ConvNeXt for Breast Cancer Classification in Infrared Images

Víctor Emmanuel Olmi Louzada[1], **Matheus de Freitas Oliveira Baffa**[2]

[1]Department of Informatics (DPI)
Federal University of Viçosa (UFV)
Viçosa, MG – Brazil

[2]Institute of Mathematical Sciences and Computing (ICMC)
University of São Paulo (USP)
São Carlos, SP – Brazil

`victor.louzada@ufv.br, mbaffa@usp.br`

Abstract. *Breast cancer is a major public health concern, highlighting the need for accurate and accessible diagnostic tools. Infrared imaging offers a non-invasive alternative for detecting physiological abnormalities associated with tumor development. In this study, we evaluate the use of ConvNeXt architectures for breast cancer classification based on static infrared images. A patient-level stratified cross-validation protocol was adopted to ensure robust and unbiased assessment. Among the evaluated models, ConvNeXt-Large achieved the best performance, reaching an accuracy of 90.43% and an AUC of 98.29%. These results demonstrate the potential of modern deep learning architectures for supporting breast cancer detection using infrared imaging.*

1. Introduction

Breast cancer is a disease characterized by the uncontrolled proliferation of abnormal cells within breast tissue, leading to the formation of tumors that may invade surrounding structures and potentially metastasize to distant organs. It represents a major global public health concern, having caused an estimated 670,000 deaths worldwide in 2022. In the same year, breast cancer was the most frequently diagnosed cancer among women in 157 out of 185 countries, highlighting its widespread impact across diverse populations [World Health Organization 2024].

In Brazil, the burden of the disease is also significant, with an estimated 73,610 new cases expected annually for the 2023 to 2025 triennium, corresponding to an adjusted incidence rate of 41.89 cases per 100,000 women [Instituto Nacional do Câncer 2024]. These statistics underscore the critical importance of effective screening, early detection, and accurate diagnostic methods to improve patient outcomes.

Breast cancer may present with a variety of physical changes, although the most commonly recognized sign is the appearance of a new lump or localized swelling in the breast, chest, or underarm region. However, the disease is not limited to this manifestation and can also involve alterations in breast morphology, such as changes in size or shape, as well as modifications in skin texture, including dimpling or redness [Koo et al. 2017]. Additional warning signs may include abnormal nipple discharge unrelated to breastfeeding, inversion or distortion of the nipple, and persistent pain in the breast or axillary area [National Health Service 2026].

The early disease detection relies mostly on imaging techniques that enable the identification of abnormalities before the first symptoms. Screening methods, particularly mammography, are widely used to detect early signs of the disease, even in asymptomatic individuals, improving the chances of effective treatment. When abnormalities are identified, additional diagnostic procedures such as diagnostic mammography, ultrasound, magnetic resonance imaging (MRI) or even biopsy, are employed to further characterize suspicious regions and confirm the presence of cancer [Evans et al. 2018].

In this context, infrared thermal imaging is a non-invasive technique that captures the heat naturally emitted by the human body, allowing the visualization of temperature distributions across the skin surface [D'Alessandro et al. 2024]. Since physiological processes such as blood flow, metabolism, and inflammation directly influence thermal patterns, this modality has been increasingly explored in biomedical applications for the detection of abnormalities. In recent years, the integration of computational analysis, particularly machine learning and deep learning methods, has significantly enhanced the interpretation of thermal images by enabling the automatic extraction of complex patterns that may not be visually perceptible [Mashekova et al. 2022, D'Alessandro et al. 2024].

Recent advances in deep learning have led to a growing interest in Vision Transformers (ViTs) for the analysis of thermal images in medical applications. Unlike traditional convolutional neural networks (CNNs), ViTs are capable of modeling long-range dependencies and capturing global contextual information, which is particularly relevant for thermographic data where subtle and spatially distributed temperature variations may indicate underlying physiological changes. This trend is reflected in recent studies that explore transformer-based approaches for breast cancer detection in thermal imaging, including the use of ViTs for vasodilation pattern analysis [Rodd 2026], multimodal vision-language frameworks [Pinto et al. 2025], Swin Transformer architectures [Hayat 2023], and multi-view thermal analysis [Cihan and Ceylan 2025].

