heavily on access to large-scale, diverse, and well-annotated datasets.

In practice, assembling such datasets is challenging due to strict data privacy regulations (e.g., HIPAA, GDPR), ethical constraints, and institutional data governance policies that limit data sharing.

Federated Learning (FL), introduced by Google in 2016, has emerged as a promising alternative to centralized ML. FL enables collaborative training of models across multiple institutions without the need to exchange raw patient data [3]. Instead, model parameters are transferred to a central server for aggregation, maintaining data locality and safeguarding patient confidentiality. This approach benefits from the collective knowledge of geographically distributed datasets while reducing privacy risks.

In the context of breast cancer diagnostics, FL offers the potential to match or even exceed the performance of centralized models, especially when trained on heterogeneous datasets from diverse institutions. Kaissis et al. [4] demonstrated notable improvements in cross-institutional generalization for tumor segmentation using FL. Similarly, Wu et al. [5] achieved near-centralized baseline AUC scores for mammogram classification without compromising data privacy. Beyond privacy, FL's ability to incorporate demographically and geographically diverse datasets helps mitigate model bias and overfitting issues common in homogeneous, locally trained datasets [6].

A typical FL architecture for breast cancer detection consists of multiple local client nodes (e.g., hospitals or research labs) that train models on-site using their private datasets. Clients then send encrypted model updates to a central server, where techniques such as Federated Averaging (FedAvg) or Federated Optimization (FedOpt) are applied to produce a global model [7]. This model is redistributed to clients for further refinement in iterative training cycles until convergence.

Privacy-preserving enhancements such as Differential Privacy (DP), Homomorphic Encryption (HE), and Secure Multi-Party Computation (SMPC) have been integrated into FL frameworks [8]. These methods prevent data reconstruction from gradients, resist adversarial attacks, and ensure robustness against malicious updates through Byzantine fault tolerance. In medical imaging, particularly mammography and histopathology, FL can be implemented with 2D or 3D convolutional architectures, with local preprocessing steps like normalization, resizing, and augmentation. ResNet, DenseNet, and EfficientNet are common backbone architectures, and hybrid strategies combining FL with transfer learning have been shown to improve efficiency, especially in resource-constrained settings [9].

While FL offers significant benefits, real-world adoption in healthcare faces challenges including non-independent and identically distributed (non-IID) data, system heterogeneity, communication latency, and synchronization issues. Techniques like asynchronous updates, adaptive learning rates, and client sampling have been proposed to address these limitations [10].

In this work, we propose a novel FL-based architecture tailored for breast cancer detection using multi-institutional datasets. Our system integrates advanced privacy-preserving mechanisms and federated optimization algorithms to achieve high accuracy while safeguarding patient data. Using public datasets such as the UCI Breast Cancer Wisconsin dataset and Kaggle mammography images, we evaluate our model's performance using standard metrics including accuracy, precision, recall, and AUC. By maintaining data sovereignty and enabling collaborative intelligence, this framework lays the groundwork for scalable, secure, and ethical AI-based diagnostic systems.

II. STATE OF THE ART

Research on Federated Learning in healthcare has increasingly focused on balancing privacy, security, and computational efficiency. Early studies [11] revealed that FL, while avoiding direct data sharing, could still leak sensitive information through shared model updates. This vulnerability spurred the development of privacy-preserving strategies aimed at strengthening FL's security guarantees.

The appeal of FL in healthcare lies in its ability to enable collaborative model training without centralizing sensitive patient data. However, decentralization introduces new risks such as privacy leakage, susceptibility to data poisoning, model inversion attacks, and increased communication overhead [12–16]. Additional challenges

include handling non-IID data distributions, ensuring system scalability, and maintaining regulatory compliance. Addressing these issues is essential for deploying FL in medical applications.

Although FL avoids transmitting raw data, model parameters sent to central servers can still reveal private information. Differential Privacy (DP) provides a formal framework to quantify and control privacy loss by injecting controlled noise into model updates, limiting the risk of reconstructing sensitive data [21]. In medical AI applications, DP effectively safeguards privacy but can also degrade model accuracy, creating a trade-off that must be managed carefully. DP has been applied in various FL settings, including Vertical FL (VFL), Horizontal FL (HFL), and Transfer FL (TFL), using centralized, local (LDP), and shuffle-based implementations [22].

