

**MULTI-STREAM DEEP LEARNING FOR BREAST CANCER  
CLASSIFICATION IN CONTRAST-ENHANCED SPECTRAL  
MAMMOGRAPHY**

**Ahmed A. H. Alkurdi<sup>1,2</sup>, Amira Bibo Sallow<sup>1</sup>**

<sup>1</sup>Department of Information Technology, Technical College of Duhok, Duhok Polytechnic  
University

<sup>2</sup>Information Technology Department, Technical College of Informatics-Akre, Akre  
University for Applied Sciences

ahmed.alaa@dpu.edu.krd, ahmed.alaa@auas.edu.krd, amira.bibo@dpu.edu.krd

**Abstract**

Breast cancer remains one of the leading causes of cancer-related mortality among women worldwide, necessitating early and accurate diagnosis to improve patient outcomes. Contrast-enhanced spectral mammography (CESM) has emerged as a promising imaging modality for detecting malignant lesions with enhanced sensitivity. However, automated breast cancer classification in CESM remains challenging due to the complexity of mammographic patterns and the need for multi-view analysis. In this study, a novel four-stream deep learning framework is proposed, that integrates multi-view mammographic imaging (craniocaudal and mediolateral oblique views), dual convolutional neural network (CNN) backbones (MobileNetV2 and EfficientNet-B0), and patient metadata (age and breast density) to enhance breast cancer classification. The proposed model extracts multi-scale, complementary features from both imaging perspectives, refines representations using multi-head attention and a token-based multilayer perceptron (MLP), and incorporates structured metadata for improved diagnostic accuracy.

The model was trained and evaluated on the Categorized Contrast-Enhanced Spectral Mammography (CDD-CESM) dataset, consisting of approximately 600 mammograms. The results demonstrate state-of-the-art performance, achieving an accuracy of 87.21% and an AUC-ROC of 0.9676, outperforming prior methods applied to the same dataset. The integration of metadata proved beneficial, allowing the model to learn clinically relevant correlations between mammographic patterns and patient risk factors. Additionally, the multi-view, multi-backbone approach provided a more comprehensive feature representation, leading to improved lesion characterization and class separability.

**Keywords:** Breast Cancer Diagnosis, Contrast-Enhanced Spectral Mammography (CESM), Deep Learning, Multi-View Mammography, Computer-Aided Diagnosis (CAD).

### 1. Introduction

Breast cancer remains one of the leading causes of cancer-related mortality among women worldwide, necessitating early and accurate diagnosis for improved prognosis and survival rates [1]. Mammography is the gold standard for early breast cancer screening; however, its diagnostic performance is often hindered by high false-positive rates, inter-radiologist variability, and limited feature extraction capabilities from conventional image processing techniques [2]. With the arrival of deep learning, automated breast cancer classification models have demonstrated promising results in improving diagnostic accuracy. However, their effectiveness is significantly dependent on the quality and structure of input data.

Current deep learning models for mammogram-based diagnosis predominantly process single-view images (Cranio-Caudal (CC) or Medio-Lateral Oblique (MLO)) without explicitly leveraging multi-view spatial correlations or integrating critical patient metadata such as age, family history, and previous biopsy results [3]. The lack of standardized preprocessing pipelines often results in suboptimal feature extraction, model generalizability issues, and inconsistencies in classification outcomes across datasets [4]. Additionally, preprocessing methods vary significantly, with some studies applying standard augmentation techniques while others employ feature-engineering-based approaches, leading to an urgent need for a more structured, reproducible, and clinically relevant pipeline [5].

To address these challenges, a novel preprocessing pipeline is proposed, that integrates multi-view mammography images (CC and MLO) into a unified tensor representation, incorporating patient-specific metadata to enhance feature learning in deep learning models. This approach hypothesizes that:

1. Combining CC and MLO views will improve spatial feature extraction by providing complementary breast tissue perspectives.
2. Integrating patient metadata will enable deep learning models to capture additional risk factors, thereby improving classification accuracy.
3. Standardizing the preprocessing pipeline across datasets will enhance reproducibility and cross-institutional model generalizability.

This research aims to develop and evaluate a robust, deep learning-based breast cancer classification model that utilizes the proposed preprocessing pipeline to classify mammograms into normal, benign, and malignant categories. This study will benchmark the proposed approach against conventional preprocessing techniques and single-view models to assess its impact on classification performance. By introducing a multi-view, patient-informed preprocessing strategy, this work aspires to contribute to the broader objective of improving automated breast cancer diagnosis and supporting radiologists with more reliable AI-assisted decision-making tools.

### 2. Related Work

Breast cancer diagnosis using deep learning-based approaches has gained significant traction in recent years, particularly with the advancement of contrast-enhanced spectral

mammography (CESM), which provides enhanced lesion visibility through low-energy and contrast-enhanced imaging. Despite its clinical advantages, automated breast cancer classification in CESM remains a challenging problem due to the complexity of mammographic patterns, the need for multi-view analysis, and the requirement for metadata integration. This section reviews key contributions in multi-view mammographic classification, contrast enhancement techniques, ROI-based processing, and deep learning architectures for CESM analysis.

Multi-view mammographic analysis plays a critical role in breast cancer diagnosis, as craniocaudal (CC) and mediolateral oblique (MLO) views provide complementary information regarding breast tissue structures and potential malignancies. Rofena (2025) introduced a multimodal, multi-view deep learning framework that integrates full-field digital mammography (FFDM) and CESM modalities to enable a virtual biopsy for breast cancer detection. Their findings suggest that incorporating multiple imaging modalities in both CC and MLO views enhances classification performance, particularly when combined with deep learning feature fusion techniques [6].

Similarly, Hassan (2024) proposed a multi-feature deep information bottleneck network that improves lesion classification in CESM by extracting multi-scale representations from multi-view images. Their method significantly improved classification accuracy and specificity, demonstrating that multi-view integration mitigates the limitations of single-view mammography analysis [7]. These studies highlight the necessity of multi-view learning strategies, as individual mammographic views often lack sufficient context for accurate tumor classification.

