Journal of Health Informatics

J Health Inform. 2026;18:e1836
DOI: https://doi.org/10.59681/2175-4411.v18.iEspecial.2026.1836

ARTIGO COMPLETO
CBIS 2026

A comparative study of deep learning models for classification of histopathological images of breast cancer

Estudo comparativo de modelos de aprendizado profundo para classificação de imagens histopatológicas de câncer de mama

Arthur Moreira Nogueira de Melo[1], João Marcelo de Souza Motta[1], Maria Karoline da Silva Andrade[1], Juliana Carneiro Gomes Cassemiro[1], Giselle Machado Magalhães Moreno[1], Cecília Cordeiro da Silva[1], Wellington Pinheiro dos Santos[1], Maíra Araújo de Santana[1] ✉

ORCID: 0009-0005-4169-2941 · 0009-0006-9769-656X · 0000-0003-0727-6707 · 0000-0002-0785-0767 · 0000-0003-4076-3494 · 0000-0002-7061-8477 · 0000-0003-2558-6602 · 0000-0002-1796-7862

[1] Universidade Federal de Pernambuco, Recife (PE), Brasil.

ARTICLE INFORMATION

TYPE
Artigo completo · CBIS 2026

IDENTIFIER
e1836

DOI
10.59681/2175-4411.v18.iEspecial.2026.1836

RECEIVED
17 May 2026

ACCEPTED
20 July 2026

REVIEW PROCESS
Double Blind Review

HANDLING EDITOR
Juliano de Souza Gaspar

ACCESS
Open, CC BY 4.0

CORRESPONDING AUTHOR
Maíra Araújo de Santana
maira.araujosantana@ufpe.br

LICENSE

This is an open access article published under the Creative Commons CC BY 4.0 license, which permits use, distribution, and reproduction in any medium, provided the original work is properly cited.

CC BY 4.0

ABSTRACT

Introduction: Accurate diagnosis of breast cancer through histopathological images is fundamental for therapeutic decision-making; however, manual analysis is time-consuming and dependent on the pathologist's experience. **Objective:** To present a comparative study of four Deep Learning architectures: ResNet50, EfficientNetV2, Vision Transformer (ViT), and ConvNeXt, all applied to the multiclass classification of eight lesion subtypes (benign and malignant) from the BreaKHis database. **Method:** An optimized Transfer Learning strategy with automatic hyperparameter selection via Optuna and two-phase training (feature extraction and fine-tuning) was employed. **Results:** ConvNeXt achieved the highest overall accuracy (94.18%), surpassing ViT (92.67%) and ResNet50 (89.13%). EfficientNetV2 also stood out for its balance between performance (91.40%) and computational cost. **Conclusion:** Convolutional networks incorporating hierarchical Transformer principles preserve important inductive bias for identifying histopathological patterns in moderately sized datasets.

Keywords: breast neoplasms; histopathology; artificial intelligence.

RESUMO

Introdução: O diagnóstico preciso do câncer de mama por imagens histopatológicas é fundamental para a decisão terapêutica, porém a análise manual é demorada e depende da experiência do patologista. **Objetivo:** Apresentar um estudo comparativo de quatro arquiteturas de Deep Learning: ResNet50, EfficientNetV2, Vision Transformer (ViT) e ConvNeXt, todas aplicadas à classificação multiclasse de oito subtipos de lesões (benignas e malignas) da base BreaKHis. **Método:** Empregou-se uma estratégia otimizada de Transfer Learning com seleção automática de hiperparâmetros via Optuna e treino em duas fases (extração de características e ajuste fino). **Resultados:** A ConvNeXt obteve a maior acurácia global (94,18%), superando o ViT (92,67%) e a ResNet50 (89,13%). A EfficientNetV2 também se destacou pelo equilíbrio entre desempenho (91,40%) e custo computacional. **Conclusão:** Redes convolucionais que incorporam princípios hierárquicos dos Transformers preservam viés indutivo importante para a identificação de padrões histopatológicos em datasets de tamanho moderado.

Descritores: neoplasias da mama; histopatologia; inteligência artificial.

CBIS 2026 XXI CONGRESSO BRASILEIRO DE INFORMÁTICA EM SAÚDE

Paper presented at the 21st Brazilian Congress of Health Informatics (CBIS 2026), held from September 23 to 25, 2026, in Brasília (DF), Brazil, and published in the CBIS 2026 section of the Journal of Health Informatics.

HOW TO CITE THIS ARTICLE: de Melo AMN, Motta JMS, Andrade MKS, Cassemiro JCG, Moreno GMM, da Silva CC, et al. A comparative study of deep learning models for classification of histopathological images of breast cancer. J Health Inform 2026;18:e1836. doi: https://doi.org/10.59681/2175-4411.v18.iEspecial.2026.1836

1 INTRODUCTION

Breast cancer represents a global public health challenge, with an estimated incidence of 11.7% of new annual cancer cases worldwide [1]. Although rare before the age of 35, its incidence increases progressively from that age group, especially after the age of 50 [1]. In Brazil, an estimated 73,610 new annual cases were expected between 2023 and 2025, corresponding to a rate of 66.54 per 100,000 women. The main risk factors include germline mutations in the BRCA1 and BRCA2 genes, reproductive characteristics associated with prolonged estrogen exposure (late first pregnancy, nulliparity), family or personal history of the disease, as well as obesity, alcohol consumption, smoking, and exposure to ionizing radiation.

Current diagnosis is based on multimodal approaches, combining mammography, ultrasound, and immunohistochemical tests for molecular classification, which are fundamental for prognosis and therapeutic definition. However, despite diagnostic and therapeutic advances, the psychological and social impact of breast cancer and its treatments remains substantial, demanding not only greater technical accuracy but also humanization in care.