Despite the growing adoption of ViTs in thermal imaging, the exploration of hybrid architectures that combine convolutional inductive biases with modern transformer-inspired design remains limited in this domain. In particular, ConvNeXt has emerged as a strong alternative, redesigning conventional convolutional neural networks using principles derived from Vision Transformers, such as large kernel sizes, improved normalization strategies, and simplified architectural blocks. These modifications enable ConvNeXt to achieve competitive performance with transformer-based models while maintaining the efficiency and stability of convolutional operations, especially in scenarios with limited data [Liu et al. 2022]. However, its application to infrared breast imaging has not been thoroughly investigated. This gap motivates the present study, which evaluates ConvNeXt as a potential architecture for breast cancer classification, aiming to assess whether its design can effectively capture the thermal patterns associated with malignancy.

Therefore, this study aims to evaluate the effectiveness of ConvNeXt architectures for breast cancer classification using infrared thermal images. Specifically, we investigate three variants in order to analyze the impact of model capacity on classification performance: ConvNeXt-Small, ConvNeXt-Base, and ConvNeXt-Large. By comparing these architectures under a patient-level validation protocol, the study seeks to determine whether increasing representational complexity leads to improved discrimination of malignant and healthy patterns in thermographic breast imaging.

2. Related Work

Over the literature, infrared imaging has been extensively explored by diverse methods, including traditional computer vision approaches based on feature engineering, as well as custom and pretrained CNNs. This section presents recent approaches for breast cancer classification using thermographic images and their respective performance.

[Baffa and Conci 2022] investigated a radiomics-based approach, extracting handcrafted features from thermographic images and combining them with a fully-connected neural network for classification. Their method achieved high performance, with 98.61% accuracy, precision, and recall, as well as an AUC of 99.14%, reinforcing the relevance of radiomic features in breast cancer detection.

In 2023, [Hayat 2023] explored the use of transformer-based architectures by applying a Swin Transformer model to thermal breast images. Their method achieved competitive results, including 96.55% accuracy, 95.50% precision, 95.76% recall, 95.43% F1-score, 97.34% specificity, and an AUC of 96.21%, demonstrating the viability of transformer models for this task.

Later on, [Nguyen Chi et al. 2025] proposed a lightweight hybrid framework combining a ResNet34 backbone for feature extraction with a support vector machine classifier. This method achieved outstanding performance, with accuracy, precision, recall, and F1-score all exceeding 99%, indicating the effectiveness of combining deep features with classical machine learning classifiers.

In 2025, [Alzahrani et al. 2025] introduced a hybrid method that integrates convolutional neural networks with particle swarm optimization to enhance feature selection and model performance. Their approach achieved an accuracy of 97.5%, with high precision (96.8%), recall (98.4%), and F1-score (97.6%), demonstrating the potential of optimization techniques in improving CNN-based models.

Recently, [Aggarwal et al. 2026] proposed an ensemble-based classifier that combines multiple models to improve robustness and predictive performance. Their approach achieved strong results, reporting an accuracy of 97.10%, sensitivity of 97.52%, specificity of 96.73%, and an AUC of 99.40%, highlighting the effectiveness of ensemble learning in thermographic breast cancer detection.

Despite the strong performance reported in the literature, most existing approaches rely on either handcrafted radiomic features, conventional CNN-based pipelines, or more complex hybrid and ensemble strategies. While these methods achieve high accuracy, they often involve multiple processing stages or require additional optimization procedures, which may limit their generalizability and reproducibility. Moreover, although transformer-based models have recently been introduced, the exploration of modern convolutional architectures inspired by transformer design, such as ConvNeXt, remains limited in the context of breast infrared imaging. In particular, there is a lack of systematic evaluation of different ConvNeXt variants under a patient-level validation protocol.