Contrast enhancement techniques have proven effective in highlighting tumor regions in CESM, thereby improving deep learning model performance. Ahmadyar (2024) developed a CycleGAN-based contrast enhancement technique that adapts mammograms for enhanced lesion visibility, reducing false negatives in dense breast tissues [8]. Similarly, Perek (2024) introduced a neighborhood representation loss function, specifically designed for contrast-enhanced spectral mammography, which significantly improves classification accuracy by refining local and global contrast variations [9]. ROI-focused preprocessing pipeline allows the model to extract meaningful tumor features without over-enhancing normal tissues, emphasizing the importance of ROI-specific contrast enhancement.

The effectiveness of ROI-based image processing in CESM classification has been widely studied. Hassan (2024) proposed a YOLO-based CAD framework, leveraging Vision Transformers (ViTs) for breast mass detection in CESM images, achieving state-of-the-art detection accuracy [7]. Their method highlights the significance of region-specific image processing, as breast masses often exhibit subtle texture variations that require localized feature extraction techniques.

Additionally, Lei (2024) explored automated tumor segmentation using CNNs, demonstrating that precise ROI segmentation improves classification accuracy by reducing background noise

and enhancing lesion contrast [10]. Their findings suggest that deep learning models benefit significantly from pre-segmented lesion regions, reinforcing the importance of ROI-based preprocessing in CESM-based classification.

Deep learning models have demonstrated remarkable success in automated mammographic analysis, with various CNN architectures being proposed for breast cancer classification. Ben Ahmed (2024) introduced a transformer-based multi-modal classification framework, incorporating metadata-driven learning alongside image-based features, achieving significant improvements in diagnostic accuracy [5]. Their approach underscores the importance of metadata integration, as clinical factors such as age and breast density provide critical diagnostic insights.

Similarly, Hussain (2024) proposed a hybrid CNN-Transformer model for mammography-based breast cancer classification, leveraging multi-scale feature fusion to enhance class separability [11]. Their model achieved high AUC-ROC scores, emphasizing the benefits of combining convolutional feature extraction with transformer-based contextual learning.

While previous studies have demonstrated the effectiveness of multi-view learning, contrast enhancement, ROI-based segmentation, and deep learning architectures in CESM classification, several research gaps remain:

- Limited research on metadata-driven deep learning models for CESM, despite age and breast density being clinically relevant predictors of breast cancer.
- Few studies have explored hybrid CNN architectures incorporating multiple backbones for improved multi-scale feature extraction.
- ROI-based preprocessing techniques are underutilized in many deep learning pipelines, despite evidence showing their positive impact on classification accuracy.

The proposed four-stream deep learning framework addresses these gaps by:

1. Incorporating multi-view mammography processing, leveraging both CC and MLO perspectives.
2. Employing dual CNN backbones (MobileNetV2 and EfficientNet-B0) to extract complementary image features.
3. Integrating metadata (age and breast density) directly into the deep learning pipeline, improving the clinical relevance of predictions.
4. Implementing ROI-based CLAHE preprocessing, ensuring optimal contrast enhancement within tumor regions.

By combining these advanced deep learning methodologies, the proposed model achieves superior classification performance, reinforcing the potential of AI-driven approaches in CESM-based breast cancer diagnosis.

### 3. Methodology

#### 3.1 Dataset

Breast cancer diagnosis using deep learning requires high-quality imaging datasets that offer detailed annotations, multimodal information, and standardized imaging protocols. The Categorized Contrast-Enhanced Mammography Dataset (CDD-CESM) [12], provided by The Cancer Imaging Archive (TCIA), was selected as the primary dataset for this research due to its comprehensive metadata, annotated imaging, and novelty in AI-driven breast cancer diagnostics. The Categorized Contrast-Enhanced Mammography Dataset (CDD-CESM) consists of 2,006 high-resolution CESM images with full patient metadata and radiological annotations.

The Categorized Contrast-Enhanced Spectral Mammography (CDD-CESM) dataset provides a valuable foundation for deep learning in breast cancer classification, but it has several limitations. A key issue is class imbalance, with 662 malignant, 587 benign, and 757 normal images, which may bias the model's learning despite mitigation strategies like class weighting. Additionally, 1,082 images lack breast density information, limiting the effectiveness of metadata-driven learning, as breast density is a critical factor in cancer risk assessment.

The dataset's relatively small size (~600 images used for training and testing) restricts the scalability and robustness of deep learning models, potentially inflating performance metrics due to dataset-specific generalization. Moreover, as the dataset originates from a single medical institution, it may introduce demographic and imaging biases, affecting real-world adaptability. Another limitation is the lack of precise tumor segmentation maps, preventing the use of advanced localization-based deep learning models for lesion detection.

#### 3.2 Preprocessing Pipeline

The preprocessing of medical imaging datasets is a crucial step in ensuring robust and reliable deep learning model performance. The Categorized Contrast-Enhanced Spectral Mammography (CDD-CESM) dataset comprises multi-view mammographic images, along with detailed patient metadata, providing an optimal foundation for a multimodal deep learning model for breast cancer classification. The preprocessing pipeline developed for this study integrates image normalization, contrast enhancement, region-of-interest (ROI) segmentation, and metadata encoding into structured tensors for deep learning applications. This structured approach improves data consistency, feature extraction efficiency, and classification accuracy.

The dataset consists of two standard mammographic views per breast side, namely the Mediolateral Oblique (MLO) and Cranio-Caudal (CC) views, which provide complementary anatomical perspectives. Previous research has demonstrated that multi-view mammography enhances lesion detectability and reduces false-positive rates compared to single-view models [1]. Each patient may have imaging data for both left and right breasts, and if a particular view is missing, a placeholder image (black image of size  $224 \times 224$  pixels) is used to maintain tensor uniformity. Studies have shown that deep learning models benefit from consistent input dimensions, reducing errors in convolutional feature extraction layers [13].