While DP is widely used, its limitations have led researchers to explore cryptographic techniques such as Secure Multi-Party Computation (SMPC) and Homomorphic Encryption (HE). SMPC enables multiple parties to jointly train models without exposing their local data, whereas HE supports direct computation on encrypted data [23]. Despite their strong privacy assurances, these methods often incur substantial computational overhead, limiting their practicality in real-time medical workflows. In comparative studies, DP has emerged as the more practical choice for balancing privacy protection with computational feasibility, though HE and SMPC remain viable for specific use cases.

Applications of DP in healthcare have extended to genomic analysis, wearable sensor data, and multi-modal medical datasets [24]. While DP provides significant privacy protection, it does not eliminate all vulnerabilities—FL models remain susceptible to adversarial manipulation, malicious updates, and data integrity threats [25]. Therefore, integrating DP with additional safeguards such as Byzantine fault tolerance and anomaly detection is essential for maintaining both accuracy and trustworthiness.

In summary, existing literature underscores FL’s potential for privacy-preserving, collaborative AI in healthcare, but also highlights its limitations and security vulnerabilities. Our proposed work builds upon these foundations by combining FL with advanced DP mechanisms to enhance privacy while ensuring high diagnostic accuracy in breast cancer detection. This approach addresses the dual objectives of regulatory compliance and clinical performance, positioning FL as a practical, secure, and scalable solution for AI-driven medical diagnostics.

III. METHODOLOGY

The proposed system is a Decentralized Federated Learning (DFL) framework for breast cancer detection, leveraging textual medical records and pathology reports, is depicted in Figure 1. The key innovation lies in enabling privacy-preserving, collaborative learning across institutions while using Natural Language Processing (NLP) techniques to model free-text clinical narratives.

A. Data Collection and Preprocessing

The raw data consists of unstructured text such as radiology reports, diagnostic summaries, treatment plans, and pathology annotations, collected independently by each institution. Preprocessing involves cleaning (lowercasing, punctuation removal, and stop-word filtering), tokenisation into tokens or sub-word units for transformer models, named entity recognition (NER) to identify and remove protected health information (PHI), and embedding, where tokens are converted into dense numerical vectors using Eq. (1).

$$TF - IDF: tf_{i,j}idf_{i,j} = tf_{i,j} \cdot \log\left(\frac{N}{df_i}\right) \quad (1)$$

- Word2Vec: Learn word vectors via skip-gram or CBOW models.
- Transformers (e.g., BERT): Contextualised embeddings using self-attention mechanisms.

B. Model Initialization

A global base model is initialized and distributed to local clients, defining the architecture, parameters, and data format for collaborative training. Given the unstructured textual nature of medical records, the architecture must capture semantic, contextual, and syntactic patterns.

The proposed work incorporates Recurrent Neural Networks (RNNs) to model sequence dependencies in clinical narratives, preserving temporal context essential for medical text. An RNN processes an input sequence $[x_1, x_2, \dots, x_T]$ with hidden state h_t at each step, as defined in Eq. (2).

$$h_t = \sigma(W_h h_{t-1} + W_x x_t + b) \quad (2)$$

Let the input sequence be denoted as in eqn (3):

$$X = [x_1, x_2, \dots, x_T] \quad (3)$$




Figure 1. Decentralised Federated Architecture of the Proposed System

The model maps this to a prediction $\hat{y} = f(X; \theta)$ where θ denotes model parameters. The initial parameters are broadcast to all participating nodes for federated training.

In FL, each client receives a globally consistent model architecture and initial parameters. For weight initialization, Random Initialization is adopted as it is suitable for training from scratch as given in (4).

$$W \sim U \left(-\sqrt{\frac{6}{n_{in} + n_{out}}}, \sqrt{\frac{6}{n_{in} + n_{out}}} \right)$$

Beyond weight initialization, federated learning performance depends on careful hyperparameter tuning. The learning rate (η) controls optimization step size—0.001 is common with Adam for steady convergence, while 0.01 suits SGD for faster starts. Batch sizes of 16–32 balance efficiency and generalization for textual medical data. Dropout (0.1–0.3) mitigates overfitting in small local datasets, and local epochs (E) determine how many passes over local data occur before server communication, with higher values boosting local learning but risking divergence in non-IID settings. Implementation typically uses PyTorch or TensorFlow, extended with frameworks like Flower, PySyft, or TensorFlow Federated for communication, secure aggregation, and synchronization. Tokenizer consistency is critical for transformer models, requiring uniform use across clients. Before training, clients perform weight synchronization checks to ensure alignment of structure, dimensions, and parameters with the global model, preventing aggregation errors and ensuring stable collaborative learning.