3. Materials and Methods

This section describes the materials and methods employed to evaluate the proposed approach for breast cancer classification using infrared images. Initially, the dataset and its characteristics are presented, followed by the definition of the classification methodology, including data augmentation strategies, adaptation of the classification layer, and the optimization and training procedures. In addition, the performance evaluation framework is detailed, incorporating both standard metrics and threshold calibration techniques to ensure a robust assessment. Figure 1 illustrate the pipeline.

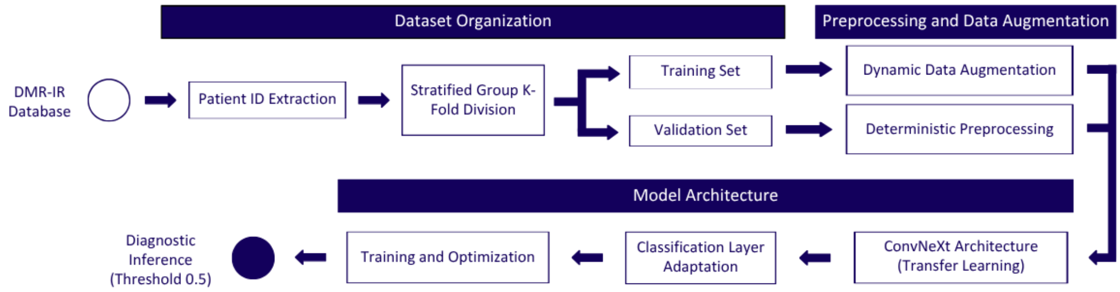

Figure 1. Overview of the proposed classification pipeline.

3.1. Dataset Description

The dataset used in this study is the publicly available Database for Mastology Research on Infrared Imaging (DMR-IR), developed by Federal Fluminense University (UFF) [Silva et al. 2014]. This database contains breast infrared images acquired under both static and dynamic protocols (Fig. 2).

In the static acquisition protocol, five images are captured per patient, consisting of one frontal view and four lateral views obtained at angles of 45° and 90° to the left and right. In the dynamic protocol, a sequence of 20 images is acquired over time following a thermal stimulus, enabling the analysis of temperature variations.

This work focuses on the use of frontal colored images from the static protocol for breast cancer classification, with the size of 640×480 pixels. This choice provides a simplified and consistent representation, using a single image per patient while preserving relevant thermal patterns. The dataset comprises a total of 146 images, with 73 samples from cancer cases and 73 from healthy individuals, resulting in a balanced binary classification scenario.

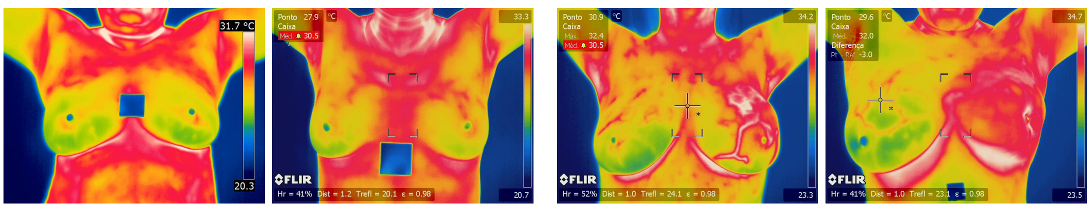

(a) (b)

Figure 2. Examples of infrared breast images from the dataset: (a) healthy cases and (b) cancer cases.

3.2. Classification Methodology

To conduct an analysis focused on the efficacy of thermal feature extraction, the study established a systematic evaluation of a state-of-the-art architecture based on transfer learning, pre-trained on ImageNet: ConvNeXt. Aiming to investigate the relationship between parametric complexity and the model's generalization capacity for the task, the architecture was tested across three scales of depth and width: small, base, and large.