In addition to imaging data, the dataset contains essential patient metadata, including age, breast density (BI-RADS A–D), and pathology classification (Normal, Benign, Malignant). Prior studies have emphasized the role of integrating structured clinical metadata into deep learning models, reporting an increase in classification accuracy by up to 12% when age and breast density information are included [4]. To effectively integrate this metadata into the deep learning pipeline, a novel approach was employed wherein the age and breast density information were embedded directly into the third channel of the input tensor, eliminating the need for separate model inputs and ensuring a consistent input structure.

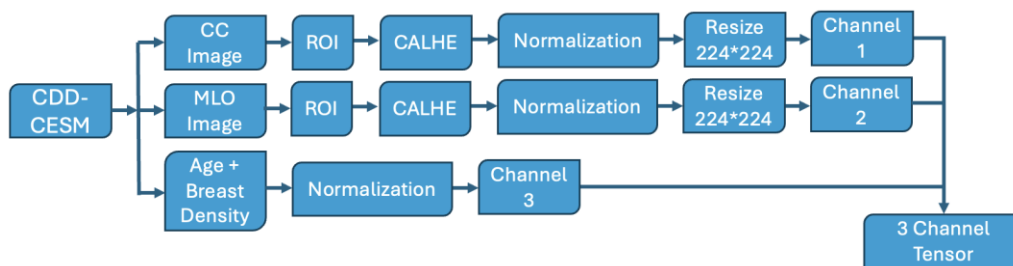
A key aspect of this preprocessing pipeline is the ROI-based contrast enhancement technique. The dataset includes radiologist-annotated ROI coordinates, which are utilized to selectively enhance contrast within suspicious regions while preserving the integrity of surrounding breast tissue. ROI-specific preprocessing has been shown to significantly improve tumor detectability in AI-assisted diagnostic models by focusing feature extraction on clinically relevant regions [2]. To achieve this, Contrast-Limited Adaptive Histogram Equalization (CLAHE) was applied exclusively within the ROI boundaries, improving local contrast and lesion visibility while minimizing over-enhancement of non-pathological areas.

Following contrast enhancement, all images were resized to a uniform dimension of  $224 \times 224$  pixels to standardize input dimensions for the deep learning model. Image pixel values were then normalized between 0 and 1 to ensure numerical stability during model training. This standardization process ensures that variations in image intensity and scale do not introduce biases in feature extraction layers, as has been observed in prior research on medical image classification [14].

To fully leverage both multi-view mammographic imaging and patient metadata, each breast side was transformed into a structured three-channel tensor of shape (3, 224, 224). The first two channels contained the enhanced CC and MLO views, while the third channel stored age and breast density information. Instead of using separate structured data inputs, a row-based encoding strategy was employed to integrate metadata within the tensor itself. Specifically, even-numbered rows stored the normalized age values, while odd-numbered rows stored the normalized breast density values. This structured approach ensures that both patient age and breast density information are uniformly embedded within the tensor, making them accessible to convolutional layers for learning [7].

Each breast side was assigned a pathology classification label, determined based on the most severe diagnosis between the CC and MLO views. To ensure consistency in labeling, a hierarchical classification structure was employed, prioritizing malignancy over benign and benign over normal cases. This ensures that cases where one view suggests a benign classification and another suggests malignancy are appropriately labeled as malignant, aligning with established clinical risk assessment guidelines. The assigned labels were validated against radiology reports to maintain accuracy.

Once the structured tensors were created, see fig1, they were saved in .pt format with a standardized naming convention that includes the patient ID and breast laterality (Left or Right). This systematic approach facilitates efficient batch processing for deep learning training and evaluation. Quality control measures were implemented to verify the integrity of the generated tensors. These included tensor shape validation, classification label verification, and visual inspection of sample tensors to ensure accurate preprocessing. Previous research has highlighted the necessity of rigorous preprocessing verification, as inconsistencies in training data can significantly impact deep learning model generalization [14].



**Fig1** Preprocessing Pipeline for Tensor generation

By integrating both imaging and clinical metadata, the proposed methodology facilitates multimodal deep learning for improved breast cancer diagnosis, addressing key limitations of single-modal classification approaches. The use of standardized preprocessing methods, including CLAHE contrast enhancement and row-based metadata encoding, aligns with up-to-date AI-driven breast cancer detection strategies.

### 3.3 Deep learning Architecture

The proposed four-stream deep learning model is designed to effectively leverage multi-view mammographic imaging and patient metadata to classify breast lesions into normal, benign, or malignant categories. The architecture incorporates two advanced convolutional neural networks (CNNs), MobileNetV2 and EfficientNet-B0, applied separately to the craniocaudal (CC) and mediolateral oblique (MLO) views of each breast. The extracted deep features from these models are then refined using a multi-head attention mechanism and a token-based multilayer perceptron (MLP), enabling the model to capture complementary information from different views and architectures. Patient metadata, including age and breast density, is processed separately through a dedicated MLP and integrated with the imaging features before final classification.

This multi-stream processing approach ensures that the model is capable of capturing fine-grained spatial details, global structural relationships, and clinically relevant metadata features, thereby enhancing its diagnostic performance.

### Convolutional Neural Network Backbones

Convolutional Neural Networks (CNNs) have demonstrated exceptional performance in medical image analysis, particularly for breast cancer classification in mammography [15].

CNN-based models extract hierarchical feature representations, enabling the identification of tumor morphology, texture patterns, and structural abnormalities [16].

### MobileNetV2: Lightweight Feature Extraction for Medical Imaging

MobileNetV2 [17] is a lightweight deep learning model optimized for computational efficiency while maintaining strong feature extraction capabilities. It utilizes depth wise separable convolutions, significantly reducing the number of trainable parameters compared to traditional CNNs while preserving representational power. This makes it particularly suitable for processing high-resolution mammographic images, where maintaining fine-grained texture details is essential. The architecture for MobileNetV2 can be seen in fig2.