C. Decentralised Federated Learning Process

The decentralized federated learning process allows healthcare institutions to collaboratively train a breast cancer detection model without sharing sensitive patient data. Each institution trains locally on its textual medical records and pathology reports using algorithms like FedAvg, exchanging only model updates for secure aggregation into a global model. This preserves privacy, meets regulatory requirements, and improves accuracy by leveraging diverse datasets. Decentralization also enhances robustness by avoiding single points of failure.

Algorithm Selection

Each participating node retains its local dataset of unstructured clinical narratives, pathology reports, and diagnostic notes. A globally initialized model is distributed to all nodes, where local training is performed independently using optimizers such as SGD or Adam. Tables 1–2 outline the FedAvg process on server and client sides, while Tables 3–4 present the corresponding Federated Optimizer workflows.

TABLE I. ALGORITHM FOR SERVER-SIDE FEDAVG

Algorithm: Server side FedAvg	
Input : Number of clients K , total rounds T , local dataset sizes n_k for each client k	
Output: Final global model weights θ_T	
1	Initialize global model weights θ_0
2	for $t = 1$ to T do
3	Select a subset $S_t \subseteq \{1, \dots, K\}$ of participating clients
4	foreach <i>client</i> $k \in S_t$ in parallel do
5	Send current global weights θ_t to client k
6	Client k trains locally and returns updated weights θ_t^k
7	//Aggregate weights:
8	$n \leftarrow \sum_{k \in S_t} n_k$
9	$\theta_{t+1} \leftarrow \sum_{k \in S_t} \frac{n_k}{n} \theta_t^k$

TABLE II. ALGORITHM FOR CLIENT-SIDE FEDAVG

Algorithm: Client-Side Fed Avg	
Input : Global weights θ_t from server, local dataset D_k , learning rate η , number of epochs E	
Output: Updated local weights θ	
Receive global weights θ_t from server	
Set local model $\theta \leftarrow \theta_t$	
for <i>epoch</i> = 1 to E do	
foreach <i>mini-batch</i> b in D_k do	
Compute gradient: $\nabla \mathcal{L}(\theta; b)$	
Update weights: $\theta \leftarrow \theta - \eta \cdot \nabla \mathcal{L}(\theta; b)$	
Return updated weights θ to server	

TABLE III. ALGORITHM FOR SERVER-SIDE FEDOPT

Algorithm: Server-Side Federated Optimization Procedure	
Input : Number of clients K , total rounds T , learning rate α , momentum coefficient β	
Output: Final global model weights θ_T	
Initialize global model weights θ_0	
Initialize server optimizer state	
for $t = 1$ to T do	
Select a subset $S_t \subseteq \{1, \dots, K\}$ of participating clients	
foreach <i>client</i> $k \in S_t$ in parallel do	
Send current global weights θ_t to client k	
//Receive local updates from each client:	
$\Delta \theta_k = \theta_k - \theta_t$	
//Compute average update:	
$\Delta \theta_{\text{avg}} = \frac{1}{ S_t } \sum_{k \in S_t} \Delta \theta_k$	
$v_t = \beta \cdot v_{t-1} + (1 - \beta) \cdot \Delta \theta_{\text{avg}}$	
$\theta_{t+1} = \theta_t + \alpha \cdot v_t$	

TABLE IV. ALGORITHM FOR CLIENT-SIDE FEDAVG

Algorithm: Client-Side Federated Optimization Procedure	
Input : Global weights θ_t from server, local dataset D_k , learning rate η , number of epochs E	
Output: Local update $\Delta \theta_k$	
Receive global weights θ_t from server	
Set local weights $\theta \leftarrow \theta_t$	
for <i>epoch</i> = 1 to E do	
foreach <i>mini-batch</i> b in D_k do	
Compute gradient: $\nabla \mathcal{L}(\theta; b)$	
Update weights: $\theta \leftarrow \theta - \eta \cdot \nabla \mathcal{L}(\theta; b)$	
Send local update $\Delta \theta_k = \theta - \theta_t$ back to server	

D. Local model training and Updation

The training objective at each client is to minimise a binary cross-entropy loss function given by as in eqn (5):

$$\mathcal{L}(y, \hat{y}) = -y \log(\hat{y}) - (1 - y) \log(1 - \hat{y}) \quad (5)$$

