EXPLORING MACHINE LEARNING FOR BREAST CANCER DETECTION IN X-RAY IMAGING

[1]WEITONG LIAO, [2]LIGANG HE, [3]MOHAMMED H. ALGHAMDI, [4]HAMMAM M. ALGHAMDI

[1,2] Department of Computer Science, University of Warwick
[3,4] College of Computer Science and Engineering, University of Jeddah
[3]College of Computer Science, King Khalid University
E-mail: [1]weitong.liao@warwick.ac.uk, [2]ligang.he@warwick.ac.uk, [3]mhalghamdi@uj.edu.sa, [4]hmsalghamdi@uj.edu.sa

Abstract - Breast X-ray imaging is a primary method for early breast cancer screening. However, applying machine learning techniques to this task presents significant challenges, particularly in capturing fine-grained details from high-resolution medical images. This study explores multiple machine learning approaches for breast cancer detection, including image-level and patch-level CNN-based classification, diffusion model-based anomaly detection, and federated learning (FL) for training models across distributed datasets. We evaluate these methods on real-world breast cancer X-ray datasets. Our findings suggest that FL offers a promising balance between performance and privacy preservation, underscoring its potential for real-world medical AI applications.

Keywords - Breast Cancer Detection, Convolutional Neural networks (CNN), Diffusion Models, Federated Learning (FL), Medical Imaging.

I. INTRODUCTION

According to the GLOBOCAN 2020 study, an estimated 2.3 million new cases of female breast cancer were reported in 2020, accounting for 11.7% of all cancer cases, surpassing lung cancer as the most common cancer in women [1]. The commonly used medical imaging techniques to detect breast cancer in clinical practice include X-ray imaging, magnetic resonance imaging (MRI), nuclear medicine imaging (NMI), and ultrasound imaging (UI)[2]. Due to its low cost and ease of detection, breast X-ray imaging has an advantage over other diagnostic methods for early breast cancer detection [3], [4]. However, it also has limitations, such as the inability to accurately detect cancer in patients with high breast density, as dense breast tissue often appears similar to cancerous tissue in imaging [3]. This highlights the need for computer-aided systems or machine learning models to assist doctors in making more accurate diagnoses.

This paper explores the feasibility and potential of applying different kinds of machine learning algorithm to breast cancer detection using X-ray images. Given the unique characteristics of different dataset, such as the imbalance in class distributions and the heterogeneity in data sources, we selected appropriate models and training approaches to optimize the performance of breast cancer detection tasks. By constructing real-world scenarios with different datasets across various institutions, our federated learning model achieved performance comparable to local training, even under conditions of data imbalance, non-IID distributions, and institutional data heterogeneity.

II. RELATED WORK

A. Mammography

Mammography is also referred to as breast X-ray imaging. In traditional clinical scenario, doctors manually assess X-ray images of patients, with the accuracy of diagnosis influenced by various factors such as the doctor's experience, image quality, and the patient's age. This diagnostic approach increases the risk of missed detections or misdiagnoses, potentially delaying early treatment or causing unnecessary burden on patients[5]. Furthermore, as patients' imaging data is considered private and not shared through medical institutions, this limits further research based on patient clinical data. Existing representative breast X-ray datasets such as CBIS-DDSM, CMMD, and VinDr-Mammo provide limited samples, with 2,620, 1,775, and 5,000 patients, respectively[5], [6], [7]. The limited sample size poses a significant challenge for training machine learning models, as it may lead to overfitting and limit the generalization ability of the models.

Clinical data and the imaging reports are commonly annotated according to the standards set by the American College of Radiology (ACR) [8], using the BI-RADS (Breast Imaging-Reporting and Data System) categories. These categories are as follows: BI-RADS 0 (incomplete, additional imaging evaluation required), BI-RADS 1(normal), BI-RADS 2 (typically benign), BI-RADS 3 (probably benign, follow-up recommended), BI-RADS 4 (suspicious), BI-RADS 5 (Highly suggestive of malignancy), BI-RADS 6 (known biopsy-proven) [5], [8]. The standards also define specific mammographic features, which are described in the BI-RADS lexicon,

including densities and masses, micro or macrocalcifications, architectural distortions, special cases (such as ductal ectasia, intramammary lymph nodes, or focal asymmetric density), and associated findings (including skin or nipple retraction, skin thickening, cutaneous lesions, and axillary lymph nodes) [8]. However, the aforementioned datasets provide limited descriptions of each patient. For example, the CBIS-DDSM dataset only records benign (BI-RADS 2) and malignant cases (above BI-RADS 3), with each case's abnormality type classified as either a mass or a calcification.