The continuous increase in the incidence of the disease and the challenges around this issue reinforce the need for faster, more accurate, and scalable screening and diagnostic methods. In this context, the digitization of pathology and advances in artificial intelligence (AI) have revolutionized the analysis of breast images. Deep learning architectures demonstrate high performance in clinically relevant tasks, such as automated quantification of biomarkers, detection of mitotic figures, identification of lymph node metastases, and prediction of therapeutic response and prognosis.

However, the clinical implementation of these technologies still faces technical, regulatory, and ethical challenges, such as the need for large databases, interoperability between systems, high computational requirements, certification for clinical use, and privacy and accountability issues in automated decisions. Overcoming these obstacles requires rigorous validation and multidisciplinary collaboration among computer scientists, pathologists, health managers, and public policy makers, aligning with the principles of ethical and sustainable innovation in Digital Health.

Previous studies have compared deep learning architectures for histopathological breast cancer classification, primarily focusing on conventional convolutional neural networks (CNNs) under standard transfer learning settings [2,3]. In contrast, the present study evaluates representative architectures from different design paradigms using a unified two-phase training protocol and automated hyperparameter optimization. Although pathology-specific foundation models such as CONCH and Virchow2 currently represent the state of the art, they were not included because they address a different experimental setting (whole-slide image classification) and require substantially greater computational resources [4]. Instead, this work provides a practical, reproducible benchmark of widely accessible general-purpose architectures for supervised patch-level classification.

This paper is structured as follows: Section 2 describes the BreaKHis database and the experimental methodology, including the deep learning architectures compared (ResNet50, EfficientNetV2, Vision Transformer, and ConvNeXt), the hyperparameter optimization strategy with Optuna, and the two-phase training protocol. Section 3 presents and analyzes the results obtained, emphasizing accuracy, sensitivity, specificity, and computational cost metrics. Finally, Section 4 presents the conclusions of the work.

2 METHODS

The methodology adopted in this study is illustrated in Figure 1 and it was structured in the following main stages: (A) preparation and organization of the database, (B) exploratory data analysis (EDA), (C) preprocessing and definition of data augmentation procedures, (D) construction and (E) training of the deep learning architecture, (F) optimization and (G) evaluation of the model. The entire pipeline was implemented in a Google Colab environment using the PyTorch library. These steps are further explained in this section.

Figure 1. General methodology steps of the study

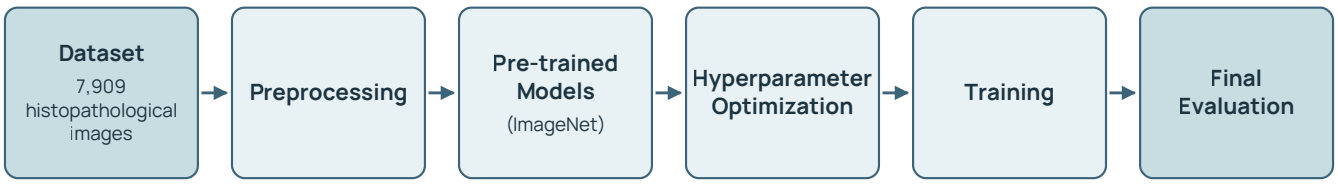

Source: Prepared by the authors (2026).

Figure 2. Random sample of each lesion subtype of the image database

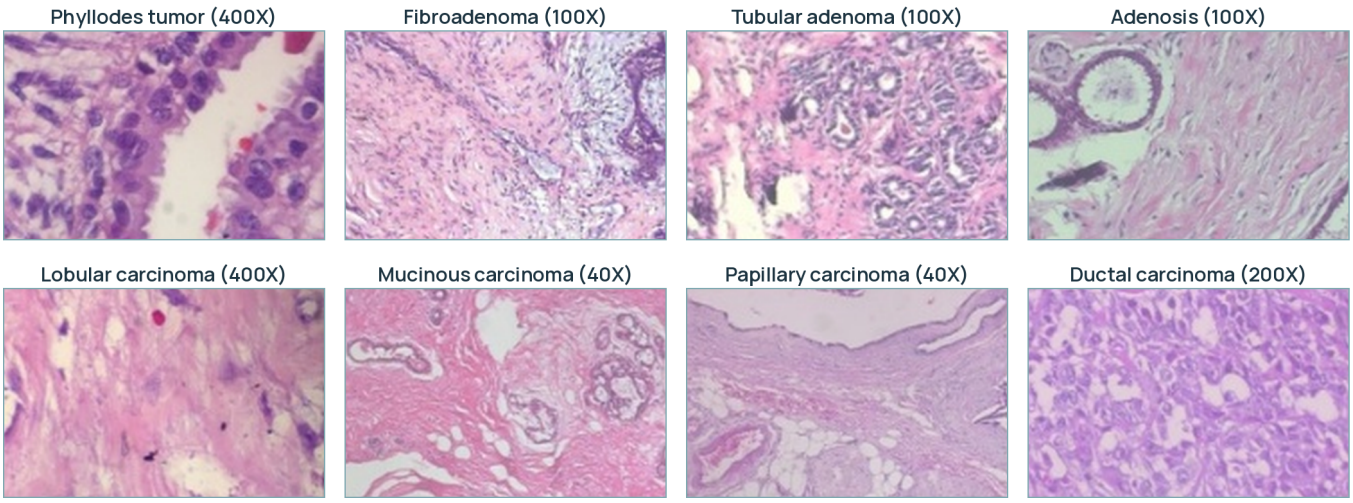

Source: Prepared by the authors (2026), adapted from [5].

