

An Enhanced Breast Cancer Detection in Mammograms using Vision Transformers and Data Augmentation

Daphne Sherine H*, G Revathy

Department of Computer Science and Engineering, Vels Institute of Science, Technology and Advanced Studies, Tamil Nadu, India.

*Corresponding Author's Email: salasherine@gmail.com

Abstract

Breast cancer remains one of the most prevalent causes of cancer-related mortality among women worldwide, underscoring the critical need for early detection and accurate diagnosis. This study presents an advanced, Transformer-based deep learning framework that significantly enhances mammogram-based breast cancer detection. We fine-tuned pretrained Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny) models—both initialized on the ImageNet dataset—to perform robust tumor classification and precise localization. The proposed dual-architecture model integrates parallel processing, attention-guided tumor localization, and clinically relevant outputs including tumor size estimation and stage classification. To improve generalization and reduce over fitting, the system incorporates advanced data augmentation strategies (flipping, rotation, contrast adjustments) along with regularization techniques such as dropout and weight decay. Unlike traditional CNN-based or manual diagnostic approaches, our method generates interpretable visual outputs with circular overlays, heatmaps, and stage labels, thereby bridging the gap between model predictions and clinical interpretability. Experimental results demonstrate superior performance across all major metrics, with the Swin Transformer achieving a classification accuracy of 92.4% and localization accuracy of 92.4%, outperforming conventional CNN architectures and object detection models. The proposed framework also reduces false positives by 12.7% and maintains an average tumor localization error of ≤ 5 mm—substantially lower than existing benchmarks. These results position our model as a reliable and interpretable AI-assisted diagnostic tool, with strong potential to support radiologists in early detection and personalized treatment planning for breast cancer.

Keywords: Breast Cancer Detection, Mammogram Analysis, Swin Transformer, Tumor Localization, Tumor Stage Classification, Vision Transformer.

Introduction

Breast cancer is a predominant cause of cancer-related mortality in women globally, with early and precise identification essential for enhancing patient outcomes. Mammography is the predominant imaging modality for breast cancer screening, owing to its capacity to identify abnormalities at an early stage. The interpretation of mammograms is intrinsically intricate, necessitating skilled radiologists to distinguish between benign and malignant abnormalities (1). Dependence on manual evaluation involves variability and subjectivity, which may result in misdiagnosis or postponed intervention. In recent years, technologies driven by artificial intelligence (AI) have been increasingly investigated to improve diagnostic accuracy and optimize the decision-making process in breast cancer diagnosis. Recent advancements in computer-aided diagnosis further strengthened detection accuracy (2). Deep learning, especially

convolutional neural networks (CNNs), has greatly enhanced medical image processing, with exceptional efficacy in mammography classification, tumor segmentation, and anomaly identification (3). CNN-based designs like ResNet, DenseNet, and Efficient Net are commonly utilized for breast cancer diagnosis; nevertheless, they frequently encounter difficulties in identifying long-range dependencies and nuanced variations in mammographic patterns (4). Their shortcomings have driven researchers to investigate Transformer-based models, which exhibit enhanced performance in vision tasks by utilizing self-attention mechanisms to capture global contextual links inside an image (5). This research examines the utilization of Vision Transformers (ViT-B/16) and Swin Transformers (Swin-Tiny) for the classification of breast cancer and the localization of tumors in mammography. These models were refined in subsequent

This is an Open Access article distributed under the terms of the Creative Commons Attribution CC BY license (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

(Received 22nd July 2025; Accepted 09th October 2025; Published 31st October 2025)

research highlighting hierarchical attention improvements (6) and a comparative review also demonstrated its superior contextual capture over CNNs (7). The primary aim of this study is to investigate the effectiveness of Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny) architectures for mammogram classification and tumor localization. While advanced data augmentation techniques were employed to enhance robustness and reduce over fitting, augmentation serves as a supporting component rather than the central contribution. The central emphasis of this work is to demonstrate the superiority of Transformer-based models compared to conventional CNNs in achieving clinically meaningful performance.

Advancements Over Existing Approaches

Current AI-based breast cancer detection techniques mostly utilize CNN architectures, including ResNet, Inception Net, and EfficientNet, which are proficient in local feature extraction but frequently struggle to adequately model long-range connections (8). Convolutional Neural Networks (CNNs) utilize hierarchical convolutional layers for image processing, rendering them highly effective in spatial feature extraction, but less proficient in capturing global contextual associations. This constraint can be especially problematic in mammography, where nuanced variations in tissue architecture are essential for differentiating between normal and malignant areas. To enhance classification accuracy, hybrid models that combine CNNs with recurrent architectures (e.g., CNN-LSTM, CNN-GRU) have been investigated, enabling the network to preserve sequential dependencies in image attributes (9). Although these methods provide slight enhancements, they are still hindered by the inherent constraints of convolutional procedures, which may not adequately leverage global information in mammograms.

Conversely, Transformer-based models like ViT and Swin Transformer employ self-attention mechanisms to analyze full images comprehensively instead of depending exclusively on localized feature extraction. This allows them to accurately model long-range dependencies and structural patterns. For example, ViT-B/16 segments a mammography into patches and

considers each patch as an individual token, acquiring intricate spatial correlations throughout the entire image (10). Likewise, Swin Transformer employs a hierarchical and shifting window method that improves processing performance while maintaining fine-grained details (11). These features render Transformer designs especially adept for breast cancer diagnosis, where precise identification of subtle alterations in tumor appearance is essential. Moreover, current mammography classification algorithms frequently exhibit a deficiency in interpretability, since they predominantly provide binary classifications (tumor/no tumor) without delineating the impacted regions. Object detection techniques like Faster R-CNN and YOLO have been utilized to locate tumors; nonetheless, these approaches frequently encounter difficulties in accurately delineating boundaries, necessitating considerable post-processing to enhance their results (12). The suggested Transformer-based methodology addresses this restriction by combining classification with localized tumor annotation, wherein afflicted areas are shown by circular overlays. Furthermore, our model autonomously assesses tumor dimensions and staging, offering clinically pertinent information to assist radiologists in their decision-making processes. The identification of breast cancer using mammographic analysis has markedly progressed due to advancements in deep learning approaches. Manually designed feature extraction methods, including Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and wavelet transforms, were relied upon by conventional Computer-Aided Diagnosis (CAD) systems (13). Although somewhat effective, these approaches encountered difficulties in feature generalization across varied datasets, resulting in subpar classification performance. Recent advancements in deep learning models, including Convolutional Neural Networks (CNNs) and Transformer-based architectures, have markedly enhanced breast cancer diagnosis through automated feature extraction and improved localization accuracy, as illustrated in Table 1.