As detailed in Table 1, the intrinsic dimensionality of the output tensors from these networks varies substantially, producing spatial embeddings that range from 768 to 1536 channels. Due to this variation, the methodology incorporated a dynamic instantiation strategy. Using predefined mappings, the algorithm automatically replaces the final layer of the neural network so that it adjusts to the dimensions of the chosen base model (backbone). This ensures proper alignment between the features extracted by the convolutional layers and the final classification stage. The training of the models was executed on graphics processing units (GPUs) with CUDA support. To avoid data transfer bottlenecks between the CPU and the GPU during training, memory pinning (`pin_memory`) was enabled in the data loaders.

Table 1. Output dimensions of the ConvNeXt architecture variants.

Architecture Family	Model Variant	Technical Nomenclature	Output Dimension
ConvNeXt	Small	`convnext_small`	768
	Base	`convnext_base`	1024
	Large	`convnext_large`	1536

3.2.1. Sample Selection and Preparation Criteria

To ensure the statistical independence of the samples and prevent the model's learning from being biased, the indexing algorithm was programmed to process all valid images in the class directories, extracting the unique patient identifier (*patient ID*) directly from the original file nomenclature. Each selected image was mapped to this identifier, forming sample groups. This tracking was the central mechanism that allowed the stratified group k-fold algorithm ($k = 10$) to perform the dataset division, ensuring that all images from the same individual were mutually exclusively assigned to either the training or validation set, mitigating the risk of data leakage. Specifically, in each fold, 90% of the data is used for training and 10% for testing.

3.2.2. Data Augmentation

To increase the model's robustness against noise, handle positional variations, and mitigate overfitting, we applied a stochastic data augmentation pipeline exclusively to the training set. It is noteworthy that these transformations were generated dynamically at runtime, keeping the physical size of the dataset unchanged.

Initially, the images were resized to an intermediate resolution of 256×256 pixels, followed by a random resized crop to the final resolution of 224×224 pixels, with the preserved area scale ranging between 85% and 100%. The morphological and color transformations included: horizontal flipping with a 50% probability; subtle affine transformations applied to 45% of the samples (involving rotations up to 7°, translations up to 3% on the axes, and a scaling factor between 0.97 and 1.03); and color variations (color jitter) with a 35% probability (altering brightness and contrast by up to 12%, saturation by 8%, and hue by 2%).

Additionally, the random erasing technique was employed in 20% of the iterations, occluding rectangular portions with random noise (covering 1% to 6% of the total image area, with an aspect ratio between 0.4 and 2.4), forcing the neural network to extract pathological features from multiple anatomical regions rather than memorizing a single local pattern. In contrast, the validation set images were subjected only to the preliminary resizing (256×256), deterministic central cropping (center crop) to 224×224, and standardized statistical normalization using the ImageNet RGB channel values (means: 0.485, 0.456, 0.406; standard deviations: 0.229, 0.224, 0.225).

3.2.3. Classification Head Adaptation

An essential methodological step of this work consisted of adapting the final classification head of the ConvNeXt models. Originally, these transfer learning-based architectures are pre-trained on large databases of generic images (such as ImageNet) and structured with an output layer sized to predict thousands of distinct classes. To adapt the models to the specific context of binary thermographic diagnosis, the network topology required a minor architectural modification.

Specifically, the original linear prediction layer of each ConvNeXt variant was replaced by a new linear layer configured to return a single scalar value (logit). This architectural simplification allows the network to act directly as a binary classifier. The model relies on the global average pooling native to its own architecture to condense the spatial feature maps extracted from the thermal images into a representative one-dimensional vector.

The optimization of this new topology was conducted using the binary cross-entropy with logits loss function. This approach was selected for ensuring greater numerical stability during training, as it integrates the sigmoid activation function and the binary cross-entropy into a single optimized mathematical operation. Therefore, the atypical thermal gradients extracted by the deep convolutional layers are mapped directly to a diagnostic probability, enabling the network to robustly differentiate the thermal signatures associated with healthy tissue from those indicative of anomalies.