B. Medical Image Analysis

Machine learning methods are widely applied in the field of medical imaging analysis, covering various tasks such as classification [9], [10], [11], [12], anomaly detection [13], [14], segmentation [15], [16], [17], and annotation[16].

Most research in the field of medical image analysis focuses on tissue-level analysis. For example, [9] uses datasets such as CAMELYON-16[18] and TCGA-BRCA [19], which consists of tissue slices. Due to the high resolution of tissue images, entire tissue slides are typically divided into smaller patches and then fed into a pre-trained CLIP model [20] to generate embeddings, which helps save computational resources during training. The downstream models can utilize these dimensionality-reduced embeddings for tasks such as prediction, segmentation, and others. However, for breast X-ray images, which are generally smaller in size and require consideration of the entire image for accurate diagnosis, dividing them into patches may lead to the loss of valuable information.

In the area of breast X-ray image analysis, [11] focused on the CBIS-DDSM breast X-ray dataset, using complete images and applying Convolutional Neural Network (CNN) for patient classification. [10] proposed case-level breast cancer prediction by combining the patient's craniocaudal view (CC view) and mediolateral oblique view (MLO view) X-ray images for more accurate prediction. [12] applied multiple machine learning models on the CBIS-DDSM dataset, providing a transparent baseline for subsequent research.

[13] and [14] approach the diagnostic task as an anomaly detection problem, where the data distribution of normal cases differs from that of abnormal cases. Machine learning models such as Variational Autoencoders models (VAE) and Diffusion models are trained to capture these differences, enabling the identification of anomalies or the reconstruction of abnormal regions. [13] applied a Diffusion Model on the BRATS2020 brain tumor MRI dataset to reconstruct abnormal regions into normal tissue, enabling precise segmentation of tumor areas and demonstrating the potential of diffusion models in medical imaging. However, for breast X-ray images, abnormal features are often less distinct than brain tumors and can be small in size, such as tiny calcifications. Therefore, the effectiveness of this approach in breast cancer detection requires further investigation.

C. Medical Applications of Federated Learning

In real-world clinical scenarios, the issue of data silos is significant, as medical institutions are restricted from sharing data due to privacy regulations [21]. Additionally, data heterogeneity is prevalent—benign cases naturally outnumber malignant cases, leading to highly imbalanced labels. Moreover, differences in data collection and annotation protocols across institutions further complicate the application of machine learning [22]. As a result, individual institutions often struggle to train high-performance models due to limited and biased data availability. Federated learning offers a solution to these challenges by enabling each medical institution to train local models on their private data without sharing it [23]. After each training round, a central server aggregates the locally trained models to produce a more robust global model [23]. This approach allows for effective utilization of distributed data while preserving patient privacy and addressing issues related to data silos and heterogeneity. [24] proposed an Inside and Outside model Personalization Federated Learning (IOP-FL) framework, which enhances personalized federated learning by distinguishing between global and local gradients achieved performance improvements of 2.63% and 2.79%, respectively, compared to FedAvg [23].

Most existing studies artificially partition the same dataset across different clients to evaluate federated learning performance. However, in real-world scenarios, such an approach fails to fully capture the heterogeneity of medical data. Research on federated learning with truly heterogeneous datasets distributed across different clients remains limited, particularly in the domain of breast cancer detection.

III. METHODOLOGY

This study applies various machine learning approaches, including both image-level and patch-level CNN-based classification, diffusion model-based anomaly detection and federated learning, to conduct breast cancer detection using real-world datasets. We used two publicly available datasets, VinDr-Mammo [5] and CMMD [7], both annotated by multiple expert radiologists. A comprehensive analysis was conducted using these machine learning approaches.

Additionally, due to performance constraints, the input to neural networks is often resized, which results in a

loss of fine details compared to the original images. To investigate whether the approach of partitioning images into smaller patches, commonly used in tissue analysis, is applicable to breast cancer detection, this study leverages pre-annotated suspicious regions from the CBIS-DDSM dataset [6]. These regions are considered as small patches and used for classification tasks with machine learning models. Then, to simulate a real-world scenario, we distributed the VinDr-Mammo dataset and the CMMD dataset across separate clients and evaluated the performance of the aggregated global model after federated training.

A. Convolutional Neural Network Based Classification Task

1) Whole Image-Level Classification

For image classification tasks, the most straightforward approach is to feed the entire image into a neural network, minimizing dataset requirements. Non-end-to-end methods, however, require manual annotation of regions of interest (ROIs), which is time-consuming, costly, and dependent on doctors' expertise[11]. Additionally, the limited availability of annotated datasets and the high cost of annotations make transferring models to other datasets more challenging [11].