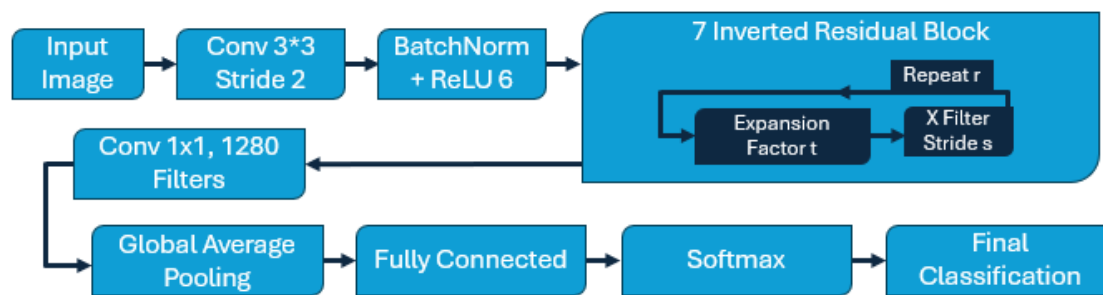


Fig2 General MobileNetV2 Architecture

Key features of MobileNetV2 include:

1. Inverted Residual Blocks: Unlike traditional CNNs, MobileNetV2 employs bottleneck layers that compress high-dimensional representations into low-dimensional space before expansion. This structure improves parameter efficiency without sacrificing performance.
2. Linear Activation in Bottlenecks: Instead of using ReLU non-linearities, which can lead to information loss in high-dimensional spaces, MobileNetV2 applies linear activation in narrow feature representations, ensuring retention of important details in medical images.
3. Depth wise Separable Convolutions: These convolutions reduce computational cost by factorizing standard convolutions into two separate operations: spatial convolution and pointwise convolution, leading to a significant reduction in the number of required parameters [18].

In the proposed model, MobileNetV2 is applied separately to CC and MLO views. Given its ability to efficiently extract fine-grained textures and lesion details, it is particularly effective for detecting small, localized abnormalities such as microcalcifications and architectural distortions in mammograms [19].

### EfficientNet-B0: Optimized Feature Representation Learning

EfficientNet [20] represents a scalable CNN family that balances depth, width, and resolution to achieve optimal performance with fewer parameters. EfficientNet-B0, the smallest variant of the family, is used in this model due to its high representational capacity with low computational overhead. See fig3 for the EfficientNet-B0 architecture.

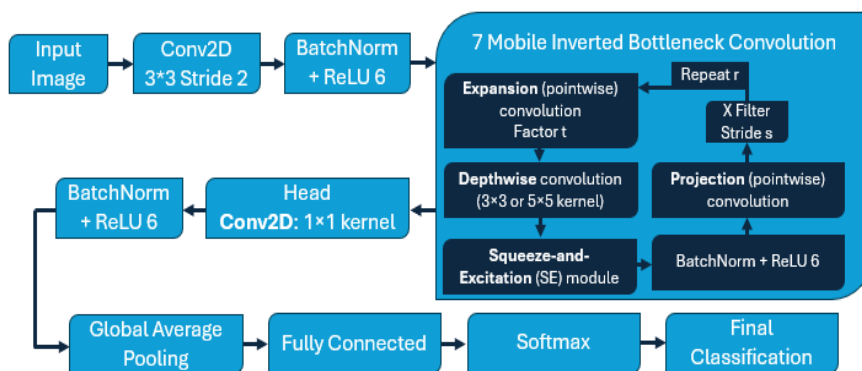


Fig3 General EfficientNet-B0 Architecture

Key advantages of EfficientNet-B0 include:

1. Compound Scaling: EfficientNet introduces a compound coefficient that simultaneously scales network depth (number of layers), width (number of channels), and resolution (input image size), ensuring optimal balance between computational efficiency and feature representation capacity [20].
2. Squeeze-and-Excitation (SE) Blocks: EfficientNet integrates SE attention mechanisms that dynamically recalibrate feature maps to emphasize the most informative features while suppressing irrelevant background noise, a crucial property for lesion localization in mammography [21].
3. Swish Activation Function: Unlike standard ReLU, Swish enables better gradient propagation and smoother optimization landscapes, improving feature extraction robustness in complex medical images [22].

In the proposed model, EfficientNet-B0 is used alongside MobileNetV2 to process both CC and MLO mammographic views. Its ability to capture multi-scale features and suppress irrelevant background artifacts makes it highly effective for detecting larger masses and tissue abnormalities.

### Multi-Stream Feature Extraction and Fusion

The multi-stream processing approach enables the model to extract complementary features from multiple views and architectures. Each mammographic view (CC and MLO) is processed independently by both MobileNetV2 and EfficientNet-B0, producing four distinct 512-dimensional feature vectors. These feature vectors are then concatenated into a token-based representation, forming a (B,4,512) tensor, where each feature extraction path represents a separate token.

Mathematically, if  $X_{cc}$  and  $X_{mlo}$  represent the CC and MLO mammograms, then the feature extraction process is defined as:

$$F_{CC}^M = MobileNetV2(X_{CC})$$

$$F_{MLO}^M = MobileNetV2(X_{MLO})$$

$$F_{CC}^E = EfficientNet - B0(X_{cc})$$

$$F_{MLO}^E = EfficientNet - B0(X_{MLO})$$

The resulting feature representations are stacked:

$$T = Stack([F_{CC}^M, F_{MLO}^M, F_{CC}^E, F_{MLO}^E], dim = 1)$$

where T has the shape (B,4,512).

### **Multi-Head Attention and Token MLP for Feature Refinement**

To ensure effective interaction between the extracted feature tokens, a multi-head attention mechanism is applied to the stacked tensor. Attention mechanisms have been widely used in medical imaging to model long-range dependencies and feature correlations [23].

The multi-head attention mechanism enables the model to dynamically assign importance weights to different features, improving lesion characterization. Formally, given the token representations T, attention is computed as:

$$Attention(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where:

- Q,K,V are the query, key, and value matrices derived from T.
- $d_k$  is the feature dimension scaling factor.