3.2.4. Optimization and Training Strategy

The weight update during training was conducted by the AdamW optimizer. The choice of this optimizer is justified by its decoupled weight decay formulation. Unlike traditional Adam, which couples L2 regularization to the adaptive gradient calculation, attenuating its effect, AdamW applies the penalty directly to the weights. This mathematical separation is essential for the fine-tuning of architectures with high parametric capacity, such as ConvNeXt, on restricted medical datasets, as it maximizes the efficacy of the regularization and reduces the risk of overfitting. To ensure stable convergence toward the global minimum of the loss function, the learning rate was dynamically modulated throughout the epochs using a cosine annealing scheduler. Finally, we employed the binary cross-entropy with logits loss function, as previously described, due to its numerical stability.

4. Experiments and Results

The experiments were implemented in Python 3.14 using PyTorch 2.11 as the primary deep learning framework. Image processing operations were performed with OpenCV 4.12, while numerical computations and data handling were supported by NumPy 2.4 and scikit-learn 1.8. All experiments were conducted on a high-performance workstation equipped with an Intel Core i9 processor, 128 GB of DDR5 RAM, an 8 TB SSD for data storage, and an NVIDIA RTX 4090 GPU, ensuring efficient training and evaluation of the models.

Model performance was evaluated using accuracy, precision, F1-score, sensitivity, specificity, and the area under the receiver operating characteristic curve (ROC-AUC). Accuracy measures overall correctness, while precision and sensitivity assess the detection of positive cases and control of false positives; the F1-score balances these metrics, and specificity quantifies the correct identification of negative cases. ROC-AUC summarizes the model's discriminative ability across all decision thresholds. For performance estimation, a stratified group k-fold cross-validation protocol with 10 folds ($k = 10$) was adopted, preserving class proportions while enforcing patient-level grouping to prevent data leakage between training and validation sets.

4.1. Results

The quantitative results indicate that the ConvNeXt-Large variant achieved the most consistent overall performance among the evaluated models, obtaining the highest values in key metrics such as accuracy (90.43%), F1-score (90.25%), and sensitivity (90.30%). In addition to superior predictive performance, this variant also demonstrated greater stability across the cross-validation folds, as reflected by its lower standard deviations, particularly in accuracy ($\pm$ 6.22%). These results suggest that increasing model capacity contributes not only to improved classification performance but also to more robust generalization.

Table 2. Mean performance ($\pm$ standard deviation) of ConvNeXt variants across folds.

Variant	ACC	SEN	SPE	PRE	F1-Score	AUC
Small	86.24% $\pm$ 9.61	83.33% $\pm$ 19.69	88.75% $\pm$ 17.82	91.99% $\pm$ 11.09	85.10% $\pm$ 11.10	0.981 $\pm$ 0.023
Base	83.00% $\pm$ 13.03	80.83% $\pm$ 18.15	84.64% $\pm$ 29.45	89.80% $\pm$ 15.21	82.43% $\pm$ 12.68	0.986 $\pm$ 0.021
Large	90.43% $\pm$ 6.22	90.30% $\pm$ 11.64	90.18% $\pm$ 11.10	91.79% $\pm$ 8.55	90.25% $\pm$ 6.81	0.983 $\pm$ 0.023

The analysis of model scaling revealed a non-linear performance trend across the ConvNeXt variants. The ConvNeXt-Base model presented lower accuracy (83.00%) and higher variability across folds, as indicated by the increased standard deviation in accuracy (13.03%) and specificity (29.45%), when compared to the ConvNeXt-Small variant (86.24% accuracy). This behavior suggests that the intermediate-capacity model may be more sensitive to data variability, particularly under limited sample conditions. Such instability can be observed in Figure 3, which illustrates the accuracy evolution across folds, highlighting a pronounced performance drop for the Base model in Fold 4. In contrast, the ConvNeXt-Large variant exhibits more stable performance throughout the validation process.