Several CNN models are explored by this research. The Vision Transformer (ViT) [25], which incorporates an attention mechanism [26], is one of the options. ViT splits the entire image into patches, generates features through an embedding layer, then processes them through transformer encoder layers with attention mechanisms, ultimately generating predictions [25]. However, as shown in [25] and supported by our experiments, although ViT performs better on large datasets due to its complex multi-head attention mechanism and a greater number of parameters, it lacks the inherent inductive bias of CNNs, making it less effective on smaller datasets [25]. Therefore, ViT is not suitable for this study. Instead, a ResNet [27] variant with residual connection mechanism, ResNet-101, is chosen as the model for the classification task.

For each sample in the dataset, two views of the same side breast for each patient are captured, namely the MLO view and the CC view. To maximize the use of the dataset, both views are treated as separate samples and fed into the neural network during training. Due to its unique angle, the MLO view provides a comprehensive view of both the breast tissue and the underlying muscle tissue, often demonstrating a stronger correlation with the patient's actual condition compared to the CC views [5], [28]. Therefore, for evaluation, we combine the two views to leverage the additional information provided by both perspectives. Let $p_{CC}^{(1)}$ and $p_{CC}^{(0)}$ represent the predicted probabilities for class 1 (malignant) and class 0 (benign) from the Craniocaudal (CC) view, respectively. Similarly, let $p_{MLO}^{(1)}$ and $p_{MLO}^{(0)}$ represent the predicted probabilities for class 1 (malignant) and class 0 (benign) from the Mediolateral Oblique (MLO) view. To combine the predictions for class 1 and class 0, we use a weighted average with a factor α, where $0 \leq \alpha \leq 1$:

For class 1 (malignant):
$p_{final}^{(1)} = \alpha \cdot p_{CC}^{(1)} + (1-\alpha) \cdot p_{MLO}^{(1)}$.
For class 0 (benign):
$p_{final}^{(0)} = \alpha \cdot p_{CC}^{(0)} + (1-\alpha) \cdot p_{MLO}^{(0)}$.

2) Patch-Level Classification

In the tissue classification task, an image is divided into multiple patches. The image-level label is determined by the labels of its patches under the following condition. Let an image I be composed of N patches, where each patch i has an individual classification label y_i, where $y_i \in \{0,1\}$(0 for negative, 1 for positive). The image-level classification label Y is given by:

$$Y = \max_{i \in \{1,2,\ldots,N\}} y_i$$

where: $y_i = 1$ if patch i is classified as positive, $y_i = 0$ if patch i is classified as negative [29]. Thus, the image is classified as positive ($Y = 1$) if at least one patch is positive, otherwise, it is classified as negative ($Y = 0$) [29].

To explore the feasibility of this approach for breast X-ray images, we applied CNNs at the patch level to examine whether the model could distinguish between benign and malignant tissue without access to global contextual information. To achieve this, we utilized the CBIS-DDSM dataset [6], which includes expert-annotated ROI regions for each case. Each ROI is labeled with pathology (benign or malignant) and abnormality type (calcification or mass).

B. Diffusion Model Based Anomaly Detection Task

Diffusion models have shown strong performance in brain tumor MRI tasks [13]. Since diffusion models inherently model data distributions, they learn the distribution of normal samples and can generate normal samples from random noise. This suggests their potential for transforming unseen abnormal samples into normal ones, enabling anomaly detection by comparing differences. However, as generative models, their reliability and stability in medical applications remain a concern [30].

In this research, the Denoising Diffusion Probabilistic Model (DDPM) proposed by [31] is applied to the breast cancer detection domain, and its performance is evaluated. The diffusion process in this study follows the DDPM formulation, where Gaussian noise is progressively added to the sample using the forward process:

$$q(x_t|x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{\alpha_t}x_{t-1}, (1-\alpha_t)I\right) \quad (1)$$

where t is a hyperparameter chosen from the range$[0, T]$, controlling the noise level. $x_t$is the image at time step t, obtained by progressively adding noise to the original image over t steps. The model learns to denoise by estimating the noise at a given timestep t, following the reverse process:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}\left(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)\right) \quad (2)$$

where θ represents the parameters of the noise prediction model (U-Net). The U-Net takes the noisy image x_t and the current time step t as input, and generates a predicted noise $\epsilon_\theta(x_t, t)$[31]. The goal is for the model to predict the noise such that subtracting it from the noisy image $x_t$results in a cleaner image. The diffusion loss is computed as defined in DDPM:

$$L(\theta) = E_{t,x_0,\epsilon}\left[|\epsilon - \epsilon_\theta\left(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, t\right)|^2\right]. \quad (3)$$

This loss function compares the predicted noise ϵ_θ with the actual noise ϵ that was added during the forward process.