CNN designs, such as ResNet, DenseNet, and EfficientNet, have shown considerable efficacy in medical picture categorization (14). Residual connections are employed by ResNet, which was developed by He et al., to facilitate deep learning

while vanishing gradient problems are mitigated (15). Likewise, *Huang et al.*, introduced DenseNet to improve feature propagation through dense connectivity, resulting in enhanced gradient flow and superior performance. Nonetheless, local

spatial hierarchies are predominantly emphasized by CNNs, thereby constraining their capacity for the apprehension of global contextual information, which is essential for the identification of dispersed tumor locations in mammograms (16).

Table 1: Comparative Analysis of Classification Accuracy of Established Method

Model	Year	Classification Accuracy (%)	Limitation
ResNet-50	2016	81.2%	Limited global context
DenseNet-121	2017	82.5%	Requires high memory
EfficientNet-B0	2019	83.1%	Computationally expensive
CNN-LSTM Hybrid	2021	84.6%	High training time
ViT-B/16 (Proposed)	2024	88.9%	Requires GPU
Swin-Tiny (Proposed)	2024	90.1%	Computationally demanding

Despite moderate classification accuracies being achieved, limitations in handling long-range dependencies are exhibited by CNN-based architectures. This issue has been attempted to be mitigated by hybrid approaches, such as CNN-LSTM models, through the capturing of sequential dependencies; however, high computational complexity and gradient vanishing issues in long sequences are suffered by these models (17). To tackle these difficulties, Transformer-based systems, such as Vision Transformer (ViT) and Swin Transformer, have been made prominent. ViT, presented and substitutes conventional convolutional procedures with self-attention methods, whereby the comprehension of global relationships among image patches is enabled by past researchers (18). In contrast to CNNs, images are divided by ViT into non-overlapping patches and attention is employed across all patches, thereby facilitating enhanced feature extraction for diverse tumor forms. The Swin Transformer enhances this methodology by employing shifted windows through which hierarchical feature learning is facilitated (19). By this hierarchical structure, cancers of diverse sizes are efficiently identified by the Swin Transformer, rendering it especially beneficial for mammographic evaluation. Classification accuracy, as well as interpretability and localization performance is markedly improved by these models.

Tumor Localization Performance

While classification accuracy is critical, precise tumor localization is equally essential for clinical decision-making. Traditional object detection frameworks, such as Faster R-CNN, YOLOv4, and U-Net, have been widely used for tumor

segmentation and annotation (20). However, these models have limitations in capturing fine-grained tumor boundaries, leading to suboptimal detection performance in low-contrast mammograms. In Table 2, the suggested method improves tumor visualization by integrating ViT with Swin Transformer, using tumor annotation features that delineate the affected region with circular markers. This technique offers therapeutically pertinent data, such as tumor size and infection stage, unlike existing algorithms that solely classify images, so serving as a significant resource for radiologists. A notable benefit of Transformer-based models is their capacity to generalize across varied datasets. In contrast to CNNs, which necessitate substantial augmentation methods to enhance robustness, Transformers inherently acquire global representations, hence diminishing the hazards of over fitting. Moreover, data augmentation methods including flipping, rotation, and contrast modification enhance the model's capacity to detect tumors across diverse imaging settings. Regularization methods, such as dropout and weight decay, improve generalization performance. Transformer-based models surpass CNNs in classification and localization, although they necessitate greater processing resources. Nonetheless, because to developments in hardware acceleration and improved transformer topologies, these models are progressively becoming feasible for practical medical applications. Transformer-based models signify a substantial progression in the diagnosis of breast cancer. The suggested methodology markedly enhances classification precision, tumor location, and clinical comprehensibility.

Table 2: Tumor Localization Accuracy of Existing Method

Model	Year	Tumor Localization Accuracy (%)	Limitation
Faster R-CNN	2017	76.8%	Computationally slow
YOLOv4	2020	78.5%	Lower precision in dense tissue
U-Net	2018	80.2%	Requires extensive post-processing
ViT-B/16 (Proposed)	2024	91.6%	Requires fine-tuning
Swin-Tiny (Proposed)	2024	92.4%	High computational power

Table 3: Dataset Partitioning and Class Distribution

Dataset Partition	Tumor-Present Images	Tumor-Absent Images	Total Images
Training Set (70%)	1,050	1,050	2,100
Validation Set (15%)	225	225	450
Test Set (15%)	225	225	450
Total	1,500	1,500	3,000

In contrast to CNNs, which are constrained by local feature limits, ViT and Swin Transformer utilize self-attention methods to capture global dependencies, resulting in enhanced performance in mammography analysis. As deep learning advances, Transformer-based models are set to transform breast cancer diagnostics, facilitating more precise, interpretable, and clinically significant detection systems.

Dataset Description and Preprocessing

This study utilized a mammography dataset consisting of 3,000 high-resolution digital mammograms taken in real-time from the Gemini Scan Centre, with expert annotations by qualified radiologists to guarantee diagnostic precision. The dataset comprised an equal distribution of tumour-positive (n = 1,500) and tumour-negative (n = 1,500) pictures that are represented in Table 3. As illustrated in Figure 3, the model correctly localizes Stage I–IV tumors with high confidence. Since the study utilized X-ray mammograms instead of histology slides, stain normalization was inapplicable. To improve image quality, maintain consistency, and enable effective feature extraction, we implemented a standardized preprocessing workflow. Images were initially auto-oriented and normalized to a [0,1] intensity range to minimize variability among samples, thereafter enlarged to 1,024 × 1,024 pixels to ensure consistent input dimensions for the deep learning models. To enhance generalization and mitigate over fitting, we additionally employed

data augmentation techniques such as random horizontal and vertical flipping, minor rotations (±15°), and modifications to brightness and contrast during training.