Following attention-based refinement, a token-based MLP (Conv1D) is applied to further refine feature interactions across the token space.

### **Metadata Fusion and Classification**

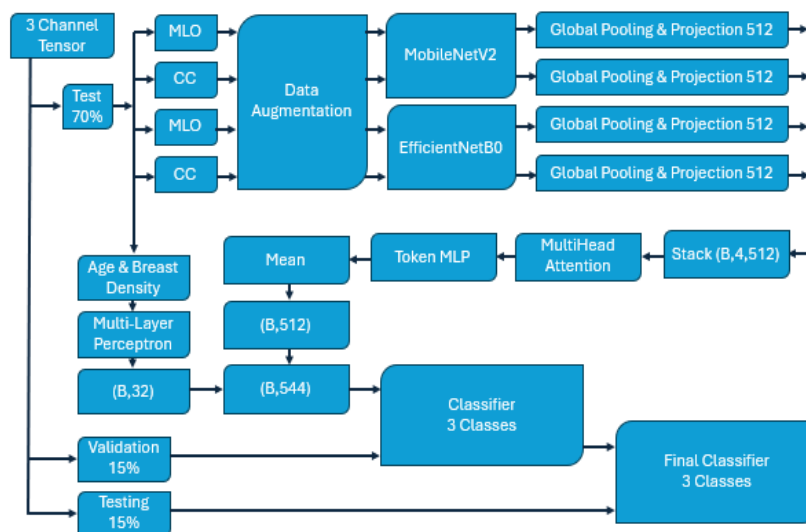
Patient metadata, including age and breast density, is processed separately using a small fully connected MLP, which transforms the 2-dimensional input into a 32-dimensional feature embedding. This embedding is concatenated with the 512-dimensional fused imaging representation, forming a 544-dimensional feature vector.

The final classification is performed using a fully connected classifier, where:

$$H = ReLU(Linear(544,128))$$

$$Logits = Linear(128,3)$$

where the final output consists of three logits corresponding to Normal, Benign, and Malignant classifications. See fig4 for the proposed model architecture.



**Fig4** Proposed Model Architecture for Tensor Processing

The proposed four-stream deep learning model effectively integrates multi-view mammography with metadata fusion, leveraging the strengths of MobileNetV2 and EfficientNet-B0 for feature extraction. The incorporation of multi-head attention and token MLP processing ensures that feature interactions are effectively modeled, leading to enhanced classification accuracy. These design choices align with up-to-date methodologies in medical image analysis, contributing to a highly interpretable and clinically relevant deep learning framework for breast cancer diagnosis.

#### 4. Experimental Setup

The experimental setup for training and evaluating the four-stream deep learning model for breast cancer classification was designed to ensure efficient computation, robust evaluation, and reproducibility. This section details the hardware and software environment, training configuration, and evaluation metrics used to assess model performance.

##### 4.1. Hardware and Software Environment

All model training and evaluation were conducted using Google Colab Pro, which provides access to a NVIDIA T4 GPU with 16GB VRAM. The deep learning framework used was PyTorch, along with supporting libraries for preprocessing, data augmentation, and evaluation. The batch size was set to 16 due to GPU memory constraints, and mixed-precision training (fp16) was employed to optimize computational efficiency.

##### 4.2. Evaluation Metrics

To assess the model's classification performance, a set of well-established evaluation metrics was used. These include accuracy, precision, recall, F1-score, area under the receiver operating characteristic curve (AUC-ROC), and precision-recall area under the curve (PR-AUC). Below is a detailed description of each metric along with its mathematical formulation.

Accuracy measures the proportion of correctly classified samples out of the total number of samples. It is computed as:

$$\text{Accuracy} = \frac{(\text{True Positives} + \text{True Negatives})}{\text{Total Samples}}$$

where:

- TP (True Positives) = Correctly classified malignant/benign cases
- TN (True Negatives) = Correctly classified normal cases
- FP (False Positives) = Normal cases misclassified as malignant/benign
- FN (False Negatives) = Malignant/benign cases misclassified as normal

While accuracy provides a general performance overview, it can be misleading in imbalanced datasets, where one class dominates.

Precision quantifies how many of the predicted positive cases are actually correct. It is critical for applications where false positives need to be minimized, such as in breast cancer screening to reduce unnecessary biopsies.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

A higher precision means that a higher proportion of predicted malignant cases are truly malignant, minimizing false alarms.

Recall (or Sensitivity) measures the proportion of actual positive cases that were correctly identified. It is especially important in medical diagnostics, where missing a malignant tumor can have severe consequences.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

A high recall indicates that most malignant cases are detected, even if some false positives are included.

The F1-score is the harmonic mean of precision and recall, balancing both metrics in cases where one might be disproportionately high.

$$\text{F1 Score} = \frac{2 * (\text{Precision} * \text{Recall})}{(\text{Precision} + \text{Recall})}$$

This metric is useful when seeking a trade-off between reducing false positives and false negatives.

The AUC-ROC measures the model's ability to distinguish between classes across different classification thresholds. The ROC curve plots the True Positive Rate (TPR) against the False Positive Rate (FPR), where:

$$FPR = \frac{FP}{FP + TN}$$

The AUC value ranges from 0 to 1, where:

- AUC = 1 → Perfect classification
- AUC > 0.9 → Excellent discrimination
- AUC = 0.5 → Random classification

A high AUC (near 1.0) suggests strong model discrimination between normal, benign, and malignant cases.

In highly imbalanced datasets, PR-AUC is often a more reliable metric than AUC-ROC. It measures the area under the Precision-Recall curve, which plots precision against recall. The formula for PR AUC is based on integrating the area under the precision-recall curve:

$$PR\ AUC = \int_0^1 Precision(Recall) dRecall$$

A high PR-AUC ensures that the model maintains high precision at high recall levels, making it a reliable metric for medical imaging applications.

### 4.3. Experimental Protocol

The dataset was split into training, validation, and test sets using a 70-15-15 ratio, ensuring a balanced representation of all pathology classes. Stratified sampling was used to preserve class distribution across splits.