2.1 A. Database

As the main data source we used the BreaKHis database [5] built from Brazilian data collected by a group from the Federal University of Paraná. This dataset contains 7,909 microscopic images of breast lesion tissues collected from 82 patients. The images were captured at different magnifications (40x, 100x, 200x, and 400x). The samples are divided into two categories: benign (2,480 samples), which are non-cancerous lesions and generally do not represent a life-threatening risk, and malignant (5,429 samples), characterized by the uncontrolled growth of cells with the potential to invade other tissues and organs, and can be fatal. The images have a resolution of 700x460 pixels, are in 3-channel RGB, with a depth of 8 bits in each channel, and are in PNG format. Furthermore, they are classified into four types of benign lesions: adenosis, fibroadenoma, phyllodes tumor, and tubular adenoma; and four types of malignant tumors: ductal carcinoma, lobular carcinoma, mucinous carcinoma, and papillary carcinoma.

To conduct the experiments, we divided this dataset into three subsets: 80% for training, 10% for validation, and 10% for testing, preserving the original distribution of the subtypes, which is further shown in the following section.

2.2 B. Exploratory Data Analysis (EDA)

Exploratory Data Analysis was performed on the complete set of images to obtain a deep understanding of the distribution, structure, and visual characteristics of the BreaKHis dataset. Before analyzing the dataset, a random sample visualization of each subtype was performed to assess the complexity of the images (Figure 2).

The class distribution revealed a notable imbalance. At the binary level (Figure 3), the Malignant class (5,429 samples) predominates over the Benign category (2,480 samples). This imbalance becomes more substantial when assessing the data distribution among the eight lesion subtypes (Figure 4), where the Ductal Carcinoma class has more than 3 times the data amount of the other categories, displaying the need to handle this issue in order to mitigate the risk of model bias.

We addressed this class imbalance by using a Weighted Random Sampler in the training dataloader, ensuring that samples from minority classes received greater weight, based on an inverse relationship between the number of images in the class and the weight assigned to it. Figure 5 displays this relationship.

Regarding microscopic magnification, the dataset has images in four scales (40x, 100x, 200x, and 400x) and the distribution of images by magnification is completely balanced (Figure 6a). It means that each subtype has the same number of samples at each magnification level. This uniformity highlights the complexity of the dataset, requiring the architecture to be able to learn relevant features at different scales of detail, without relying on a prevalence of magnifications per class. Finally, considering the color importance to differentiate lesion subtypes in pathology assessment, we conducted a statistical distribution analysis of the pixels involved in the images by displaying the color

Figure 3. Distribution of data between Benign and Malignant

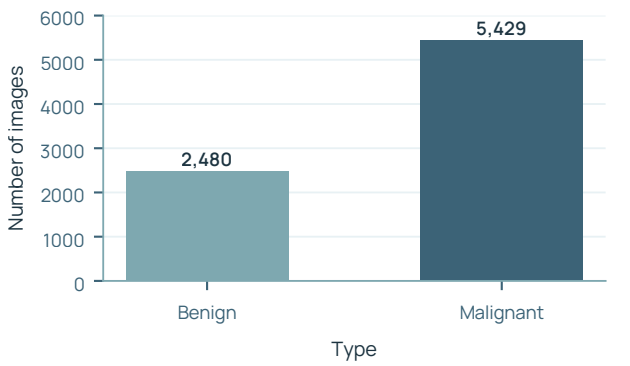

Source: Prepared by the authors (2026).

Figure 4. Data distribution among the eight lesion subtypes

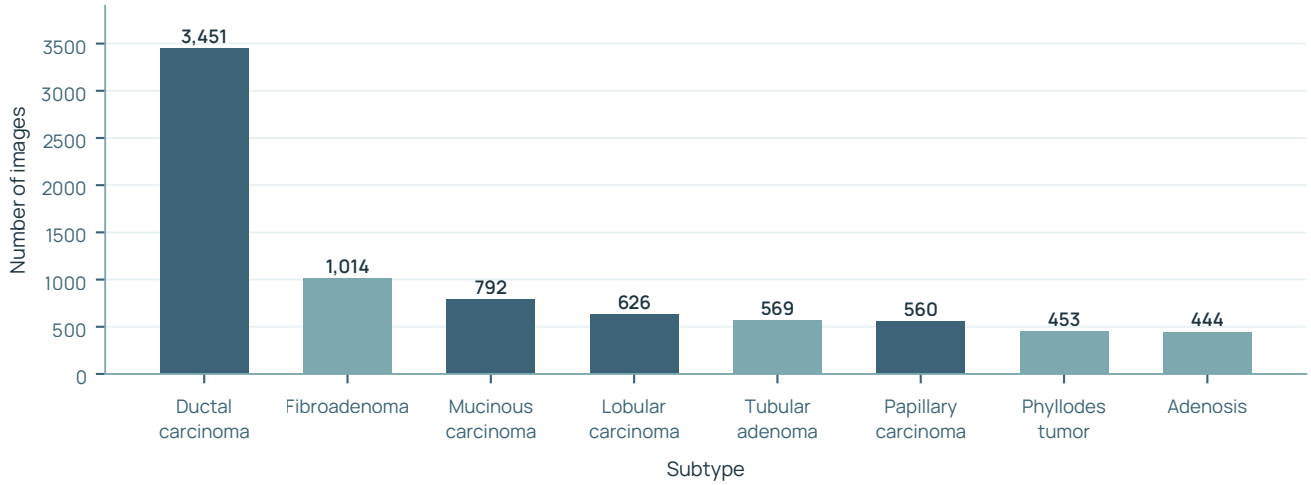

Source: Prepared by the authors (2026).