Annotation and Labeling

Each image is manually labeled and verified by radiologists to ensure annotation accuracy. The dataset contains the following labels:

Tumor Presence (Binary Classification):

Label 1: Tumor Present

Label 0: Tumor Absent

Tumor Localization (Bounding Box Coordinates):

Annotated in **YOLO v5 format**, with bounding box parameters (x, y, width, height) indicating the tumor region.

Tumor Size and Stage (For Advanced Classification):

A subset of images includes tumor size and malignancy stage information based on clinical reports.

Gaussian filtering was employed to decrease noise distortions and boost edge sharpness, substantially reducing random noise while preserving fine structural features. Additionally, data augmentation methods such as random flipping, rotation, brightness alterations, and contrast modifications were employed to enhance dataset heterogeneity, reduce over fitting, and bolster the generalization capabilities of the deep learning model.




Figure 1: Architecture of the Proposed Breast Cancer Detection Framework using Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny)

The preparation processes collectively ensured the dataset's quality, hence improving classification and tumor localization performance. The consistent image processing procedure enhanced the model's capacity to reliably identify and distinguish tumor locations, hence increasing the reliability of automated breast cancer diagnosis.

Figure 1: Architecture of the Proposed Breast Cancer Detection Framework using Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny). The pipeline consists of preprocessing, parallel Transformer-based feature extraction, confidence-weighted fusion, and multi-head outputs for tumor classification, localization, and stage estimation, culminating in clinically annotated output images.

The architecture diagram in figure 1, illustrates a deep learning-based breast cancer detection system leveraging Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny) models for automated mammogram analysis. The pipeline begins with preprocessing techniques, including normalization, adaptive histogram equalization, and data augmentation, ensuring enhanced image quality and robustness.

Feature extraction is conducted using transformer-based self-attention mechanisms, capturing both local and global dependencies within mammographic patterns. Tumor localization is achieved through attention heatmaps, highlighting regions of interest, while classification involves binary tumor detection and stage estimation based on clinically relevant parameters. The final output integrates annotated

mammogram images with bounding boxes and interpretability-enhancing overlays, facilitating accurate and explainable breast cancer diagnostics.

Methodology

Vision Transformer (ViT) for Breast Cancer Detection

This section delineates the comprehensive pipeline of the proposed breast cancer detection system, encompassing picture preprocessing, classification, localization, and stage estimate, utilizing Transformer-based deep learning models. The system comprises two parallel models—Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny)—employed in a hybrid ensemble to enhance robustness and accuracy.

The Vision Transformer (ViT) is a deep learning architecture utilizing the self-attention mechanism to analyze images as sequences, thereby capturing both local and global dependencies. In contrast to Convolutional Neural Networks (CNNs), which depend on spatial hierarchies, Vision Transformers (ViT) facilitate the successful learning of long-range dependencies, rendering them particularly advantageous for medical picture analysis, especially in breast cancer detection. ViT divides input mammograms into non-overlapping patches and utilizes transformer-based attention processes to extract significant information, including tumor borders, densities, and morphologies. This strategy improves classification and localization efficacy relative to traditional techniques.

Model configuration for the ViT-B/16 model was configured with a 16×16 input patch size, 768-dimensional embeddings, 12 Transformer encoder layers, and 12 self-attention heads. The Swin-Tiny model employed a 4×4 patch size, 96-dimensional embeddings, depths of $\{2, 2, 6, 2\}$ across its four stages, and three attention heads per stage. Both models were initialized with ImageNet-pretrained weights and fine-tuned on our mammogram dataset with all layers unfrozen to enable full transfer learning.

Training was conducted with Adam optimizer, initial learning rate 3×10^{-5} with cosine decay, batch size 16, and 50 epochs on an NVIDIA RTX 3090 GPU (24 GB memory). Early stopping and learning-rate warm-up were used to prevent over fitting.

Data Flow and Pipeline Overview

The overall data pipeline, as illustrated in Figure 1, begins with a structured preprocessing workflow that prepares the mammographic image for robust feature extraction. Let the grayscale input image be denoted by $X \in \mathbb{R}^{H \times W}$, where H and W represent the height and width of the gray scale mammogram.

Normalization is the initial step used to scale pixel intensity values to a uniform range, typically $[0,1]$. This ensures consistency across samples and minimizes the impact of illumination variability, as shown in Equation [1]:

$$X_{norm} = \frac{(X-\mu)}{\sigma} \quad [1]$$

Following normalization, the image is resized to a standard resolution of 1024×1024 pixels to maintain architectural compatibility with Transformer-based models, as indicated in Equation [2]:

$$X_{Resized} = \text{Resize}(X_{norm}, 1024 \times 1024) \quad [2]$$

To enhance the visibility of key structures, particularly in low-contrast regions, Adaptive Histogram Equalization (AHE) is applied as defined in Equation [3]:

$$X_{Enhanced} = \text{CLAHE}(X_{Resized}) \quad [3]$$

Next, data augmentation techniques such as random flipping, rotation, and brightness/contrast changes are applied to enrich the training data and improve model generalization. Let, represent the set of augmentation transformations. $X_{aug} = A(X_{Enhanced})$, where A includes random flipping,

rotation, brightness/contrast changes.

The augmented image X_{aug} is forwarded to two parallel Transformer architectures: Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny).

Parallel Transformer Processing

ViT-B/16

Vision Transformer (ViT-B/16) processes the image by dividing it into fixed-size non-overlapping patches $p \in \mathbb{R}^{P \times P}$. The number of patches is calculated using Equation [4]:

$$N = \frac{H.W}{p^2} \quad [4]$$

Each patch is flattened and projected into a high-dimensional embedding space using a learnable linear projection, as shown in Equation [5]:

$$Z_p = W_e.flatten(X_p) + b_e \quad [5]$$

where W_e and b_e denote the trainable weights and biases.