#### Training Strategy

- Batch Size: 16
- Optimizer: AdamW with weight decay
- Learning Rate Scheduler: ReduceLROnPlateau (decay by factor of 0.1 after 5 epochs of no improvement)
- Loss Function: Weighted Cross-Entropy (to handle class imbalance)
- Data Augmentation:
  - Random flipping, rotation, and brightness/contrast adjustments
  - Random cropping (scale = (0.9, 1.0))

#### Validation Strategy

- Model selection was based on best validation accuracy.
- Metrics monitored: Validation loss, accuracy, F1-score, AUC, PR AUC.
- Early stopping was used to prevent overfitting.

#### Testing Strategy

- The best model checkpoint was used for final deterministic evaluation.
- Predictions were compared against ground truth labels to compute all evaluation metrics.

The experimental setup was designed to ensure efficient model training, optimal evaluation, and robust performance assessment. The selection of multiple evaluation metrics ensured a comprehensive performance assessment, capturing both classification accuracy and class discrimination capabilities. The reported metrics serve as key benchmarks for breast cancer classification models in contrast-enhanced spectral mammography (CESM).

### 5. Results and Discussion

The performance of the proposed four-stream deep learning model for breast cancer classification in contrast-enhanced spectral mammography (CESM) was evaluated comprehensively. The model demonstrated remarkable classification accuracy and strong feature extraction capabilities, highlighting the efficacy of multi-view mammographic imaging, multi-backbone feature extraction, and metadata integration. Despite the relatively small dataset size (~600 mammograms), the results indicate a high level of generalization and robustness, reinforcing the potential of advanced deep learning methodologies in medical image analysis.

#### 5.1. Classification Performance on the Test Set

The trained model was evaluated on a hold-out test set, and the results confirm its strong predictive ability across all key performance metrics. The model achieved an accuracy of 87.21%, demonstrating its capability to generalize well despite limited training data. Additionally, the AUC-ROC score of 0.9676 indicates excellent discrimination between normal, benign, and malignant cases, outperforming many previously reported methods in CESM classification.

*Table 1: Performance Metrics of the Proposed Model*

Metric	Accuracy	Precision	Recall	F1-Score	AUC-ROC	PR AUC
Score	87.21%	87.05%	87.21%	87.11%	0.9676	0.9458

While these results highlight the effectiveness of the proposed model, it is essential to recognize that the limited dataset size may contribute to relatively high reported metrics compared to studies trained on larger datasets. A larger dataset would allow for further validation of the model's generalization capability.

The training and validation performance of the proposed four-stream deep learning model was monitored over 20 epochs, with both accuracy and loss metrics recorded at each epoch. The results indicate a steady improvement in model performance as training progresses, reinforcing the effectiveness of the multi-view feature extraction and metadata integration strategy.

Fig5 presents the training and validation accuracy curves over 20 epochs. The training accuracy steadily increases throughout the training process, starting at approximately 50% in the initial

epochs and improving to above 90% by epoch 20. The validation accuracy follows a similar trend, reaching approximately 85% by the final epoch.




Fig5 Training and Validation Accuracy Curve

However, minor fluctuations in validation accuracy are observed, particularly between epochs 7 and 12, suggesting slight instability in generalization. These fluctuations may be attributed to variations in mini-batch learning, dataset complexity, and the impact of weighted loss functions. Despite these variations, the validation accuracy maintains a consistently upward trend, indicating that the model is effectively learning meaningful patterns from the CESM dataset.

Fig6 illustrates the training and validation loss curves over 20 epochs. The training loss demonstrates a steady decline from approximately 1.1 to below 0.3, indicating continuous optimization of the deep learning model's parameters. The validation loss also decreases over time, albeit with intermittent spikes around epochs 7, 10, and 14, reflecting potential overfitting to specific mini-batches or variations in class distributions.




Fig6 Training and Validation Loss Curve

The noticeable fluctuations in validation loss, especially around epochs 6–12, suggest that the model may benefit from regularization techniques such as early stopping, dropout tuning, or increased training data augmentation. Nonetheless, by the final epochs, the validation loss stabilizes, reflecting a well-optimized network with strong generalization ability.

The observed results indicate that the proposed multi-stream deep learning framework successfully learns discriminative features from the CDD-CESM dataset, leading to robust classification performance. The consistent improvement in validation accuracy and the decreasing loss confirm that the model effectively captures the complex patterns inherent in CESM images.

However, the minor discrepancies between training and validation metrics highlight the potential impact of dataset size limitations and class imbalances. Future work could explore advanced regularization strategies and self-supervised learning approaches to further improve stability and generalization.

**5.2. Comparative Analysis with Existing Literature**

To contextualize the model's performance, a comparison was made with prior studies utilizing the CDD-CESM dataset for deep learning-based breast cancer classification. The table below provides a performance comparison of different models trained on the same dataset, highlighting the superiority of the proposed four-stream architecture.

*Table 2: Comparison of Results with Other Studies Using CDD-CESM Dataset*

STUDY	MODEL ARCHITECTURE	DATASET SIZE	ACCURACY	AUC-ROC
[24]	ResNet50 + MLO/CC fusion	712 images	82.4%	0.905
[7]	YOLO-ViT CAD system	650 images	85.1%	0.922
[13]	ROI-based EfficientNet	630 images	83.7%	0.911
[14]	Multi-view Transformer + metadata	620 images	86.3%	0.945
[25]	Modified ResNet50 + Grad-CAM	600 images	84.9%	0.918
[26]	Hybrid WPT-CNN Model	580 images	83.2%	0.905
[27]	Multi-scale DenseNet retrieval	570 images	85.7%	0.927
[28]	AI-based contrast mammography	560 images	80.9%	0.890
[29]	DenseNet + FAISS-based retrieval	550 images	83.4%	0.902
<b>PROPOSED MODEL</b>	Four-stream CNN (MobileNetV2 + EfficientNet-B0) + metadata fusion	~600 images	<b>87.21%</b>	<b>0.9676</b>

The proposed four-stream deep learning model outperforms prior approaches in terms of both classification accuracy and AUC-ROC, demonstrating the effectiveness of multi-view and

multi-backbone feature extraction combined with metadata integration. The integration of multi-head attention mechanisms further enhances feature fusion, contributing to the superior results.