Figure 5. Distribution of samples by class and their respective weighting in the weight function

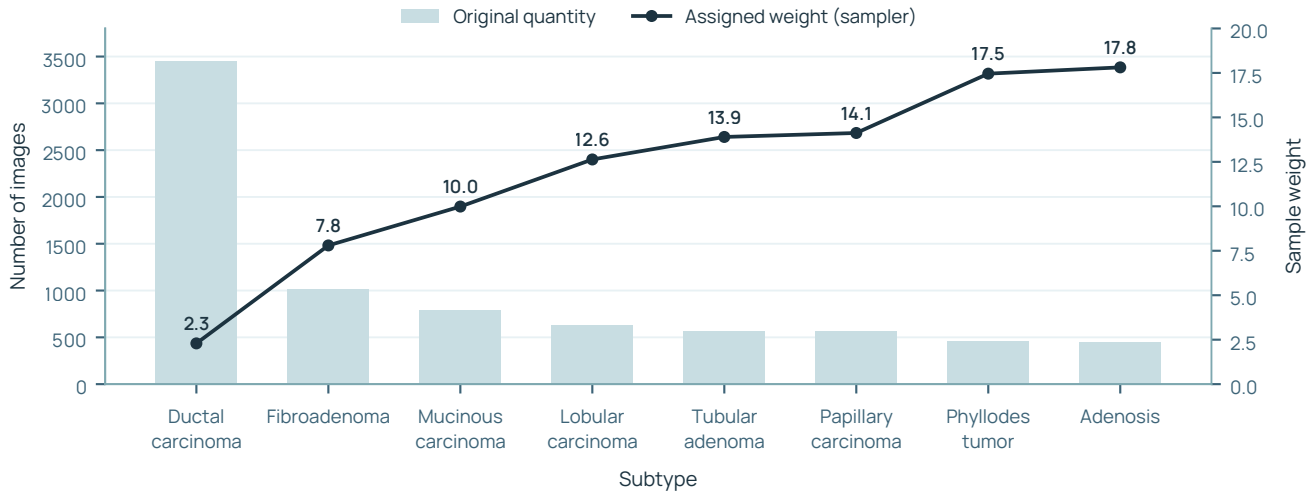

Source: Prepared by the authors (2026).

Figure 6. (a) Distribution of subtypes by magnification factor; (b) pixel intensity distribution by RGB channel

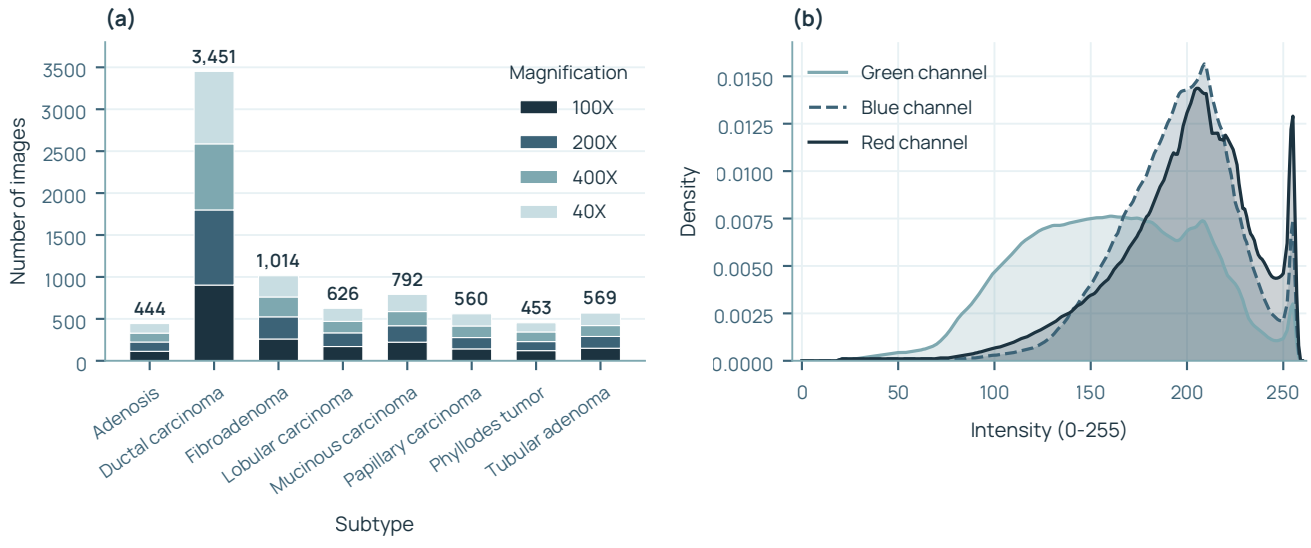

Source: Prepared by the authors (2026).

intensity distribution in each RGB channel (Figure 6b). This visualization allowed us to see aspects related to the histopathological nature of the images. The high peak at the right side of the graph on Figure 6b confirms that a large part of the images have the white background of the microscope, where the pixel intensity approaches the maximum value (255, 255, 255 in RGB). Furthermore, the predominance of the Red and Blue channels at high intensities reflects the H&E (Hematoxylin and Eosin) chemical staining, which generates the pink and purple tones used to visualize tissues and nuclei. In contrast, the Green channel is more spread out and without a large peak, due to its lower presence in the purple and pink pigments that dominate the staining of the samples. This EDA phase, therefore, provided a clear understanding of the quantitative and visual characteristics of BreaKHis database, serving as a basis for defining the procedures in the next steps.

2.3 C. Preprocessing and Data Augmentation

The data preparation workflow was structured to adapt the histopathological samples to the requirements of deep neural networks and optimize feature learning. Initially, we resized all images to the spatial dimension of 224x224 pixels, ensuring the necessary dimensional compatibility with the architectures, which were pre-trained in ImageNet. Next, the image tensors were normalized using the mean and standard deviation of the ImageNet dataset. This procedure is fundamental to align the distribution of the new data with the numerical scale of the pre-trained weights, facilitating numerical stability and accelerating gradient convergence during training.