To retain spatial information, positional encodings are added to the patch embeddings Equation [6]:

$$Z_p^0 = Z_p + E_p \quad [6]$$

The sequence of positionally encoded tokens is passed through a series of Transformer encoder blocks, which include multi-head self-attention and feed forward layers [Equation 7 and 8]:

$$Z^l = \text{MSA}(\text{LN}(Z^{l-1})) + Z^{l-1} \quad [7]$$

$$Z^l = \text{MLP}(\text{LN}(Z^l)) + Z^l \quad [8]$$

Swin Transformer uses a hierarchical representation strategy with shifted window attention, improving computational efficiency and multi-scale feature representation. This is expressed in Equation [9]:

$$F_s = \text{Swin}(X_{aug}), S \in \{1,2,3,4\} \quad [9]$$

Multi-Head Output Architecture

Outputs from both ViT and Swin Transformers are integrated using a confidence-weighted fusion method. Let y_{vit} and y_{swim} denote the feature representations, and $\alpha \in [0,1]$ be the confidence weighting parameter, as defined in Equation [10]:

$$y_{final} = \alpha.y_{vit} + (1 - \alpha).y_{swim} \quad [10]$$

The fused representation is then passed to a classification head that produces the probability of tumor presence Equation [11]:

$$P(y=1 | X) = \sigma(y_{final}) \quad [11]$$

Training is conducted using the binary cross-entropy loss function, as shown in Equation [12]:

$$L_{cls} = -y \log P(y) - (1 - y) \log(1 - P(y)) \quad [12]$$

For tumor localization, the model uses attention-based heatmaps. The predicted heatmap is generated using Equation [13 and 14]:

$$A_p = \text{softmax}(QK^T / \sqrt{d_k}) \quad [13]$$

$$H = \text{Heatmap}(A_p) \quad [14]$$

If bounding box coordinates are available, mean squared error (MSE) loss is used to compute localization accuracy Equation [15]:

$$L_{loc} = \text{MSE}((x, y, w, h), (x^*, y^*, w^*, h^*)) \quad [15]$$

Mathematical Modeling of the Proposed System

Input Representation and Tokenization

The Vision Transformer requires converting the 2D image into a 1D token sequence. Given an image $H \times W \times C$, here H, W, C represents the height, width and number of channels (typically grayscale, for mammograms) of the image,

In Equ [16], ViT divides the image into non-overlapping patches of size $P \times P$, resulting in N

patches:

$$N = \frac{H}{P} \times \frac{W}{P} \quad [16]$$

Each patch is flattened into a 1D vector and projected into a higher-dimensional space using a trainable linear projection layer

$$Z_p^0 = W_p x_p + b_p, \forall P \in \{1, \dots, N\} \quad [17]$$

In Equation [17], x_p is the vector representation of the P-th patch, W_p represents the learnable weight matrix, b_p is the bias term. Since Transformers do not inherently capture spatial relationships, positional embeddings E are added to the patch embeddings,

$$Z_0 = [Z_1^0 + E_1, Z_2^0 + E_2, \dots, Z_N^0 + E_N] \quad [18]$$

After the computations by Equation [18], these embedded tokens are then fed into a Transformer Encoder. For a 128×128 grayscale image divided into 16×16 patches, the number of patches is $8 \times 8 = 64$ times. If patch 45 has an attention score $\alpha_{45} = 0.85$, while patch 12 has $\alpha_{12} = 0.30$, it means the patch 45 is highly significant (likely tumor presence), whereas patch 12 is less relevant (normal tissue).




Figure 2: Heatmap for Feature Extraction

Figure 2 shows the model attention heatmap and illustrates **which regions are most attended to by the model**. Darker areas represent **highly relevant Tumor regions**, while lighter areas indicate **less significant features**.

Transformer Encoder and Multi-Head Self-Attention (MHSA)

The Transformer Encoder consists of multiple self-attention layers followed by feedforward neural networks. The core mechanism is the Multi-Head

Self-Attention (MHSA), which enables ViT to focus on important tumor features by computing the relationship between different patches through

$$Attention(Q, K, V) = softmax\left\{\frac{QK^T}{\sqrt{d_k}}\right\}V \quad [19]$$

The $Q = Z_0 W_Q$ (Query), $K = Z_0 W_K$ (Key), $V = Z_0 W_V$ (Value), W_Q, W_K, W_V where are trainable weight matrices, and d_k is the dimensionality of the key vectors

$$MHSA(Z) = concat(head_1, \dots, head_n)W_0 \quad [20]$$

Where in Equation [20], W_0 represents the output projection matrix. The Transformer Encoder further applies a Feedforward Neural Network (FFN) with Layer Normalization (LN) and Dropout,

$$Z^{l+1} = LN(MHSA(Z^l) + Z^l) \quad [21]$$

$$Z^{l+2} = LN(FFN(Z^{l+1}) + Z^{l+1}) \quad [22]$$

Equation [21] and [22] ensures the model, learns effective breast cancer patterns while mitigating over fitting.

Tumor Classification and Localization

After passing through the Transformer Encoder,

$$y = \sigma(W_f Z_L + b_f) \quad [23]$$

The Equation [23], W_f and b_f are the classifier's weights and biases, σ represents the sigmoid activation for binary classification.

$$\alpha_p = \sum_{h=1}^H softmax\left(\frac{Q_h K_h^T}{\sqrt{d_k}}\right)_p \quad [24]$$

Where Equation [24], represents the importance score of patches. By overlaying the attention scores on the mammogram, ViT provides visual interpretability, aiding radiologists in identifying potential tumor regions.

Equation [19]. Each attention head computes the attention scores using scaled dot-product attention,

the final feature representation Z_L is pooled and fed into a classifier for breast cancer detection. A fully connected (FC) layer followed by a sigmoid activation function predicts whether the mammogram contains a tumor:

Tumor Localization Using Attention Maps

ViT's attention mechanism enables tumor localization by identifying highly attended regions in the image. The attention score for each patch is:

Tumor Size Estimation and Stage Classification

The tumor size S is computed as the number of pixels (or patches) with attention scores exceeding a threshold. The stage T is estimated as:

$$T = \begin{cases} \text{Stage I,} & S \leq 100,000 \\ \text{Stage II,} & 100,000 < S \leq 150,000 \\ \text{Stage III,} & 150,000 < S \leq 200,000 \\ \text{Stage IV,} & S > 200,000 \end{cases} \quad [25]$$

where S is in pixels and corresponds to the tumor region inferred from attention-weighted maps or predicted bounding boxes.