In the evaluation phase, a hyperparameter N is set to control the number of noise steps added to the test samples. This helps transform the potential malignant case distribution into a neutral distribution. The model, trained on benign cases, then uses the reverse process to restore the image regions. The complete process can be seen on Fig. 1.

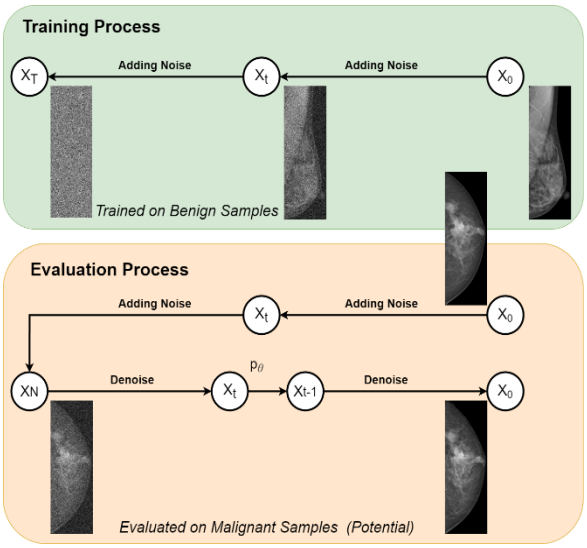


Fig. 1. Overview of the diffusion process for anomaly detection in breastcancer imaging

C. Federated Learning for Breast Cancer Detection

Finally, this research examines the use of federated learning to train models on data collected from multiple institutions. The study incorporates two datasets, VinDr-Mammo and CMMD, each assigned to a separate client. Each client only trains a local model using its respective dataset. The specific FL process is outlined in Algorithm 1.

At the beginning of the FL process, a global model is initialized and distributed to clients as the local model for round 0. As discussed in SectionIII-A1, we also selected ResNet-101 as the model used by each client. During each communication round in the FL process, the server aggregates the local models from the clients according to the aggregation algorithm and sends the updated model back to all clients for the next round of training. Additionally, the VinDr-Mammo and CMMD datasets exhibit differences in label distributions, which is to be explained in SectionIV-A. To address the imbalance in data distribution, FedProx[32] is used to optimize the training process. FedProx incorporates a proximal term into each client's local updates to mitigate the impact of data distribution discrepancies [32].

Algorithm 1 Federated Learning for Breast Cancer Detection

Require: Initial global model θ_0, number of communication rounds T

Require: Clients A and B with datasets D_A and D_B

for each round $t = 1, 2, \ldots, T$ **do**
  for each client $i \in \{A, B\}$ **do**
    Sample dataset D_i from client i
    Train local model $\theta_i^{(t)}$ by solving:

$$\theta_i^{(t)} = \arg\min_{\theta_i}\left(\mathcal{L}_i(\theta_i) + \frac{\mu}{2}\|\theta_i - \theta_{t-1}\|^2\right)$$

    Send local model update $\theta_i^{(t)}$ to the server
  end for
  Aggregate local updates at the server:

$$\theta_t = Agg(\theta_A^{(t)}, \theta_B^{(t)})$$

  Update global model $\theta_t \leftarrow \theta_t$
end for

IV. EXPERIMENTS, RESULTS AND DISCUSSIONS

A. Dataset

1) VinDr-Mammo Dataset

The VinDr-Mammo dataset is a publicly available large-scale dataset from Vietnam, which comprises X-ray images of both side breasts (MLO and CC views) of 5,000 patients, totaling 10,000 samples [5]. Three radiologists evaluated each image to determine the level of BI-RADS, which ranges from 1 to 5 [5]. While the dataset is relatively large compared to others, it still suffers from an imbalanced label distribution. The label distribution is shown in Fig. 2. It can be seen that limited amount malignant samples are available in the dataset.

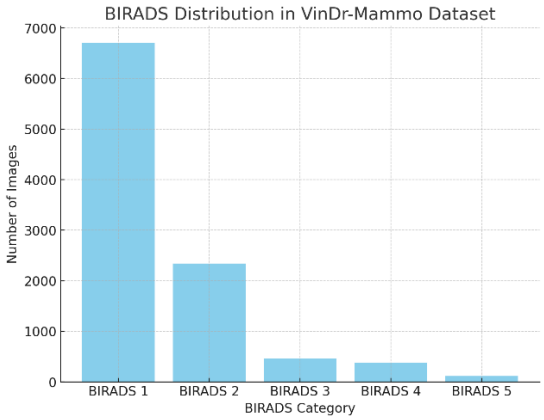


Fig. 2. The data distribution of VinDr-Mammo dataset.