Finally, specifically during the training phase, we applied a data augmentation strategy, consisting of random rotations, horizontal flips, and brightness and contrast perturbations. The inclusion of these stochastic transformations aims to mitigate overfitting, which is a common challenge in medical datasets, since these approaches may artificially increase the variability of the data, forcing the model to learn structural patterns independent of orientation or lighting conditions.

2.4 D. Model Architectures

For the comparative analysis of the AI models, four distinct architectures representing widely adopted, high-performance design paradigms in computer vision were selected. All architectures were loaded using the timm library [6] (PyTorch Image Models), ensuring a unified and consistent programming interface (API) for the configuration and fine-tuning of different cutting-edge models. We initialized all models with pre-trained weights on the ImageNet-1k dataset (Transfer Learning, 1,000-class dataset), which allows leveraging generic knowledge of visual features to overcome the limitation of labeled data in BreaKHis. To standardize the experiment, the final classification layer of all models was automatically replaced and adapted by the timm API to output 8 classes from the dataset.

ResNet50 (resnet50.a1_in1k) [7] was adopted as the baseline due to its proven robustness and the introduction of residual connections, which mitigate the problem of gradient vanishing in deep networks, allowing for effective learning of complex features. In contrast, EfficientNetV2 (tf_efficientnetv2_m.in1k) [8] was chosen to evaluate the trade-off between performance and computational cost; this network uses a neural architecture search and separable convolutions (“Fused-MBConv”) to maximize training speed and parameter efficiency without sacrificing accuracy. Seeking to explore beyond traditional Convolutional Networks, we included Vision Transformer (vit_base_patch16_224.in1k) [9]. The choice of this model is justified by its ability to capture long-range global dependencies through the Self-Attention mechanism. By processing the image as a sequence of patches, ViT can be crucial for identifying dispersed histological patterns that escape the local receptive field of CNNs. Finally, ConvNeXt (convnext_base.in1k) [10] was used, a modernized purely convolutional architecture with design choices inspired by Transformers (such as large kernels and specific normalization layers). Its inclusion aims to test whether an optimized CNN can compete with the generalization capacity of Transformers in medical images. To standardize the experiment, the original classification layer of all models was replaced by an identical dense structure.

2.5 E. Two-Phase Training

The training protocol was executed in two sequential stages. In the initial stage, called feature extraction, we froze the parameters of the feature extractor (backbone), restricting gradient calculations and weight updates exclusively to the classification layer (head). After this we performed the fine-tuning phase, where the architecture was fully thawed. The complete model was then subjected to a new training cycle, using a reduced learning rate compared to the previous phase, allowing for a global readjustment of the parameters to the BreaKHis data.

2.6 F. Optimization with Optuna

Most of the workload in setting deep network models lies in the definition of hyperparameters. Here we automated this process using the Optuna framework [11], employing the Tree-structured Parzen Estimator (TPE) to explore the search space efficiently. The process consisted of dynamically varying the learning rate on a logarithmic scale, the dropout rate (range 0.3 to 0.5), and choosing the optimizer (AdamW, SGD with momentum, or RMSprop), with the main goal of minimizing CrossEntropyLoss in the validation subset.

Table 1. Pre-Optimization Test Set Results

Architecture	Accuracy	F1-Score	Precision	Recall
EfficientNetV2	58.79%	0.59	0.65	0.58
ResNet50	57.65%	0.58	0.68	0.57
ConvNeXt	75.98%	0.76	0.77	0.75
ViT	67.38%	0.68	0.69	0.67

Source: Prepared by the authors (2026).

Table 2. Results of the Post-Optimization Test Set

Architecture	Accuracy	F1-Score	Precision	Recall
EfficientNetV2	91.40%	0.91	0.92	0.91
ResNet50	89.13%	0.89	0.91	0.89
ConvNeXt	94.18%	0.94	0.94	0.94
ViT	92.67%	0.92	0.92	0.92

Source: Prepared by the authors (2026).

At the end of the iterations, the parameter combination that resulted in the lowest loss was fixed, establishing the configuration used in the final training of the models.

2.7 G. Final Assessment

The final performance analysis was conducted on an isolated test set, preserved from any prior training or validation steps. To quantify the effectiveness of the models, the metrics of Overall Accuracy, Precision, Recall, and F1-Score were used. Given the inherent imbalance in the BreaKHis dataset, the macro-average was adopted for calculating the metrics per class. Furthermore, a Confusion Matrix was also generated to allow visualization of classification errors, discriminating the incidence of false positives and negatives among the different architectures evaluated.

3 RESULTS AND DISCUSSION

This section presents an evaluation of the performance of the proposed architectures, structured in three phases: analysis of behavior during optimization, final results after refinement (fine-tuning), and a critical discussion of the architectural and clinical implications of the obtained performances.

3.1 A. Performance of Pre-Optimization Models

In this preliminary stage, the effectiveness of pre-trained weights in ImageNet [12] acting as fixed feature extractors was evaluated. Reporting this stage is relevant because it isolates the contribution of hyperparameter optimization and fine-tuning from that of the pre-trained backbone itself, quantifying the domain gap between natural images and histopathological images before any dataset-specific adaptation is applied, and thereby providing the baseline against which the gains discussed in Section 3B are interpreted. A heterogeneity in the convergence of the loss function was observed between the architectures: the modern models, ConvNeXt and ViT, demonstrated superiority by reaching training accuracy levels close to 70%. In contrast, the ResNet50 and EfficientNet architectures showed inferior performance, ranging between 57% and 59%, as detailed in Table 1.

The analysis of the learning curves (Figure 7) reveals substantial oscillations in the training and validation metrics, with a tendency to stabilize only near the 10th epoch. However, a substantial distance is noted between the training and validation curves in all models. This behavior suggests a difficulty in generalization (overfitting), indicating that the amount of data available may be insufficient given the complexity of the networks without proper fine-tuning. In summary, although the generic characteristics of ImageNet are partially transferable, the average performance at this stage justifies the need for hyperparameter optimization and fine-tuning to adapt the models to histopathological specificity.