Gradients are computed and applied to update the parameters as in eqn (6):

$$\theta_t^k = \theta_t^k - \eta \nabla \mathcal{L} \quad (6)$$

where η is the learning rate. Clients train for several epochs before sending updates. Upon completion of local training, clients compute parameter updates $\Delta\theta^k = \theta^k - \theta$, which are exchanged with peers or a decentralised aggregation coordinator. Model updates are aggregated using the Federated Averaging (FedAvg) algorithm as in eqn (7):

$$\theta_{t+1} = \sum_{k=1}^K \frac{n_k}{n} \theta_t^k \quad (7)$$

where n_k denotes the number of training samples at client k , and $n = \sum_{k=1}^K n_k$ is the total sample size across all clients. Unlike traditional centralised federated learning, DFL eliminates the reliance on a single aggregation server, thereby reducing the risk of bottlenecks and single points of failure. This decentralised protocol not only preserves patient privacy but also enhances model robustness and generalisation across heterogeneous data distributions.

E. Secured model aggregation

To protect client updates and enhance fault tolerance, advanced aggregation strategies are incorporated into the proposed system adds Gaussian noise as in eqn (8) to protect against model inversion attacks.

$$\tilde{g} = g + \mathcal{N}(0, \sigma^2 I) \quad (8)$$

IV. RESULTS AND DISCUSSION

The proposed method was evaluated on public healthcare datasets for breast cancer detection using both text and mammogram data. Figure 2 compares classification accuracy across three FL strategies for text-based data. FedAvg achieved 89% accuracy as a solid baseline, FedOpt improved to 92.5% through adaptive server-side optimization, and the hybrid Layered FedAvg + FedOpt reached 94%, combining FedAvg's stability with FedOpt's faster convergence. These results demonstrate the advantage of hybrid aggregation in improving model performance for decentralized medical diagnosis.




Figure 2. Comparison of accuracy for text data over the proposed models




Figure 3. Loss curve for the proposed models

The loss curve shows training loss over communication rounds for three FL strategies. FedAvg and FedOpt converge steadily, while Layered FedAvg + FedOpt achieves the fastest, most stable loss reduction, indicating superior optimization for breast cancer detection.

For mammogram data (Figure 4), accuracy was compared across FL with averaging, FL with optimization, and the layered approach. After adding Gaussian noise, FL with optimization performed best at 90.1%, surpassing the other two strategies.

TABLE VI. CORRELATION ANALYSIS FOR BREAST CANCER DETECTION USING FED ALGORITHMS

	Lump	Pain	Density	Malignant	Benign
Lump	1.00	0.55	0.72	0.80	-0.45
Pain	0.55	1.00	0.48	0.30	-0.10
Density	0.72	0.48	1.00	0.66	-0.25
Malignant	0.80	0.30	0.66	1.00	-0.88
Benign	-0.45	-0.10	-0.25	-0.88	1.00

The generated heatmap, as depicted in Figure 5, visualises correlations among clinical features from textual medical records. A strong positive correlation is observed between “lump” and “malignant,” while “benign” and “malignant” exhibit a strong negative correlation. These insights highlight key diagnostic indicators, aiding in feature selection for breast cancer detection models.




Figure 4. Comparison of Accuracy for Mammogram data over proposed models




Figure 5. Heat map for cancer detection using Fed Algorithms

V. CONCLUSION

This study proposes a robust Decentralized Federated Learning (DFL) framework for breast cancer detection using textual medical records and pathology reports. Unlike centralized ML or conventional federated setups, the architecture enhances privacy, fault tolerance, and scalability by enabling local training and secure aggregation without sharing patient data, meeting HIPAA and GDPR standards. Experiments show notable gains in accuracy, precision, recall, and AUC compared to centralized methods, with strong performance on non-IID, heterogeneous data typical of clinical settings. DFL’s decentralized design addresses data silos and privacy challenges, offering a scalable solution for secure, collaborative AI in healthcare. Future work will focus on real-time deployment, interpretability, and multi-modal data integration.