For different tasks, we used varying data selection strategies. For the classification task, since the CMMD dataset only includes benign and malignant labels, we aligned the training labels with the CMMD dataset. Specifically, BI-RADS 2 was selected as benign, and BI-RADS 4 and 5 were selected as malignant. BI-RADS 3 was typically excluded in the classification task, as it is often considered unreliable, based on the guidelines provided by the VinDr-Mammo dataset [5]. Therefore, for the classification task using the VinDr-Mammo dataset, there are 494 positive samples and 2,338 negative samples. For the anomaly detection task, the model is only trained using normal samples as input, resulting in 5,362 training samples (BI-RADS 1 only) and 494 abnormal test samples (BI-RADS 4 and BI-RADS 5).

To prevent precision loss, the raw data of VinDr-Mammo dataset, which has a 16-bit depth, is loaded with the same 16-bit depth. Additionally, each image is processed using computer graphics techniques to crop the blank areas outside the tissue regions.

2) CMMD Dataset

The Chinese Mammography Database (CMMD) is collected from the Sun Yat-sen University Cancer Center in Guangzhou and the Nanhai Affiliated Hospital of Southern Medical University in Foshan [7]. For each subject, CC view and MLO view images are obtained and stored in 8-bit format [7]. The dataset includes 556 benign cases and 1,316 malignant cases. Before fed into training, each image undergoes computer graphics algorithms to remove blank areas.

3) CBIS-DDSM Dataset

The CBIS-DDSM dataset, which includes accurate region-of-interest (ROI) annotations by professional radiologists, is utilized in this study to analyze the feasibility of patch-level classification for breast cancer detection. In addition to labeling patients with suspicious anomaly as benign or malignant, the CBIS-DDSM dataset also provides annotations for abnormal region subtypes (see TableI). This allows patch-level classification to be conducted from multiple perspectives.

	Benign	**Malignant**	**Total**
Calcification	1199	673	1872
Mass	912	784	1696
Total	2111	1457	

TABLE I DATA DISTRIBUTION OF THE CBIS-DDSM DATASET.

B. Convolutional Neural Network Based Classification Task

1) Whole Image-Level Classification Results

In the experiment, we evaluated the performance of neural networks on two datasets locally, as shown in TableII. We then selected existing works in the field of breast cancer classification, namely [33] and [10], for comparison. Most of the current work primarily focuses on the AUC metric, and therefore, only the AUC values are reported in these studies. From the first two baselines in the table, it can be observed that although CNN achieves high accuracy on certain datasets, due to the imbalance in the dataset, the neural network struggles to differentiate the minority class effectively, as seen from the F1 score and AUC. The experimental results indicate that, under conditions of data imbalance, models trained by a single medical institution using only local private data may perform poorly. However, in a federated learning scenario, despite challenges such as data heterogeneity and non-IID [34] issues, the availability of a global model for training guidance has the potential to alleviate the data silo problem.

2) Patch-Level Classification Results

In the study of patch-level classification tasks on the CBIS-DDSM dataset, significant variations in model performance were observed depending on the task, as shown in TableIII. In the benign vs. malignant classification task, challenges arise due to the absence of global image context. When determining the patient's condition solely based on suspicious regions, the model may struggle. For instance, in cases of patients with high breast density, such as obese individuals, suspicious regions tend to exhibit similarly high density. This can lead the model to erroneously classify the patient as malignant based on the large density of the patch. This highlights certain limitations of using breast X-rays for breast cancer screening.

We also conducted a study on tissue subtypeclassification, where the model achieved an AUC of 0.92 which is much higher than previous task. This inspires that one potential direction for future research is to use a patch-level classifier to first distinguish tissue subtypes and then perform further classification tasks downstream. For example, when the patch-level classifier identifies regions with clear calcifications, which are relatively easier to be

classified, a more confident classifier could be applied. For regions with masses, a downstream classifier with higher redundancy could be utilized to achieve a lower false negative rate as dense tissues such as adipose may obstruct the X-ray [1], leading to missed diagnoses.