3.2 B. Performance of Post-Optimization Models

After defining the best hyperparameters via Optuna and executing the second training phase (fine-tuning with complete layer thawing), we found a huge improvement in performance metrics compared to the previous stage. Table 2 summarizes the final results obtained in the set of independent tests.

The impact of fine-tuning was substantial. ConvNeXt network achieved the best overall performance with 94.18% accuracy, which achieved accuracy of only 75.98% (severe underfitting) on the previous phase. This increase of around 18 percentage points indicates that the deep spatial filters of this network needed specific adaptation to capture the morphological nuances of the cells of different lesion subtypes, which differ substantially from the natural images of ImageNet (dataset from where the model was pre-trained).

Figure 7. Evolution of the loss function and accuracy in the training (blue line) and validation (orange line) sets for the (a) ViT, (b) ConvNeXt, (c) ResNet50, and (d) EfficientNetV2 architectures

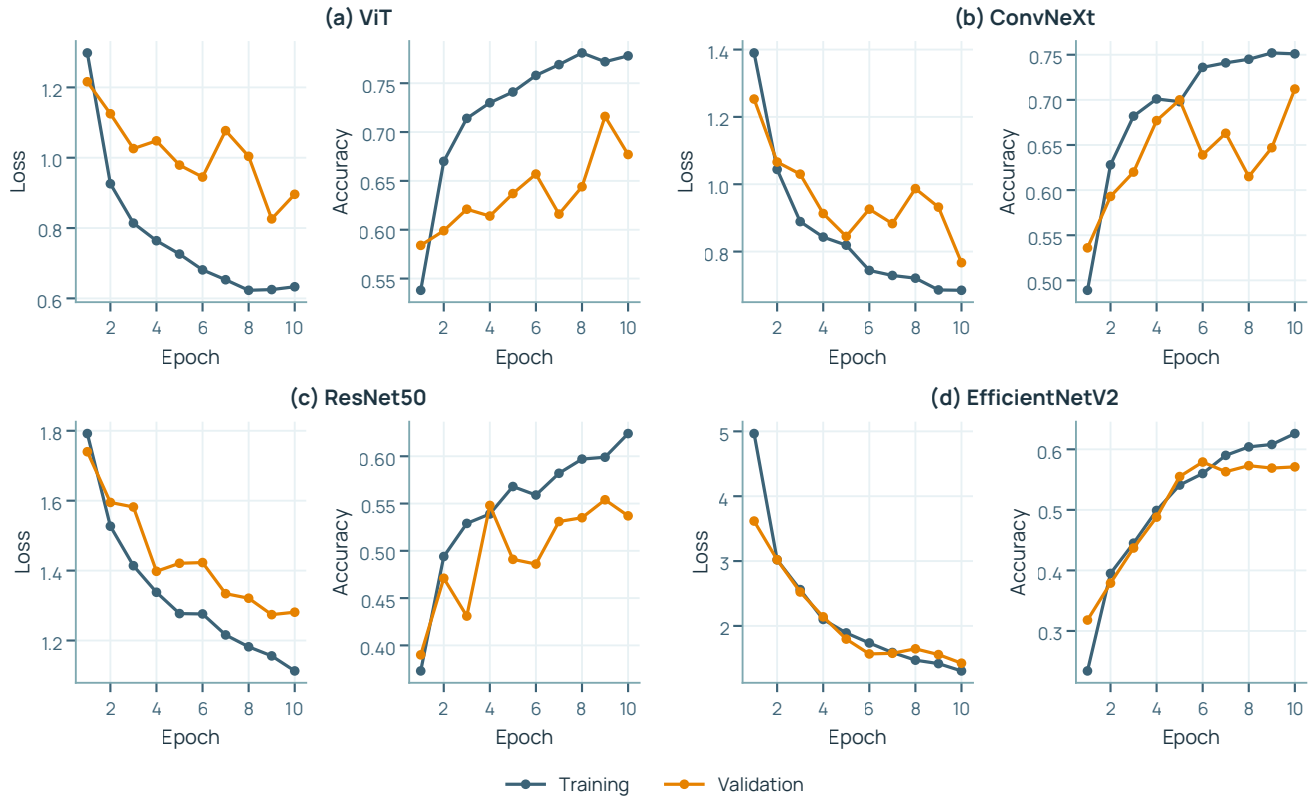

Source: Prepared by the authors (2026).

Vision Transformer (ViT) and EfficientNetV2 consolidated themselves as robust competitors, achieving close metrics, in the range of 91% to 93%. The success of ViT in this phase proves the effectiveness of the regularization found by Optuna (high Dropout rates), allowing the global attention mechanism to generalize correctly despite the limitation in data volume.

Finally, ResNet50 maintained its baseline position, with 89.13%, although it showed a substantial improvement compared to its initial 57.65%. Its inferiority compared to more modern architectures suggests that the classic residual design, while efficient, is less able to extract discriminative features in histopathology than the separable convolutions of EfficientNet or the attention blocks of ViT.

Figure 8 displays the confusion matrices for all evaluated architectures, allowing a direct comparison between the performances obtained in the pre-optimization and post-optimization phases.

3.3 C. Critical and Comparative Analysis of the Models

The superiority of ConvNeXt (94.18%) in multiclass classification might be justified by its hybrid architecture, which integrates the hierarchical structure of CNNs with design choices from Transformers. By using large kernels (7x7) and normalization layers (Layer Norm) instead of the standard Batch Norm, ConvNeXt expands the network's receptive field, allowing it to capture the global context of histology (such as cell arrangement) without losing local inductive bias. This feature was decisive in distinguishing the eight histological subtypes, where visual differences are subtle and require both texture (local) and structure (global) analysis.