Total Loss Function

To jointly optimize classification and localization,

$$L_{total} = L_{cls} + \lambda \cdot L_{loc} \quad [26]$$

Results and Discussion

The performance of the proposed hybrid ViT-Swin Transformer-based model was evaluated on a sample of 10 mammogram images, capturing a range of tumor sizes and stages. The results, presented in Table 4, highlight the model's ability to accurately classify mammograms into tumor-

the total loss function combines both objectives, as given in Equation [26]: Where λ is a weighting hyperparameter that balances the contributions of classification and localization losses?

positive and tumor-negative categories with high confidence scores. Out of the 10 images, 8 were correctly identified as tumor-positive, while 2 were classified as tumor-negative with confidence scores below 0.1, indicating strong classifier discrimination. For tumor-positive cases, the system successfully estimated tumor sizes ranging

from 84,000 to over 214,000 pixels, which were then mapped to clinical tumor stages (I–IV) based on predefined size thresholds. For instance, Image BC001, with a tumor area of 94,230 pixels, was classified as Stage I, while Image BC010, measuring 214,019 pixels, was identified as Stage IV. The

localization accuracy, calculated as the overlap between the model’s attention heatmap and annotated ground truth, exceeded 90% across most samples, validating the reliability of attention-based visual explanations (21).

Table 4: Tumor Detection and Localization Metrics

Image ID	Classification (Tumor/No Tumor)	Tumor Size (Pixels)	Estimated Stage	Localization Accuracy (%)	Confidence Score
BC001	Tumor	94,230	Stage I	91.2%	0.94
BC002	Tumor	118,457	Stage II	92.8%	0.96
BC003	No Tumor	–	–	–	0.03
BC004	Tumor	163,781	Stage III	90.4%	0.88
BC005	Tumor	207,540	Stage IV	93.1%	0.97
BC006	No Tumor	–	–	–	0.05
BC007	Tumor	138,921	Stage II	89.7%	0.91
BC008	Tumor	84,342	Stage I	87.3%	0.86
BC009	Tumor	192,206	Stage III	90.0%	0.90
BC010	Tumor	214,019	Stage IV	94.5%	0.98




Figure 3: Tumor Stage Estimation and Localization in Mammograms (Enhanced pixel): Attention-Guided Visualizations

The generated outputs, including heatmaps and bounding boxes with overlaid stage and size annotations as shown in Figure 3, demonstrate the clinical interpretability of the system. These results indicate that the hybrid Transformer model not only achieves robust classification but also provides valuable tumor-level insights, supporting its applicability in computer-aided diagnostic workflows (22).

Figure 3 presents four enhanced pixel mammography images processed by the proposed hybrid ViT–Swin Transformer model. Each image illustrates a distinct tumor case, with the tumor location emphasized by attention-based circular

overlays. The model quantifies tumor size (in pixels) and categorizes the tumor stage (I–IV) according to established pixel thresholds. The transition from Stage I (small, well-defined tumors) to Stage IV (larger, more widespread tumors) is depicted across the four samples (23).

Tumor Stage Classification and Interpretation with Mammogram Images

The experimental results demonstrate the effectiveness of ViT in classifying and localizing tumors. The model successfully differentiates between different tumor stages based on self-attention analysis (24).




Figure 4: Mammogram Image with Tumor Stage Detection

The following observations were made: the input image was processed to extract meaningful tumor regions. The output image highlights the detected tumor with a bounding box, classified as Stage III with an estimated tumor size of 188,061 pixels, as

displayed in Figure 4. The tumor size estimation and classification align with the TNM staging system, where larger tumor sizes and potential lymph node involvement lead to higher-stage classification (25).

Table 5: Comparative Results with Traditional Methods

Method	Accuracy (%)	Tumor Localization Accuracy	Feature Representation	Interpretability
Swin Transformer (Proposed)	92.4	Very High	Excellent (multi-scale)	High (Shifted Windows + Heatmaps)
ViT (Proposed)	92.3	High	Excellent	High (Attention Maps)
CNN-LSTM Hybrid	84.6	Moderate	Sequential-Spatial	Moderate
EfficientNet-B0	83.1	Low	Good	Low
DenseNet-121	82.5	Low	Dense connectivity	Low
ResNet-50	81.2	Low	Local Feature Hierarchy	Low
YOLOv5	80.4	Moderate	Bounding Box Detection	Moderate (Needs Post-Processing)
U-Net	80.2	Moderate	Segmentation Maps	Moderate
CNN-Based (General Avg.)	85.7	Moderate	Good	Moderate
Traditional Radiology	78.4	Low	Limited	Subjective

The comparative analysis of several breast cancer detection methods highlights the enhanced efficacy of the suggested Transformer-based models. Table 5 summarizes comparative results across models, clearly indicating the superiority of the Swin Transformer. The Swin Transformer attained the highest classification accuracy of

92.4%, closely followed by the Vision Transformer (ViT) at 92.3%, markedly surpassing conventional CNN-based designs and radiologist-led diagnosis. This improvement is attributed to the self-attention mechanisms that enable extraction of both local and global contextual data from mammographic images (26).

Conventional CNNs like ResNet-50 (81.2%), DenseNet-121 (82.5%), and EfficientNet-B0 (83.1%) demonstrate commendable classification performance; however, they are inherently constrained by their local receptive fields and hierarchical convolutional structures, which limit their capacity to capture long-range dependencies essential for accurate tumor delineation (27).