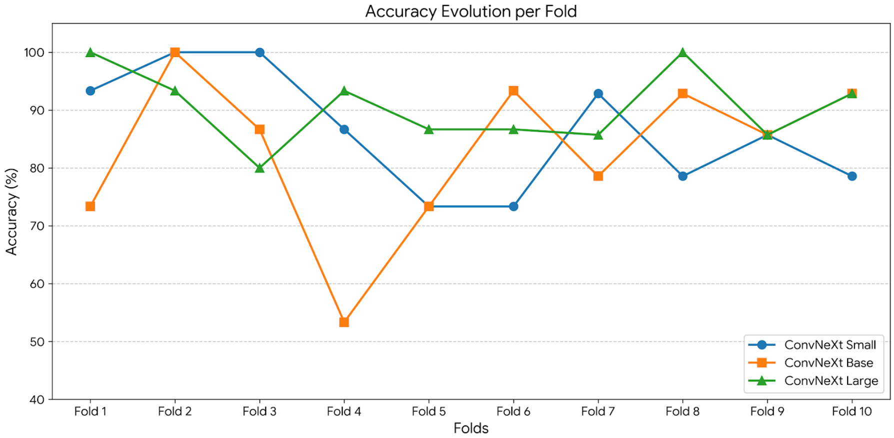

Figure 3. Accuracy evolution per fold across the ConvNeXt variants.

In contrast, the ConvNeXt-Small variant benefits from its lower model complexity, which may reduce the risk of overfitting under limited data conditions. On the other hand, the ConvNeXt-Large model appears to better leverage its higher representational capacity, capturing thermal patterns associated with malignancy with greater consistency across folds.

The effectiveness of transfer learning is reflected in the rapid convergence observed during training. Across all folds and model variants, the early stopping mechanism was consistently triggered within the initial epochs, typically between the 1st and 5th epochs. This behavior suggests that the pretrained weights provided a strong initialization, requiring only limited fine-tuning to adapt to the thermographic domain.

A representative example of this convergence behavior is shown in Figure 4, which presents the loss curves for one of the validation folds of the ConvNeXt-Large model. The validation loss stabilizes early in the training process, indicating that the network quickly adapts to the target task, reducing the computational effort required for training.

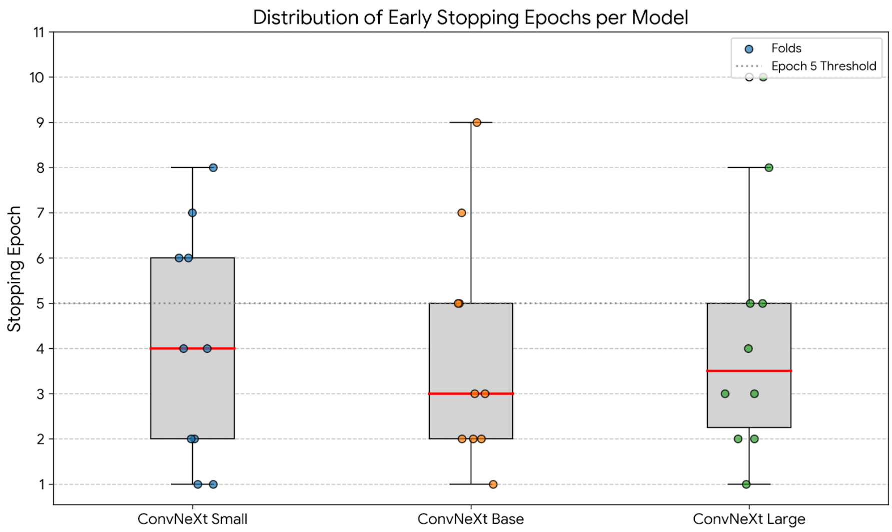

Figure 4. Distribution of early stopping epochs per model.