In contrast, the slightly lower performance of ViT (92.67%) illustrates the limitation of the pure Self-Attention mechanism in limited data regimes (7,909 images). ViT lacks locality inductive bias (the innate ability to understand that neighboring pixels are correlated). To learn these spatial relationships from scratch and distinguish complex boundaries between 8 classes, ViT would require a substantially larger volume of data. ConvNeXt outperformed ViT precisely because it eases learning through convolutions, while maintaining long-range modeling capabilities.

Finally, the comparison between EfficientNetV2 (91.40%) and ResNet50 (89.13%) reflects the evolution in feature extraction efficiency. ResNet50, while robust, uses standard residual blocks that are less effective at capturing fine-grained features than EfficientNetV2's Fused-MBConv blocks. EfficientNet optimizes the relationship between depth, width, and resolution, allowing for better convergence with fewer parameters. However, even EfficientNet did not reach the level of ConvNeXt, since its smaller kernels limit the understanding of the macroscopic tissue context, crucial for the correct differentiation between

Figure 8. Confusion matrices

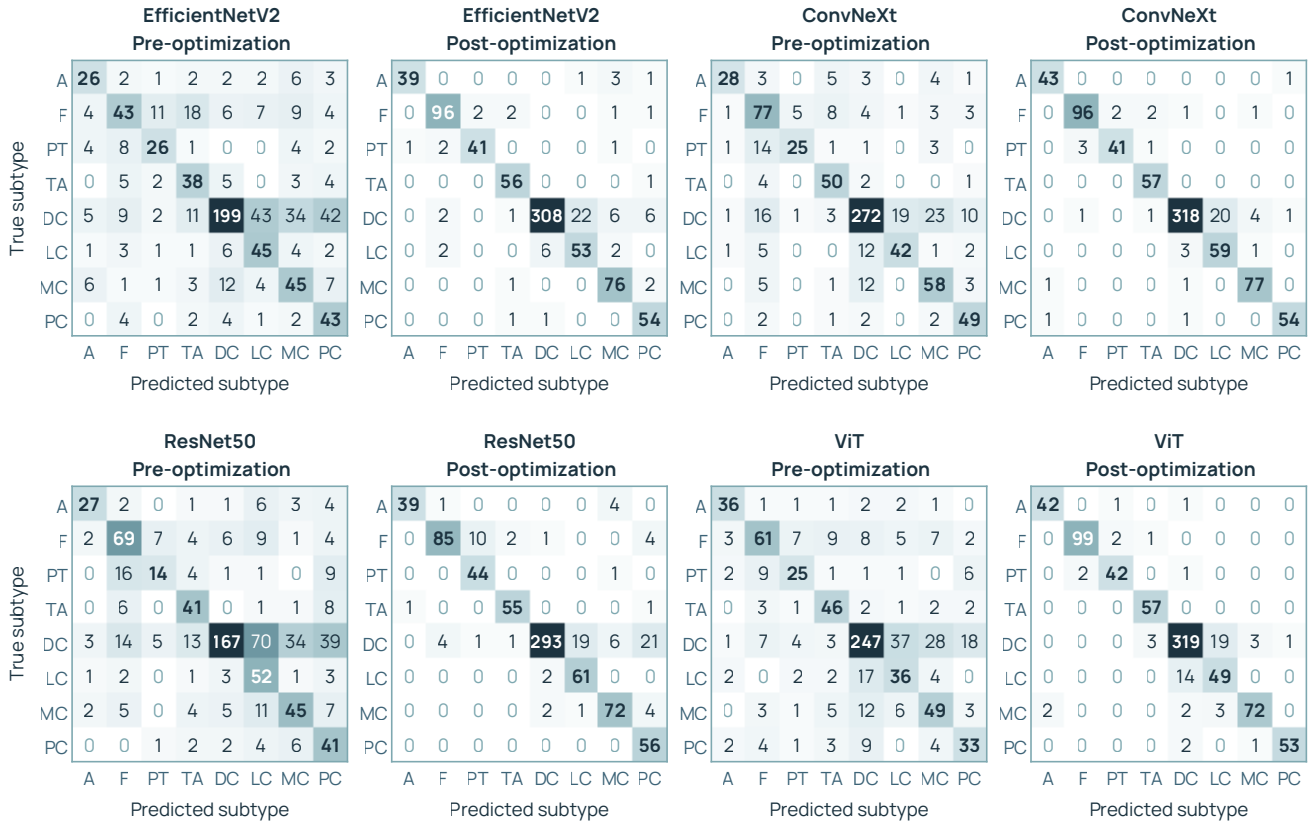

A: adenosis; F: fibroadenoma; PT: phyllodes tumor; TA: tubular adenoma; DC: ductal carcinoma; LC: lobular carcinoma; MC: mucinous carcinoma; PC: papillary carcinoma.
Source: Prepared by the authors (2026).

morphologically similar malignant subtypes.

4 CONCLUSION

This study conducted a comparative analysis between four representative Deep Learning architectures for multiclass classification of histopathological images of breast cancer, using the BreaKHis database. When faced with the challenge of distinguishing between eight histological subtypes with high intra-class variability, the experiments demonstrated that the choice of architecture and the hyperparameter optimization strategy are crucial for the effectiveness of computer-assisted classification.