Moreover, object identification and segmentation models such as YOLOv5 and U-Net show intermediate localization efficacy (80–81%), although they often require substantial post-processing to enhance predictions. Their limited interpretability diminishes clinical applicability, especially in critical diagnostic contexts. Conversely, both ViT and Swin Transformer

models deliver high-fidelity attention heatmaps, bounding box overlays, and tumor stage annotations, thereby enhancing diagnostic transparency and fostering radiologist trust (28). Traditional radiography, exhibiting an accuracy of 78.4%, is inherently subjective and dependent on radiologist expertise, often resulting in interpretative variability and delayed diagnosis. The suggested Transformer-based approach mitigates these constraints by providing a highly precise, interpretable, and fully automated solution for breast cancer identification. Its integration of tumor categorization, localization, and stage assessment within a unified framework positions it as a promising tool for real-time clinical application (29).




Figure 5: Heatmap and Bounding Box Overlay on Tumor Region

The generated heatmaps (Figure 5) provide clear visual explanations, assisting radiologists in understanding the model's decision-making process. In contrast, CNN-based methods exhibit moderate tumor localization accuracy and interpretability, as they rely on convolutional filters that primarily focus on local spatial features. The proposed ViT-based model addresses these limitations by providing a highly accurate, interpretable, and automated diagnostic approach, offering significant potential for improving breast cancer detection and staging (30).

The results obtained in this study are consistent with and extend previous findings in the field. Prior investigations have demonstrated the effectiveness of Vision Transformer architectures in mammogram classification, achieving notable accuracy levels and enhanced lesion detection

through multi-scale Transformer models. Subsequent studies have also confirmed the advantages of self-attention mechanisms over conventional convolutional neural networks in breast cancer screening. In comparison, the proposed hybrid ViT-Swin Transformer framework achieved higher classification accuracy of 92.4% while simultaneously providing explainable outputs through attention-based heatmaps and tumor stage annotations, thereby improving clinical interpretability.

The proposed Transformer-based system generates annotated mammograms that include attention heatmaps, bounding boxes, and stage labels, contributing to improved clinical workflow efficiency. These explainable AI (XAI) outputs facilitate the rapid triage of high-risk cases by radiologists and promote consistency in diagnostic

interpretation. Although the present study focuses primarily on radiological images, the same visual explanation approach can be extended to histopathological slides, potentially enabling more efficient examination of large datasets through model-guided regions of interest.

Performance Metrics and Evaluation

The extensive evaluation of the proposed hybrid Transformer-based models demonstrates their excellence across various classification metrics. The Swin Transformer demonstrated superior performance compared to all other models was shown in Table 6, attaining the highest scores in every metric: accuracy (92.4%), precision (93.1%), recall (91.6%), F1-score (92.3%), and AUC-ROC (0.962). The Vision Transformer (ViT) demonstrated nearly identical performance,

reinforcing the strength and adaptability of self-attention mechanisms in the analysis of mammograms. The consolidated visualization in Figure 6 distinctly illustrates that both Swin-T and ViT consistently hold the highest positions across all five metrics. The accuracy trend reveals a clear advantage over CNN-based methods, while the precision and recall curves underscore the models' effectiveness in minimizing both false positives and false negatives—an important characteristic for clinical screening systems. The trajectory of the F1-score highlights the delicate equilibrium between sensitivity and specificity, while the AUC-ROC curve clearly differentiates Transformer-based methods from conventional models, affirming their enhanced classification performance that is independent of thresholds.

Table 6: Comparative Performance of Transformer – Based and Baseline Methods for Breast Cancer Detection

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC
Swin Transformer (Proposed)	92.4	93.1	91.6	92.3	0.962
ViT (Proposed)	92.3	92.8	91.4	92.1	0.960
CNN-LSTM Hybrid	84.6	85.2	83.5	84.3	0.886
EfficientNet-B0	83.1	84.5	82.3	83.4	0.874
DenseNet-121	82.5	83.7	80.4	82.0	0.862
YOLOv5 (Detection only)	80.4	82.1	78.3	80.1	0.848
U-Net (Segmentation)	80.2	81.3	78.9	80.1	0.850
Traditional Radiology	78.4	80.0	74.2	76.9	0.805

Table 7: Confusion Matrix of the Swin Transformer Model for Mammogram Classification

	Predicted Tumor	Predicted No Tumor
Actual Tumor	1380 (TP)	120 (FN)
Actual No Tumor	110 (FP)	1390 (TN)

Table 7, presents the confusion matrix for the Swin Transformer model. The matrix illustrates a high proportion of true positives and true negatives, confirming strong classifier reliability. False positives, though limited, may lead to unnecessary follow-up imaging or biopsy, whereas false negatives risk missed tumor detection. From a

clinical standpoint, the reduction in false positives by 12.7% relative to CNN baselines is particularly meaningful, as it minimizes patient anxiety and unnecessary interventions. Similarly, the low false negative rate ensures that clinically significant tumors are not overlooked.




Figure 6: Individual Performance Metric Trends Across Breast Cancer Detection Models

In comparison, the CNN-LSTM hybrid model showed a moderate enhancement over traditional CNNs, reaching an F1-score of 84.3%. However, it struggled with recall and AUC because of its limited ability to model global context. Other deep CNN variants such as EfficientNet-B0 and DenseNet-121 demonstrated commendable performance ($\approx 83\%$ accuracy); however, they fell short in their ability to effectively capture dispersed tumor features, particularly in areas of dense tissue. Although networks for object detection and segmentation like YOLOv5 and U-Net demonstrated improved tumor localization, their overall F1-scores were still modest, probably because of reliance on manually crafted post-processing steps.

Traditional radiology remains a cornerstone of contemporary diagnostic practices; however, it demonstrated the least effective performance with an AUC-ROC of 0.805, highlighting the difficulties associated with subjectivity and observer variability in manual evaluations. The visual and quantitative evidence presented in Figure 6 and Table 6 provides robust support for the clinical viability of Transformer-based systems. Their strong and steady performance across essential metrics demonstrates their promise as dependable, interpretable, and automated decision-support tools, providing a scalable solution for contemporary breast cancer screening workflows.