C. Diffusion Model Results

The diffusion model has shown good performance in other medical imaging domains. However, its performance in breast cancer anomaly detection tasks is not ideal. We present both quantitative and qualitative analyses to illustrate the experimental results. In the qualitative analysis, we selected test data to reconstruct the abnormal regions and visualized the reconstruction results using a difference heatmap. Ideally, the reconstructed abnormal regions should show significant differences compared to the original image. However, the experimental results did not reveal clear differences, with only some subtle variations observed at the contours of the abnormal regions, which can be seen in Fig. 3. When using the reconstruction difference as a feature for classification, the model achieved only a 0.55 F1 score, which is significantly lower than the baseline performance. A possible reason for this is that while being capable of generating high-quality syntheticbreast X-ray samples from random noise, the U-Net model used for noise prediction is not sufficient to capture the differences between the abnormal and normal distributions. Additionally, the task of reconstructing regional details is inherently challenging. Amore feasible approach could be to focus on reconstructing the mask of the abnormal regions [13].

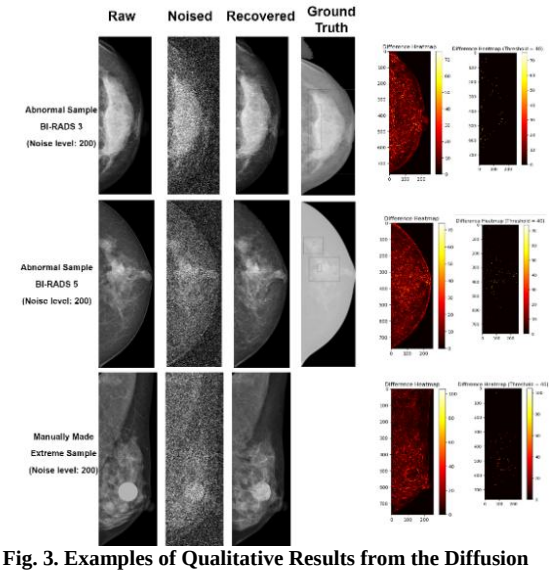


Fig. 3. Examples of Qualitative Results from the Diffusion Model.

Method	Dataset	Acc	F1	AUC
Kim et al. (2018)	CBIS	-	-	0.64 ± 0.00
ES-Attside	VinDr	-	-	0.83± 0.00
Baseline (ours)	VinDr	0.81± 0.03	0.65 ± 0.02	0.75 ± 0.03
Baseline (ours)	CMMD	0.68 ± 0.02	0.63 ± 0.01	0.77 ± 0.02
Baseline (ours)	VinDr+CMMD	0.79 ± 0.02	0.79± 0.01	0.85± 0.02
Federated Learning in Real Scenarios (ours)	VinDr+CMMD (separate client)	0.78 ± 0.02	0.77 ± 0.01	0.82 ± 0.02

TABLE II COMPARISON OF PERFORMANCE BETWEEN DIFFERENT METHODS AND DATASETS.

Task	Acc	F1	AUC
Benign vs Malignant	0.65 ± 0.03	0.60 ± 0.02	0.71 ± 0.01
Calcification vs Mass	0.83 ± 0.01	0.83 ± 0.01	0.92 ± 0.01

TABLE III PERFORMANCE OF DIFFERENT CLASSIFICATION TASKS ATTHE PATCHLEVEL.

D. Fedrated Learning for Real Scenario Breast Cancer Detection

The performance of FL in real-world scenarios, as shown in TableII, is evaluated by simulating FL execution with different datasets to mimic real-world conditions. The results indicate that when data distribution is imbalanced within a single client, the aggregated global model performs similarly to local models in breast cancer classification tasks. This suggests that FL can effectively mitigate the challenges of data non-sharing across institutions in medical image recognition.

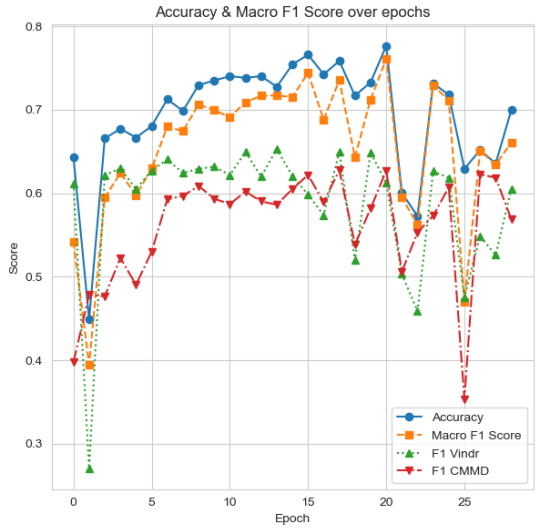


Fig. 4. Overfitting problem is observed through training process.