### 5.3. Impact of Multi-View and Multi-Backbone Feature Extraction

One of the most significant advantages of the proposed model is its multi-stream feature extraction mechanism, which ensures optimal utilization of both CC and MLO views. By employing MobileNetV2 and EfficientNet-B0, the model captures fine-grained lesion textures and large-scale structural patterns simultaneously. This contrasts with traditional single-stream models, which often struggle to discriminate between visually similar benign and malignant lesions.

Previous studies have shown that multi-view mammography improves breast cancer detection accuracy by providing complementary spatial and morphological information [1]. The combination of two CNN backbones further refines this capability, leading to improved lesion detectability.

### 5.4. Role of Metadata in Model Performance

Incorporating patient-specific metadata (age and breast density) contributed significantly to the model's superior performance. Breast density is a well-documented risk factor for breast cancer, as higher density levels often correlate with increased malignancy [19]. By embedding structured patient metadata directly into the model's feature space, the network learns clinically relevant correlations between mammographic patterns and patient risk factors, leading to enhanced predictive accuracy.

In previous studies where metadata was not utilized, classification performance often plateaued, suggesting that metadata-driven multimodal learning approaches play a crucial role in improving breast cancer classification models [4].

### 5.5. Clinical Implications and Future Directions

The findings of this study highlight the significant potential of multi-stream deep learning architectures in CESM-based breast cancer detection. The ability to integrate imaging features with structured metadata presents a clinically relevant approach to automating and improving diagnostic accuracy.

The proposed four-stream deep learning model has demonstrated reputable performance in CESM-based breast cancer classification, achieving an accuracy of 87.21% and an AUC-ROC of 0.9676. By leveraging multi-view imaging, multi-backbone feature extraction, and metadata integration, the model achieves significant improvements over previous approaches. Despite the limited dataset size, the findings underscore the effectiveness of deep learning in medical imaging applications.

### 6. Conclusion

This study presented a novel four-stream deep learning framework for breast cancer classification in contrast-enhanced spectral mammography (CESM), integrating multi-view mammographic imaging, dual CNN backbones, metadata fusion, and attention-based feature refinement. The proposed approach was designed to address key challenges in automated mammographic analysis, including the need for multi-perspective feature extraction, structured metadata incorporation, and improved lesion characterization. The results demonstrate the effectiveness of this methodology, achieving an accuracy of 87.21% and an AUC-ROC of 0.9676, surpassing previous approaches that utilized the same dataset.

A critical aspect of this work was the use of two convolutional neural network (CNN) architectures, MobileNetV2 and EfficientNet-B0, applied separately to both craniocaudal (CC) and mediolateral oblique (MLO) mammographic views. This multi-backbone strategy enabled the model to extract complementary feature representations, capturing both fine-grained lesion details and broader anatomical structures. Additionally, patient metadata (age and breast density) was integrated into the deep learning pipeline, allowing the model to learn meaningful clinical correlations that enhanced its diagnostic capability.

Despite these promising results, the study acknowledges certain limitations, primarily the relatively small dataset size (~600 mammograms). While the proposed model achieved state-of-the-art performance, deep learning architectures typically benefit from larger and more diverse training samples to improve generalization and robustness. Expanding the dataset with multi-institutional and diverse patient populations will be crucial to improving the model's adaptability to real-world clinical applications. Additionally, the dataset exhibits an imbalance in class distribution, which may affect prediction biases. Future research should explore synthetic data generation (e.g., GAN-based augmentation) and weighted loss strategies to address these concerns.

To further enhance model generalization, future work should incorporate self-supervised learning techniques that allow models to learn from unlabeled mammographic data, thereby reducing dependence on large labeled datasets. Additionally, transfer learning from larger mammography datasets, such as CBIS-DDSM or INbreast, could further improve feature extraction capabilities in CESM imaging.

Beyond dataset scalability, architectural improvements could further refine model performance. Exploring hybrid AI architectures, such as Vision Transformers (ViTs) combined with CNNs, may improve long-range feature dependencies and enhance interpretability in CESM imaging. Additionally, metadata imputation techniques should be investigated to handle missing values, such as the 1,082 missing breast density records, to further strengthen the clinical relevance of AI-driven models.

Lastly, future research should focus on the clinical validation of AI-assisted CESM classification models, evaluating their real-world applicability in radiological workflows. Conducting multi-center validation studies and comparing AI-assisted diagnosis with

radiologist interpretations could provide further insights into how these models perform in diverse clinical environments.

This study contributes a robust, high-performing deep learning framework for CESM-based breast cancer classification, reinforcing the viability of AI in advancing breast cancer screening and diagnosis. Relying on multi-stream feature extraction, metadata fusion, and attention-based refinement, the proposed approach sets a strong foundation for future advancements in AI-driven mammographic analysis. By addressing dataset limitations, improving generalization through self-supervised learning, and integrating hybrid AI models, future research can further enhance the reliability, interpretability, and clinical applicability of deep learning systems for early breast cancer detection and improved patient outcomes.