REFERENCES

- [1] World Health Organization. (2021). Breast cancer. <https://www.who.int/news-room/fact-sheets/detail/breast-cancer>
- [2] Yala, A., Lehman, C., Schuster, T., Portnoi, T., & Barzilay, R. (2019). A deep learning model using mammography data to enhance breast cancer risk assessment. *Radiology*, 292(1), 60–66.
- [3] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient deep learning on decentralized datasets. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [4] Kaissis, G., Ziller, A., Passerat-Palmbach, J., Ryffel, T., Usynin, D., Trask, A., ... & Makowski, M. (2021). Secure deep learning for medical imaging across institutions. *Nature Machine Intelligence*, 3(6), 473–484.
- [5] Wu, J., Zhang, Y., Guo, X., & Lu, W. (2022). Evaluating federated learning for classifying mammography images: A performance comparison. *IEEE Transactions on Medical Imaging*, 41(5), 1154–1164.
- [6] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Key challenges and solutions in federated learning. *IEEE Signal Processing Magazine*, 37(3), 50–60.
- [7] Kairouz, P., McMahan, H. B., et al. (2021). Recent advances and open challenges in federated learning. *Foundations and Trends in Machine Learning*, 14(1), 1–210.

- [8] Geyer, R. C., Klein, T., & Nabi, M. (2017). Federated learning with client-level privacy: An initial perspective. In *NeurIPS Workshop on Private Multi-Party Machine Learning*.
- [9] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., & Bakas, S. (2020). Enabling collaboration in medicine via federated learning without data sharing. *Scientific Reports*, 10(1), 12598.
- [10] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ... & Ramage, D. (2019). System architecture for scalable federated learning. In *Proceedings of MLSys*.
- [11] Xie, Q., et al. (2024). A brief review of efficiency optimization techniques in encrypted federated learning. *IEEE Internet of Things Journal*, 11, 24569–24580.
- [12] Xu, J., et al. (2019). Review of federated learning applications in health informatics. *Journal of Healthcare Informatics Research*, 5, 1–19.
- [13] Zheng, X., et al. (2023). Addressing non-IID data in federated learning through local and global distillation. In *2023 IEEE International Conference on Web Services (ICWS)*, 647–657.
- [14] Fu, J., et al. (2024). A systematic review of differentially private federated learning approaches. *arXiv preprint, arXiv:2405.08299*.
- [15] Heusinger, M., Raab, C., Rossi, F., & Schleif, F.-M. (2022). Surveying federated learning: Methodologies, applications, and outlook. *arXiv preprint, arXiv:2212.11729*.
- [16] McGraw, D., & Mandl, K. D. (2021). Data protection strategies for promoting digital health systems. *NPJ Digital Medicine*, 4, Article 2.
- [17] Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive classification and survey on privacy-preserving federated learning. *ACM Computing Surveys*, 54(1), 1–36.
- [18] Liu, Z., et al. (2022). Aggregation techniques for privacy in federated learning: A survey. *IEEE Transactions on Big Data*. <https://doi.org/10.1109/TBDATA.2022.3190835>
- [19] Rahaman, M., Arya, V., Orozco, S. M., & Pappachan, P. (2024). Secure multi-party computation protocols for healthcare data privacy. In *Innovations in Modern Cryptography* (pp. 190–214). IGI Global.
- [20] Xie, Q., et al. (2024). Enhancing performance in privacy-preserving federated learning via encryption methods. *IEEE Internet of Things Journal*, 11, 24569–24580.
- [21] Adnan, M., Kalra, S., Cresswell, J. C., et al. (2022). Medical image analysis using federated learning with differential privacy. *Scientific Reports*, 12, Article 1953. <https://doi.org/10.1038/s41598-022-05539-7>
- [22] Fu, J., Hong, Y., Ling, X., Wang, L., Ran, X., Sun, Z., Wang, W. H., Chen, Z., & Cao, Y. (2024). A comprehensive study of differential privacy in federated learning. *arXiv preprint, arXiv:2405.08299*. <https://doi.org/10.48550/arXiv.2405.08299>
- [23] Abubaker, A. A. M., & Mohammed, A. O. A. (2025). Technologies for preserving privacy: Homomorphic encryption and SMPC. *International Journal for Research in Applied Science and Engineering Technology*, 13(IV), 36.
- [24] Liu, W., Zhang, Y., Yang, H., et al. (2024). A review on differential privacy in medical data analytics. *Annals of Data Science*, 11, 733–747. <https://doi.org/10.1007/s40745-023-00475-3>
- [25] Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Applications of federated learning in healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19. <https://doi.org/10.1007/s41666-020-00082-4>