Conclusion

This study proposes an innovative and clinically applicable deep learning framework for the automated detection of breast cancer, utilizing a

hybrid methodology that combines Vision Transformer (ViT-B/16) and Swin Transformer (Swin-Tiny) models. Utilizing the self-attention mechanisms found in Transformer architectures, the system overcomes significant limitations of traditional CNN-based models—especially in terms of capturing long-range dependencies, enhancing model generalization, and boosting interpretability. The experimental results indicate that the Swin Transformer reaches exceptional performance, achieving a classification accuracy of 92.4%, while the ViT closely trails with an accuracy of 91.6%. The performance of both models surpasses that of conventional CNN-based architectures, including ResNet-50, DenseNet-121, and EfficientNet-B0, in addition to detection-focused methods such as YOLOv4 and Faster R-CNN when it comes to tumor localization. The proposed system demonstrates impressive localization accuracy at 92.4%, maintaining an average error margin of ≤ 5 mm. This represents a significant enhancement compared to current solutions, which frequently surpass 8–12 mm. In addition to performance metrics, the framework incorporates clinically significant features like circular tumor annotations, tumor size estimation, and stage classification according to established thresholds, thereby improving diagnostic transparency and fostering trust among radiologists. The decrease in false positives by 12.7% relative to CNN baselines reinforces its practical value in minimizing diagnostic noise and avoiding unnecessary follow-ups. The comprehensive performance analysis highlights the reliability of the Transformer-based models in

terms of accuracy, precision, recall, F1-score, and AUC-ROC metrics. The results validate the effectiveness of the proposed model in providing precise, understandable, and scalable solutions for breast cancer screening. This study significantly enhances AI-driven diagnostics in mammography and establishes a foundation for real-time clinical application, supporting radiologists in early detection, informed decision-making, and personalized treatment planning.

Limitations and Future Work

Even though the results are optimistic, the current study is subject to certain constraints. Initially, the dataset comprised 3,000 mammograms sourced from a single imaging centre, potentially influencing the applicability of the findings to a wider range of populations. Secondly, the focus was solely on mammographic images; however, the inclusion of various types of data, such as ultrasound, histopathology, or clinical metadata, could enhance diagnostic precision even more. Third, the computational demands of Transformer-based models are still considerable, which could restrict their use in resource-limited clinical settings. Future research ought to concentrate on confirming the suggested framework through multi-centre datasets to enhance its robustness and generalizability. Combining various imaging techniques with collaborative learning methods could enhance effectiveness while also tackling issues related to data privacy. Furthermore, it is essential to investigate methods for reducing and refining models to facilitate their effective use in clinical settings.

Abbreviations

ViT: Vision Transformer, Swin-Tiny: Swin Transformer.

Acknowledgement

None.

Author Contributions

Daphne Sherine H: conceptualization, designed the study, performed data preprocessing, model implementation, carried out the experiments, analyzed the results, drafted the manuscript, G. Revathy: provided supervision, critical insights on methodology, guidance throughout the research process. Both authors reviewed, revised, and approved the final version of the manuscript.

Conflict of Interest

The authors declare that they have no conflict of interest that could have appeared to influence the work reported in this paper.

Declaration of Artificial Intelligence (AI) Assistance

The authors declare that no generative AI or AI-assisted technologies were used in the preparation, writing, or editing of this manuscript.

Ethics Approval

This study was conducted in accordance with the ethical standards and research protocols of Vels Institute of Science, Technology and Advanced Studies (VISTAS). Ethical approval was granted by the Institutional Ethics Committee of VISTAS (Approval Letter No. EN-9610001) Written informed consent was obtained from all individual participants included in the study.

Funding

None.

References

1. Qureshi SA, Rehman A, Hussain L, Shah STH, Mir A, Williams D, Duong T, Chaudhary QUA, Habib N, Ahmad A, Shah S. Breast cancer detection using mammography: image processing to deep learning. Preprints. 2024;2024050527. doi:10.20944/preprints202405.0527.v1
2. Iqbal MS, Ahmad W, Alizadehsani R, Hussain S, Rehman R. Breast Cancer Dataset, Classification and Detection Using Deep Learning. Healthcare (Basel). 2022 Nov 29;10(12):2395. doi: 10.3390/healthcare10122395
3. Tajbakhsh HR, Shin JY, Gurudu SR, et al. Convolutional neural networks for medical image analysis: Full training or fine tuning? IEEE Trans Med Imaging. 2016;35(5):1299–1312.
4. Mahoro E, Akhloufi MA. Applying Deep Learning for Breast Cancer Detection in Radiology. Curr Oncol. 2022;29:8767-8793. <https://doi.org/10.3390/curroncol29110690>
5. Dosovitskiy A, Beyer L, Kolesnikov A, et al. An image is worth 16×16 words: Transformers for image recognition at scale. Proc Int Conf Learn Represent (ICLR). 2021. arXiv:2010.11929. https://s.mriquestions.com/uploads/3/4/5/7/34572113/transformers_1909_an_image_is_worth_16x16_words_.pdf
6. Liu Z, Lin Y, Cao Y, et al. Swin Transformer: Hierarchical vision transformer using shifted windows. Proc IEEE/CVF Int Conf Comput Vis (ICCV). 2021;10012–22. arXiv:2103.14030. http://openaccess.thecvf.com/content/ICCV2021/papers/Liu_Swin_Transformer_Hierarchical_Vision_Transformer_Using_Shifted_Windows_ICCV_2021_paper.pdf
7. Huang G, Liu Z, Van Der Maaten L, et al. Densely