Despite the differences observed in threshold-dependent metrics, all ConvNeXt variants achieved consistently high values in the ROC-AUC metric, ranging from 0.981 to 0.986. This result indicates that, regardless of model scale, the extracted features are effective for discriminating between healthy and pathological patterns in infrared images at the probabilistic level.

5. Discussion

The results obtained in this study demonstrate that ConvNeXt architectures are capable of effectively modeling thermal patterns for breast cancer classification, achieving consistent performance across different model scales. In particular, the ConvNeXt-Large variant presented the best overall results in threshold-dependent metrics, such as accuracy, F1-score, and sensitivity, as well as improved stability across cross-validation folds. These findings indicate that increased model capacity can enhance generalization, even in scenarios with limited data, provided that appropriate regularization and transfer learning strategies are employed.

Despite the strong performance of the proposed method, a comparison with recent works (Table 3) shows that higher classification metrics have been reported in the literature. Methods based on radiomics, hybrid CNN-classifier pipelines, and ensemble learning strategies achieve accuracy values above 97% in several cases. However, it is important to consider that many of these approaches rely on handcrafted features, multi-stage pipelines, or complex optimization procedures, which may limit their reproducibility and generalization across datasets. In contrast, the proposed method adopts a unified deep learning framework with minimal preprocessing, focusing on evaluating the representational capacity of modern convolutional architectures.

Table 3. Comparison of recent approaches for breast cancer classification using static infrared imaging.

Method	ACC	SEN	SPE	PRE	F1-Score	AUC
[Baffa and Conci 2022]	98.61%	98.61%	–	98.61%	–	99.14%
[Hayat 2023]	96.55%	95.76%	97.34%	95.50%	95.43%	96.21%
[Nguyen Chi et al. 2025]	99.62%	99.63%	–	99.63%	99.63%	–
[Alzahrani et al. 2025]	97.50%	98.40%	–	96.80%	97.60%	–
[Aggarwal et al. 2026]	97.10%	97.52%	96.73%	–	–	99.40%
Proposed Method	90.43%	90.30%	90.18%	91.79%	90.25%	98.29%

An additional experiment evaluated the impact of image segmentation as a preprocessing step. Preliminary results indicated that segmentation did not improve performance and, in some cases, introduced additional variability in the results. Given the increased methodological complexity and the absence of measurable gains, this step was removed from the final pipeline. This behavior can be partially explained by the attention mechanisms embedded in the ConvNeXt architecture, which allow the model to implicitly focus on relevant regions without requiring explicit region-of-interest extraction.

It is also important to note that transformer-based and hybrid architectures typically require larger and more diverse datasets to fully exploit their representational capacity. The relatively limited size of the dataset used in this study may have constrained the performance of the evaluated models, particularly for intermediate-capacity configurations. This limitation reinforces the need for larger annotated thermographic datasets to enable more robust training and fair comparison across architectures.

Finally, future work may explore the evaluation of transformer-based models and hybrid architectures, as well as the incorporation of more advanced data augmentation strategies. Techniques based on generative models, such as GANs, may help increase data diversity and improve model generalization. Additionally, further investigation into calibration strategies and patient-level decision modeling could contribute to enhancing the clinical applicability of thermographic breast cancer classification systems.

6. Conclusion

This study addressed the problem of breast cancer classification using infrared images by evaluating ConvNeXt architectures under a patient-level validation protocol. The results demonstrated that the ConvNeXt-Large variant achieved the best overall performance and stability, while all models exhibited strong discriminative capability as indicated by high ROC-AUC values. These findings suggest that modern convolutional architectures are effective for capturing thermal patterns associated with malignancy, supporting their potential use in computer-aided diagnosis systems based on infrared imaging.

Use of Generative AI

Generative AI tools, such as ChatGPT 5.2, were used exclusively for language revision (grammar and spelling).

Acknowledgment

M. F. O. B. is supported in party by The São Paulo Research Foundation (FAPESP).

References

Aggarwal, A. K., Alphana, and Pandey, M. (2026). Ensemble classification of breast cancer using texture and color features. *SN Computer Science*, 7(3):221.