ConvNeXt architecture overcame the others, achieving an overall accuracy of 94.18%. This result confirms the initial hypothesis that modernizing Convolutional Networks, incorporating Transformer design principles such as expanded kernels and layer normalization, offers the best compromise between extracting local features and understanding the global context. The slight superiority of ConvNeXt over Vision Transformer (ViT) (92.67%) suggests that, in limited data regimes, the inductive bias of convolutions remains crucial for training stability and for the discrimination of subtle textures in images, where pure attention mechanisms still face generalization difficulties. In addition, the study reaffirms the relevance of EfficientNetV2 as a viable alternative for environments with restricted computational resources, balancing efficiency and performance (91.40%). On the other hand, ResNet50, although useful as a baseline, showed signs of obsolescence for tasks of high histological complexity, being surpassed by more recent approaches in this scenario.

Clinically, this study shows a promising approach to help pathology specialists to overcome subjective assessments on breast lesions differentiation and staging. These findings may be further improved by adding explainability (XAI) techniques to map the regions of interest activated by ConvNeXt, increasing clinical confidence in the "black-box" model. Furthermore, the application of Ensemble Learning strategies combining ConvNeXt and ViT could be explored to mitigate the individual weaknesses of each architecture. Future work could also benchmark these general-purpose architectures against pathology-specific foundation models, such as CONCH and Virchow2 [4], to quantify the performance gap between generalist and domain-specialized feature extractors in this classification task. Studies on energy efficiency, computational complexity, and viability of clinical application should also be conducted.

Author contributions. The authors contributed equally to all stages of the article: conception and design of the study; data collection, analysis, and interpretation; drafting and critical revision of the manuscript; and approval of the final version.

Acknowledgments. The authors acknowledge all researchers of the Mamo.IA team from the Laboratório de Computação Biomédica (UPE + UFPE).

Funding. This study was supported by the Brazilian National Council for Scientific and Technological Development (CNPq) (Grant No. 383046/2024-6; Proposal No. 9646049472643444), as part of the project "Mamo.IA: Development of an Artificial Intelligence Platform for Early Diagnosis of Breast Cancer with Emphasis on Mammography, Immunohistochemistry, and Thermography" (Process No. 445896/2023-0), funded under Call No. 21/2023, Track B (Primary and Original Transdisciplinary Studies).

Data and materials availability. The authors did not provide a data and materials availability statement.

Conflicts of interest. The authors did not declare any conflicts of interest.

Use of artificial intelligence. We used ChatGPT (OpenAI) for redesigning Figure 1 and summarizing the conclusion section. No artificial intelligence was involved in the conceptualization of the study, methodological design, data collection or analysis, interpretation of results, or generation of intellectual content. All AI-assisted processes and outputs were rigorously validated and audited by the human authors, who assume full responsibility for the accuracy, originality, and scientific integrity of this work.

REFERENCES

1. Santos MO, Lima FCS, Martins LFL, Oliveira JFP, Almeida LM, Cancela MC. Estimativa de Incidência de Câncer no Brasil, 2023-2025. Rev. Bras. Cancerol. 2022;69(1). doi: https://doi.org/10.32635/2176-9745.rbc.2023v69n1.3700
2. Alom MZ, Yakopcic C, Nasrin MS, Taha TM, Asari VK. Breast cancer classification from histopathological images with inception recurrent residual convolutional neural network. J Digit Imaging. 2019;32(4):605-17. doi: https://doi.org/10.1007/s10278-019-00182-7
3. Talo M. Automated classification of histopathology images using transfer learning. Artif Intell Med. 2019;101:101743. doi: https://doi.org/10.1016/j.artmed.2019.101743
4. Neidlinger P, El Nahhas OSM, Muti HS, Lenz T, Hoffmeister M, Brenner H, et al. Benchmarking foundation models as feature extractors for weakly supervised computational pathology. Nat Biomed Eng. 2026;10(6):1113-23. doi: https://doi.org/10.1038/s41551-025-01516-3
5. Spanhol FA, Oliveira LS, Petitjean C, Heutte L. A Dataset for Breast Cancer Histopathological Image Classification. IEEE Trans Biomed Eng. 2016;63(7):1455-62. doi: https://doi.org/10.1109/TBME.2015.2496264
6. R. Wightman, "PyTorch Image Models," GitHub repository, 2019. [Online]. Available: https://github.com/rwightman/pytorch-image-models.doi:10.5281/zenodo.4414861.
7. He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). [place unknown]: IEEE; 2016. p. 770-8. doi: https://doi.org/10.1109/cvpr.2016.90
8. M. Tan and Q. V. Le, "EfficientNetV2: Smaller models and faster training," in International Conference on Machine Learning (ICML), 2021, pp. 10096-10106.
9. A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in International Conference on Learning Representations (ICLR), 2021.
10. Liu Z, Mao H, Wu CY, Feichtenhofer C, Darrell T, Xie S. A ConvNet for the 2020s. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). [place unknown]: IEEE; 2022. p. 11966-76. doi: https://doi.org/10.1109/cvpr52688.2022.01167
11. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. New York, NY, USA: ACM; 2019. p. 2623-31. doi: https://doi.org/10.1145/3292500.3330701
12. Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. [place unknown]: IEEE; 2009. p. 248-55. doi: https://doi.org/10.1109/cvpr.2009.5206848

EDITORIAL NOTES

The Journal of Health Informatics is a scientific journal maintained by the Brazilian Society of Health Informatics (Sociedade Brasileira de Informática em Saúde), a non-profit scientific association dedicated to the development of Health Informatics and Digital Health in Brazil. The editorial process is conducted with scientific independence, free from institutional, commercial, or sponsor interference, and follows international best practices in publication ethics.

This article was double-blind peer reviewed by specialists in the field. The editorial team used artificial intelligence tools, under human supervision, in the initial screening, in checking the text, references and numbers and, when needed, in redrawing figures, charts and tables, which the authors approved in the proof. The tools and safeguards are described in the journal's Policy on the use of artificial intelligence. Editorial decisions were made by people.

SBIS Sociedade Brasileira de Informática em Saúde