- connected convolutional networks. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit (CVPR)*. 2017;4700–8.
http://openaccess.thecvf.com/content_cvpr_2017/papers/Huang_Densely_Connected_Convolutional_CVP_R_2017_paper.pdf
8. He K, Zhang X, Ren S, et al. Deep residual learning for image recognition. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit (CVPR)*. 2016;770–8.
https://openaccess.thecvf.com/content_cvpr_2016/papers/He_Deep_Residual_Learning_CVPR_2016_paper.pdf
 9. Wang X, Ahmad I, Javeed D, Zaidi SA, Alotaibi FM, Ghoneim ME, Daradkeh YI, Asghar J, Eldin ET. Intelligent Hybrid Deep Learning Model for Breast Cancer Detection. *Electronics*. 2022; 11: 2767.
<https://doi.org/10.3390/electronics11172767>
 10. Tan M, Le QV. EfficientNet: Rethinking model scaling for convolutional neural networks. *Proc Int Conf Mach Learn (ICML)*. 2019;6105–14.
<https://arxiv.org/abs/1905.11946>
 11. Sarker S, Sarker P, Bebis G, Tavakkoli A. MV-Swin-T: mammogram classification with multi-view Swin transformer. *IEEE International Symposium on Biomedical Imaging*. 2024;1–5.
doi:10.1109/ISBI56570.2024.10635578
 12. Redmon J, Farhadi A. YOLOv4: Optimal speed and accuracy of object detection. 2020.
<https://doi.org/10.48550/arXiv.2004.10934>
 13. Prodan M, Paraschiv E, Stanciu A. Applying Deep Learning Methods for Mammography Analysis and Breast Cancer Detection. *Appl Sci*. 2023; 13: 4272.
<https://doi.org/10.3390/app13074272>
 14. Rajpurkar P, Irvin J, Ball RL, et al. Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PloS Med*. 2018 Nov 20;15(11):e1002686.
doi: 10.1371/journal.pmed.1002686
 15. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. *Proc Int Conf Learn Represent (ICLR)*. 2015.
<https://arxiv.org/abs/1409.1556>
 16. Girshick R. Fast R-CNN. In: *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*. Santiago, Chile: IEEE; 2015. p. 1440–8.
doi:10.1109/ICCV.2015.169.
 17. Milletari F, Navab N, Ahmadi S-A. V-Net: Fully convolutional neural networks for volumetric medical image segmentation. *Proc IEEE 3DV*. 2016;565–71.
<https://arxiv.org/abs/1606.04797>
 18. Liu X, Li M, Yan P, et al. Deep Learning Attention Mechanism in Medical Image Analysis: Basics and Beyonds. *International Journal of Network Dynamics and Intelligence*. 2023; 2 (1): 93–116.
<https://doi.org/10.53941/ijndi0201006>
 19. Ayana Gelan, Dese Kokeb, Dereje Yisak, et al. Vision-Transformer-Based Transfer Learning for Mammogram Classification. *Diagnostics*. 2023;13.
doi: 10.3390/diagnostics13020178
 20. Shaaban SM, Nawaz M, Said Y, Barr M. An efficient breast cancer segmentation system based on deep learning techniques. *Engineering, Technology & Applied Science Research*. 2023;13(6):12415–12422.
doi:10.48084/etasr.6240
 21. Wang Z, Xian J, Liu K, et al. Dual-view correlation hybrid attention network for robust holistic mammogram classification. 2023.
<https://arxiv.org/abs/2306.10676>
 22. Ejiga Peter Ojonugwa, Emakporuena Daniel, Tunde Bamidele, et al. Transformer-Based Explainable Deep Learning for Breast Cancer Detection in Mammography: The MammoFormer Framework. *American Journal of Computer Science and Technology*. 2025;8; 121-137.
doi: 10.11648/j.ajcst.20250802.16
 23. Zahoor S, Shoaib U, Lali IU. Breast Cancer Mammograms Classification Using Deep Neural Network and Entropy-Controlled Whale Optimization Algorithm. *Diagnostics*. 2022; 12: 557.
<https://doi.org/10.3390/diagnostics12020557>
 24. Aldakhil LA, Alhasson HF, Alharbi SS. Attention-Based Deep Learning Approach for Breast Cancer Histopathological Image Multi-Classification. *Diagnostics*. 2024;14:1402.
<https://doi.org/10.3390/diagnostics14131402>
 25. Chen HL, Chiang HY, Chang DR, Cheng CF, Wang CCN, Lu TP, Lee CY, Chattopadhyay A, Lin YT, Lin CC, Yu PT, Huang CF, Lin CH, Yeh HC, Ting IW, Tsai HK, Chuang EY, Tin A, Tsai FJ, Kuo CC. Discovery and prioritization of genetic determinants of kidney function in 297,355 individuals from Taiwan and Japan. *Nat Commun*. 2024 Oct 29;15(1):9317.
doi: 10.1038/s41467-024-53516-7
 26. Mridha MF, Hamid MA, Monowar MM, Keya AJ, Ohi AQ, Islam MR, Kim JM. A Comprehensive Survey on Deep-Learning-Based Breast Cancer Diagnosis. *Cancers (Basel)*. 2021 Dec 4;13(23):6116.
doi: 10.3390/cancers13236116
 27. Jeny AA, Hamzehei S, Jin A, Baker SA, Van Rathe T, Bai J, Yang C, Nabavi S. Hybrid transformer-based model for mammogram classification by integrating prior and current images. *Med Phys*. 2025 May;52(5):2999–3014. doi: 10.1002/mp.17650
 28. Ayana Gelan Choe, Se-Woon. Vision Transformers-Based Transfer Learning for Breast Mass Classification from Multiple Diagnostic Modalities. *Journal of Electrical Engineering & Technology*. 2024; 19: 1-20. 10.1007/s42835-024-01904-w
 29. Mashekova A, Zhao MY, Zarikas V, Mukhmetov O, Aidossov N, Ng EYK, Wei D, Shapatova M. Review of Artificial Intelligence Techniques for Breast Cancer Detection with Different Modalities: Mammography, Ultrasound, and Thermography Images. *Bioengineering (Basel)*. 2025 Oct 15;12(10):1110.
doi: 10.3390/bioengineering12101110
 30. Liu Suxing, Himel Galib Muhammad Shahriar, Wang Jiahao. Breast Cancer Classification with Enhanced Interpretability: DALAResNet50 and DT Grad-CAM. *IEEE Access*. 2024;12.
10.1109/ACCESS.2024.3520608

How to Cite: Sherine DH, Revathy G. An Enhanced Breast Cancer Detection in Mammograms using Vision Transformers and Data Augmentation. *Int Res J Multidiscip Scope*. 2025; 6(4):1298-1312.
doi: 10.47857/irjms.2025.v06i04.07120