Alzahrani, R. M., Sikkandar, M. Y., Begum, S. S., Babetat, A. F. S., Alhashim, M., Al-duraywish, A., Prakash, N., and Ng, E. Y. (2025). Early breast cancer detection via infrared thermography using a cnn enhanced with particle swarm optimization. *Scientific Reports*, 15(1):25290.

Baffa, M. d. F. O. and Conci, A. (2022). Radiomics for breast ir-imaging classification. In *MICCAI Workshop on Medical Image Assisted Blomarkers' Discovery*, pages 10–19. Springer.

Cihan, M. and Ceylan, M. (2025). Multi-view thermal breast imaging for malignancy detection: Performance benchmarking of cnn, transformer, and involution architectures. In *International Conference on Artificial Intelligence over Infrared Images for Medical Applications*, pages 20–35. Springer.

D'Alessandro, G., Tavakolian, P., and Sfarra, S. (2024). A review of techniques and bioheat transfer models supporting infrared thermal imaging for diagnosis of malignancy. *Applied Sciences*, 14(4):1603.

Evans, A., Trimboli, R. M., Athanasiou, A., Balleyguier, C., Baltzer, P. A., Bick, U., Camps Herrero, J., Clauser, P., Colin, C., Cornford, E., et al. (2018). Breast ultrasound: recommendations for information to women and referring physicians by the european society of breast imaging. *Insights into imaging*, 9(4):449–461.

Hayat, A. (2023). Breast cancer detection system from thermal images using swin transformer. *Journal Press India*, 3(1).

Instituto Nacional do Câncer (2024). *Controle do câncer de mama no Brasil: dados e números 2024*. INCA, Rio de Janeiro. Available at: https://www.gov.br/inca/pt-br/assuntos/gestor-e-profissional-de-saude/controle-do-cancer-de-mama/dados-e-numero. Accessed in 17 mar. 2026.

Koo, M. M., Von Wagner, C., Abel, G. A., McPhail, S., Rubin, G. P., and Lyratzopoulos, G. (2017). Typical and atypical presenting symptoms of breast cancer and their associations with diagnostic intervals: Evidence from a national audit of cancer diagnosis. *Cancer epidemiology*, 48:140–146.

Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S. (2022). A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986.

Mashekova, A., Zhao, Y., Ng, E. Y. K., Zarikas, V., Fok, S. C., and Mukhmetov, O. (2022). Early detection of the breast cancer using infrared technology – a comprehensive review. *Thermal Science and Engineering Progress*.

National Health Service (2026). Symptoms of breast cancer in women. Available at: https://www.nhs.uk/conditions/breast-cancer-in-women/symptoms-of-breast-cancer-in-women/. Accessed in 17 mar. 2026.

Nguyen Chi, T., Le Thi Thu, H., Doan Quang, T., and Taniar, D. (2025). A lightweight method for breast cancer detection using thermography images with optimized cnn feature and efficient classification. *Journal of Imaging Informatics in Medicine*, 38(3):1434–1451.

Pinto, G., León, J., Quintero, B., Villamizar, D., and Rueda-Chacón, H. (2025). Multimodal vision-language transformer for thermography breast cancer classification. In *2025 IEEE Colombian Conference on Applications of Computational Intelligence (ColCACI)*, pages 1–6. IEEE.

Rodd, B. (2026). Detecting vasodilation for early breast cancer using vit-based distribution embedding via extreme value theory on deep matrix approximation. *Biomedical Signal Processing and Control*, 119:109821.

Silva, L., Saade, D., Sequeiros, G., Silva, A., Paiva, A., Bravo, R., and Conci, A. (2014). A new database for breast research with infrared image. *Journal of Medical Imaging and Health Informatics*, 4(1):92–100.

World Health Organization (2024). Breast cancer. Available at: https://www.who.int/news-room/fact-sheets/detail/breast-cancer. Accessed in 17 mar. 2026.