However, despite efforts to simulate inter-institutional heterogeneity, the study has limitations, primarily due to the small number of available datasets, which constrains baseline model performance. As shown in Fig. 4, after a certain number of training epochs, the model's performance begins to degrade while the training loss continues to decrease, indicating severe overfitting due to the limited dataset size. This highlights the critical need for larger and more diverse datasets in medical image analysis.

V. CONCLUSION

This study applies multiple machine learning approaches for breast cancer detection using multiple medical image datasets. Our experiments demonstrated that FL can improve detection performance while addressing data privacy concerns and non-IID distributions across institutions. By aggregating local models from different clients, FL enabled the creation of a global model that achieved higher performance. Nonetheless, several challenges remain in breast cancer detection. Since X-rays serve primarily as an early screening tool, suspicious tissues often lack clear signs of abnormality. Additionally, crucial diagnostic factors, such as patient age, are considered private information and are unavailable in research datasets, further impacting model performance. Despite these limitations, FL has shown significant promise in breast X-ray imaging. Given that breast X-rays are a cost-effective and widely accessible screening method [4], vast amounts of potential data exist across medical institutions. FL offers a promising solution to harness this distributed data, facilitating the development of more accurate and efficient diagnostic models for breast cancer detection.

ACKNOWLEDGMENT

This work is partly supported by the British Council under the UK-Saudi Challenge Fund.

REFERENCES

[1] A. Karellas and S. Vedantham, “Breast cancer imaging: a perspectivefor the nextdecade,”Medical physics, vol. 35, no. 11, pp. 4878–4897,2008.

[2] Y.-n. Zhang, K.-R. Xia, C.-Y. Li, B.-L. Wei, and B. Zhang, “Reviewof breast cancer pathologigcal image processing,”BioMed researchinternational, vol. 2021, no. 1, p.1994764, 2021.

[3] M. D. Fishman and M. M. Rehani, “Monochromatic x-rays: The futureof breast imaging,”European Journal of Radiology, vol. 144, p. 109961,2021.

[4] C. Coleman, “Early detection and screening for breast cancer,” inSeminars in oncology nursing, vol. 33, no. 2. Elsevier, 2017, pp. 141–155.

[5] H. T. Nguyen, H. Q. Nguyen, H. H. Pham, K. Lam, L. T. Le, M. Dao, andV. Vu, “Vindr-mammo: A large-scale benchmark dataset for computeraideddiagnosis in full-field digital mammography,”Scientific Data,vol. 10, no. 1, p. 277, 2023.

[6] R. S. Lee, F. Gimenez, A. Hoogi, K. K. Miyake, M. Gorovoy, and D. L.Rubin, “A curated mammography data set for use in computer-aideddetection and diagnosis research,”Scientific data, vol. 4, no. 1, pp. 1–9,2017.

[7] H. Cai, J. Wang, T. Dan, J. Li, Z. Fan, W. Yi, C. Cui, X. Jiang, andL. Li, “An online mammography database with biopsy confirmed types,”Scientific Data, vol. 10, no. 1, p. 123, 2023.

[8] C. Balleyguier, S. Ayadi, K. Van Nguyen, D. Vanel, C. Dromain, andR. Sigal, “Birads™ classification in mammography,”European journalof radiology, vol. 61, no. 2, pp. 192–194, 2007.

[9] W. Tang, F. Zhou, S. Huang, X. Zhu, Y. Zhang, and B. Liu, “Featurere-embedding: Towards foundation model-level performance in computationalpathology,” in Proceedings of the IEEE/CVF Conference onComputer Vision and Pattern Recognition, 2024, pp. 11 343–11 352.

[10] S. Pathak, J. Schl¨otterer, J. Geerdink, J. Veltman, M. van Keulen,N. Strisciuglio, and C. Seifert, “Case-level breast cancer prediction forreal hospital settings,”arXiv preprint arXiv:2310.12677, 2023.

[11] L. Shen, “End-to-end training for whole image breast cancer diagnosisusing an all convolutional design,”arXiv preprint arXiv:1711.05775,2017.

[12] L. Liao and E. M. Aagaard, “An open codebase for enhancing transparencyin deep learning-based breast cancer diagnosis utilizing cbisddsmdata,”Scientific reports, vol. 14, no. 1, p. 27318, 2024.

[13] J. Wolleb, F. Bieder, R. Sandk¨uhler, and P. C. Cattin, “Diffusion modelsfor medical anomaly detection,” in International Conference on Medicalimage computing and computer-assisted intervention. Springer, 2022,pp. 35–45.

[14] J. Zhu, C. Ding, Y. Tian, and G. Pang, “Anomaly heterogeneity learningfor open-set supervised anomaly detection,” in Proceedings of theIEEE/CVF conference on computer vision and pattern recognition, 2024,pp. 17 616–17 626.