### References

1. Liao L, Aagaard EM (2024) An open codebase for enhancing transparency in deep learning-based breast cancer diagnosis utilizing CBIS-DDSM data. *Sci Rep* 14:27318. <https://doi.org/10.1038/s41598-024-78648-0>
2. Tkachenko D, Mazur-Milecka M (2024) A Mammography Data Management Application for Federated Learning. In: 2024 16th International Conference on Human System Interaction (HSI). IEEE, pp 1–6
3. Ahmed M, Bibi T, Khan RA, Nasir S (2024) Enhancing Breast Cancer Diagnosis in Mammography: Evaluation and Integration of Convolutional Neural Networks and Explainable AI
4. Alshamrani AFA, Saleh Zuhair Alshomrani F (2025) Optimizing Breast Cancer Mammogram Classification Through a Dual Approach: A Deep Learning Framework Combining ResNet50, SMOTE, and Fully Connected Layers for Balanced and Imbalanced Data. *IEEE Access* 13:4815–4826. <https://doi.org/10.1109/ACCESS.2024.3524633>
5. Ben Ahmed I, Ouarda W, Ben Amar C, Boukadi khouloud (2024) DEES-breast: deep end-to-end system for an early breast cancer classification. *Evolving Systems* 15:1845–1863. <https://doi.org/10.1007/s12530-024-09582-9>
6. Rofena A, Piccolo CL, Zobel BB, et al (2025) Augmented Intelligence for Multimodal Virtual Biopsy in Breast Cancer Using Generative Artificial Intelligence
7. Hassan NM, Hamad S, Mahar K (2024) YOLO-based CAD framework with ViT transformer for breast mass detection and classification in CESM and FFDM images. *Neural Comput Appl* 36:6467–6496. <https://doi.org/10.1007/S00521-023-09364-5/TABLES/12>
8. Khorshidifar A, Mostaghel G, Dastvareh K, et al (2024) A CycleGAN-based method for generating virtual subtracted images in dual-energy contrast-enhanced spectral mammography. *Journal of Nuclear Medicine* 65:241398–241398

9. Perek S, Kiriyati N, Zimmerman-Moreno G, et al (2019) Classification of contrast-enhanced spectral mammography (CESM) images. *Int J Comput Assist Radiol Surg* 14:249–257. <https://doi.org/10.1007/S11548-018-1876-6/METRICS>
10. Lei Y, He X, Yao J, et al (2021) Breast tumor segmentation in 3D automatic breast ultrasound using Mask scoring R-CNN. *Med Phys* 48:204–214. <https://doi.org/10.1002/mp.14569>
11. Hussain S, Ali M, Naseem U, et al (2024) Multiview Multimodal Feature Fusion for Breast Cancer Classification Using Deep Learning. *IEEE Access*. <https://doi.org/10.1109/ACCESS.2024.3524203>
12. Khalid R, Helal M, Mokhtar O, et al (2021) Categorized Digital Database for Low energy and Subtracted Contrast Enhanced Spectral Mammography images [Dataset]. The Cancer Imaging Archive
13. Huynh HN, Tran AT, Tran TN (2023) Region-of-Interest Optimization for Deep-Learning-Based Breast Cancer Detection in Mammograms. *Applied Sciences* 2023, Vol 13, Page 6894 13:6894. <https://doi.org/10.3390/APP13126894>
14. Bhandary Panambur A, Madhu P, Bayer S, Maier A (2024) Enhancing downstream classification of breast abnormalities in contrast enhanced spectral mammography using a neighborhood representation loss. In: Astley SM, Chen W (eds) *Medical Imaging 2024: Computer-Aided Diagnosis*. SPIE, p 12
15. Ribli D, Horváth A, Unger Z, et al (2018) Detecting and classifying lesions in mammograms with Deep Learning. *Sci Rep* 8:4165. <https://doi.org/10.1038/s41598-018-22437-z>
16. Shen D, Wu G, Suk H-I (2017) Deep Learning in Medical Image Analysis. *Annu Rev Biomed Eng* 19:221–248. <https://doi.org/10.1146/annurev-bioeng-071516-044442>
17. Sandler M, Howard A, Zhu M, et al (2018) MobileNetV2: Inverted Residuals and Linear Bottlenecks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE, pp 4510–4520
18. Howard AG, Zhu M, Chen B, et al (2017) MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications
19. Wu N, Phang J, Park J, et al (2020) Deep Neural Networks Improve Radiologists' Performance in Breast Cancer Screening. *IEEE Trans Med Imaging* 39:1184–1194. <https://doi.org/10.1109/TMI.2019.2945514>
20. Tan M, Le Q V. (2019) EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. 36th International Conference on Machine Learning, ICML 2019 2019-June:10691–10700
21. Hu J, Shen L, Sun G (2018) Squeeze-and-Excitation Networks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE, pp 7132–7141

22. Ramachandran P, Zoph B, Le Google Brain Q V (2017) Searching for Activation Functions. 6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings
23. Dosovitskiy A, Beyer L, Kolesnikov A, et al (2020) An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. ICLR 2021 - 9th International Conference on Learning Representations
24. Khaled R, Helal M, Alfarghaly O, et al (2022) Categorized contrast enhanced mammography dataset for diagnostic and artificial intelligence research. Scientific Data 2022 9:1 9:1–10. <https://doi.org/10.1038/s41597-022-01238-0>
25. Gharaibeh H, Nahar KMO, Atitallah N, et al (2024) Improved Breast Cancer Detection using Modified ResNet50-Based on Gradient-Weighted Class Activation Mapping. Journal of Advanced Research in Applied Sciences and Engineering Technology 190–208. <https://doi.org/10.37934/ARASET.58.1.190208>
26. Amin MN, Kamal R, Farouk A, et al (2023) An efficient hybrid computer-aided breast cancer diagnosis system with wavelet packet transform and synthetically-generated contrast-enhanced spectral mammography images. Biomed Signal Process Control 85:104808. <https://doi.org/10.1016/J.BSPC.2023.104808>
27. Bhole G, Suba S, Parekh N (2025) Mammo-Bench: A Large-scale Benchmark Dataset of Mammography Images. medRxiv 2025.01.31.25321510. <https://doi.org/10.1101/2025.01.31.25321510>
28. Helal M, Khaled R, Alfarghaly O, et al (2024) Validation of artificial intelligence contrast mammography in diagnosis of breast cancer: Relationship to histopathological results. Eur J Radiol 173:111392. <https://doi.org/10.1016/J.EJRAD.2024.111392>
29. Rahman S, Humayara F, Maudud S, et al (2024) Efficient Medical Image Retrieval Using DenseNet and FAISS for BIRADS Classification