[15] A. Rahman, J. M. J. Valanarasu, I. Hacihaliloglu, and V. M. Patel,“Ambiguous medical image segmentation using diffusion models,” inProceedings of the IEEE/CVF conference on computer vision andpattern recognition, 2023, pp. 11 536–11 546.

[16] D. Ribli, A. Horv´ath, Z. Unger, P. Pollner, and I. Csabai, “Detectingand classifying lesions in mammograms with deep learning,”Scientificreports, vol. 8, no. 1, p. 4165, 2018.

[17] R. S. Lee, F. Gimenez, A. Hoogi, K. K. Miyake, M. Gorovoy, and D. L.Rubin, “A curated mammography data set for use in computer-aideddetection and diagnosis research,”Scientific data, vol. 4, no. 1, pp. 1–9,2017.

[18] P. Bandi, O. Geessink, Q. Manson, M. Van Dijk, M. Balkenhol,M. Hermsen, B. E. Bejnordi, B. Lee, K. Paeng, A. Zhong et al., "Fromdetection of individual metastases to classification of lymph node statusat the patient level: the camelyon17 challenge,"IEEE transactions onmedical imaging, vol. 38, no. 2, pp. 550–560, 2018.

[19] W. Lingle, B. J. Erickson, M. L. Zuley, R. Jarosz, E. Bonaccio,J. Filippini, J. M. Net, L. Levi, E. A. Morris, G. G. Figler et al., "Thecancer genome atlas breast invasive carcinoma collection (tcga-brca),"(No Title), 2016.

[20] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal,G. Sastry, A. Askell, P. Mishkin, J. Clark et al., "Learning transferablevisual models from natural language supervision," in Internationalconference on machine learning. PmLR, 2021, pp. 8748–8763.

[21] A. Xu, W. Li, P. Guo, D. Yang, H. R. Roth, A. Hatamizadeh, C. Zhao,D. Xu, H. Huang, and Z. Xu, "Closing the generalization gap ofcross-silo federated medical image segmentation," in Proceedings of theIEEE/CVF Conference on Computer Vision and Pattern Recognition,2022, pp. 20 866–20 875.

[22] H. Guan, P.-T. Yap, A. Bozoki, and M. Liu, "Federated learning formedical image analysis: A survey,"Pattern Recognition, p. 110424,2024.

[23] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas,"Communication-efficient learning of deep networks from decentralizeddata," in Artificial intelligence and statistics. PMLR, 2017, pp. 1273–1282.

[24] M. Jiang, H. Yang, C. Cheng, and Q. Dou, "Iop-fl: Inside-outside personalizationfor federated medical image segmentation,"IEEE Transactionson Medical Imaging, vol. 42, no. 7, pp. 2106–2117, 2023.

[25] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai,T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly et al.,"An image is worth 16x16 words: Transformers for image recognitionat scale,"arXiv preprint arXiv:2010.11929, 2020.

[26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez,Ł. Kaiser, and I. Polosukhin, "Attention is all you need,"Advances inneural information processing systems, vol. 30, 2017.

[27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for imagerecognition," in Proceedings of the IEEE conference on computer visionand pattern recognition, 2016, pp.770–778.

[28] S. Gupta, D. Zhang, M. P. Sampat, and M. K. Markey, "Combiningtexture features from the mlo and cc views for mammographic cadx,"in Medical Imaging 2006: Image Processing, vol. 6144. SPIE, 2006,pp. 1877–1885.

[29] Z. Shao, H. Bian, Y. Chen, Y. Wang, J. Zhang, X. Ji et al., "Transmil:Transformer based correlated multiple instance learning for wholeslide image classification,"Advances in neural information processingsystems, vol. 34, pp. 2136–2147, 2021.

[30] N. K. Singh and K. Raza, "Medical image generation using generativeadversarial networks: A review,"Health informatics: A computationalperspective in healthcare, pp. 77–96, 2021.

[31] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models,"Advances in neural information processing systems, vol. 33, pp. 6840–6851, 2020.

[32] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith,"Federated optimization in heterogeneous networks,"Proceedings ofMachine learning and systems, vol. 2, pp. 429–450, 2020.

[33] H. N. Khan, A. R. Shahid, B. Raza, A. H. Dar, and H. Alquhayz,"Multi-view feature fusion based four views model for mammogramclassification using convolutional neural network,"IEEE Access, vol. 7,pp. 165 724–165 733, 2019.

[34] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federatedlearning with non-iid data,"arXiv preprint arXiv:1806.00582, 2018.

★★★