Analyzing the Impact of Image Preprocessing on Tumor Detection and Classification in Mammograms*

Enol García González[1] [0000−0001−7125−9421], Mădălina Dicu[2] [0009−0001−3877−527X], José R. Villar[1] [0000−0001−6024−9527], and Camelia Chira[2] [0000−0002−1949−1298]

[1] Department of Computer Science, University of Oviedo, C. Jesús Arias de Velasco, s/n, Oviedo 33005, Spain
{garciaenol,villarjose}@uniovi.es
[2] Faculty of Mathematics and Computer Science, Babeș-Bolyai University, Str. Mihail Kogălniceanu nr. 1, Cluj-Napoca 400084, Romania
{madalina.dicu,camelia.chira}@ubbcluj.ro

Abstract. Breast cancer is one of the most common diseases affecting women around the world, with millions of new cases each year and a high rate of mortality. Although advances in detection and treatment have led to better outcomes for many, some regions still struggle with limited access to early diagnosis. Early detection of breast cancer significantly improves treatment success, emphasizing the need for accessible, accurate diagnostic tools.

This study focuses on improving the segmentation and classification of breast tumors in mammographic images by examining the impact of various preprocessing techniques on the performance of advanced object detection models. We use the Digital Database for Screening Mammography (DDSM), a benchmark dataset for mammography analysis, to evaluate three state-of-the-art object detection models: Faster R-CNN, YOLOv8, YOLOv9, and YOLOv11. These models are trained on images processed with different methods to observe how each technique affects the accuracy of tumor segmentation and classification into benign or malignant categories.

The primary contribution of this study is the comparative analysis of preprocessing effects on AI model performance in tumor detection, aiming to identify an optimal preprocessing approach that maximizes accuracy and reduces misclassification rates. We systematically assess preprocessing effects on model accuracy, focusing on improvements in tumor visibility and classification performance.

Keywords: breast cancer · image preprocessing · segmentation

* This research has been funded by the Spanish Ministry of Economics and Industry –grant PID2020-112726RB-I00–, the Spanish Research Agency –grant PID2023-146257OB-I00–. Also, by Principado de Asturias, grant IDE/2024/000734, and by the Council of Gijón through the University Institute of Industrial Technology of Asturias grants SV-23-GIJÓN-1-17, and SV-24-GIJÓN-1-05

1 Introduction

Breast cancer remains one of the most prevalent and life-threatening diseases affecting women worldwide, with millions of new cases diagnosed each year and a significant mortality rate. Despite continuous advancements in detection and treatment, disparities in healthcare infrastructure mean that early diagnosis is still a considerable challenge in many regions. Consequently, developing accurate and accessible diagnostic tools is essential for enhancing breast cancer detection rates and reducing mortality, particularly in resource-limited settings [1, 24].

Mammography is the most widely used imaging modality for breast cancer screening [3, 5], providing a non-invasive means of detecting abnormalities at an early stage. However, the effectiveness of mammographic analysis is often hindered by challenges such as low contrast, noise, and variations in tissue density, which can obscure tumor visibility. To address these issues, advanced computer vision techniques, particularly deep learning-based object detection models, have been increasingly employed for automated tumor segmentation and classification. These models have demonstrated remarkable potential in assisting radiologists by improving detection accuracy and reducing interpretation time. Nonetheless, the performance of such models is highly dependent on the quality of input images, making preprocessing techniques a critical factor in optimizing detection outcomes.

This study investigates the impact of various preprocessing techniques on the performance of state-of-the-art object detection models in breast tumor segmentation and classification. Using the Digital Database for Screening Mammography (DDSM) [8], a benchmark dataset for mammographic analysis, we evaluate three leading object detection frameworks: Faster R-CNN [19], YOLOv8 [9], and YOLOv9 [22]. Each model is trained on images subjected to different preprocessing methods to assess how these techniques influence segmentation accuracy and the classification of tumors as benign or malignant.

To better understand the role of preprocessing in mammographic image analysis, we define the following key research questions:

- **RQ1:** How does preprocessing affect the performance of YOLO and Faster R-CNN in mammographic image analysis?
- **RQ2:** Which preprocessing techniques yield the highest tumor classification accuracy?

By answering these research questions, this study aims to provide a systematic evaluation of preprocessing effects, helping to identify optimal strategies for improving tumor detection performance.

The organization of this paper is as follows. Firstly, we present a review of the relevant literature in Section 2. Then, Section 3 presents the materials and methods that were used in this study. In Section 4, the experimentation results are presented together with a discussion of the most relevant achievements of those results. Finally, the conclusions of this paper are in Section 5.

2 Related work

Automated breast tumor detection in mammographic images has been widely studied, with preprocessing techniques playing a crucial role in improving deep-learning model performance. By reducing noise, enhancing contrast, and standardizing images, preprocessing improves tumor identification accuracy. Researchers have explored various methods to refine mammographic images before feeding them into classification or segmentation models.

Noise reduction and contrast enhancement are key preprocessing techniques. Ramani and Vanitha [18] analyzed different filtering methods, showing that while some preserve tumor boundaries, others risk introducing artifacts that obscure critical details. Avcı and Karakaya [2] demonstrated that machine learning-based enhancement improves breast cancer detection rates but requires careful parameter tuning to avoid distortions. Additionally, Warren et al. [23] studied how preprocessing influences lesion visibility, emphasizing the need for adaptable techniques that enhance diagnostic features without compromising interpretability.

Beyond individual image enhancements, preprocessing also plays a fundamental role in dataset standardization for deep learning models. Beeravolu et al. [4] explored artifact removal and contrast normalization, showing improvements in dataset consistency and CNN generalization. Lu et al. [11] further emphasized that tailoring preprocessing strategies to specific architectures, such as YOLO-based models, significantly impacts classification performance.

While many studies focus on individual preprocessing techniques, this work compares multiple methods to assess their impact on tumor detection in deep learning models. As AI in medical imaging advances, preprocessing remains essential for improving accuracy while maintaining interpretability for radiologists. The challenge is to enhance image quality—through contrast adjustment, noise reduction, and artifact removal—without introducing modifications that reduce clinical trust and usability.

3 Material and methods

This section outlines the dataset, preprocessing techniques, and object detection models used in this study. We describe the Digital Database for Screening Mammography (DDSM) dataset, detail the preprocessing strategies applied to enhance tumor visibility and present the deep learning models employed for tumor detection and classification. Finally, we explain the experimental setup and evaluation metrics used for performance comparison.

3.1 The Digital Database for Screening Mammography DDSM dataset

The Digital Database for Screening Mammography (DDSM) [8] is a public dataset maintained by the University of South Florida. It contains images of mammograms performed on patients with and without tumors at different medical centers. This dataset contains about 700 cases of patients without tumors, about 900 cases with benign tumors, and about 900 cases of patients with cancerous tumors. For each case, the dataset includes two types of images of each breast: a Mediolateral Oblique image (MLO) [16] and a Craniocaudal image (CC) [15].

A drawback of this dataset is that the images are digitized using different brands of scanners with varying resolutions in the scanning process. When selecting the images to be used in this work, we selected the MLO images scanned with a HOWTEK scanner with a resolution of 43.5 microns, and only those that present any kind of tumor –benign or malign. In this way, it is standardized that all the images involved in the comparison have the same characteristics. The dataset used in the experimentation is composed of a total of 964 images, of which 547 present cases with benign tumors and 419 cases with cancers. Figure 1 provides an example of the dataset's mammographic images.

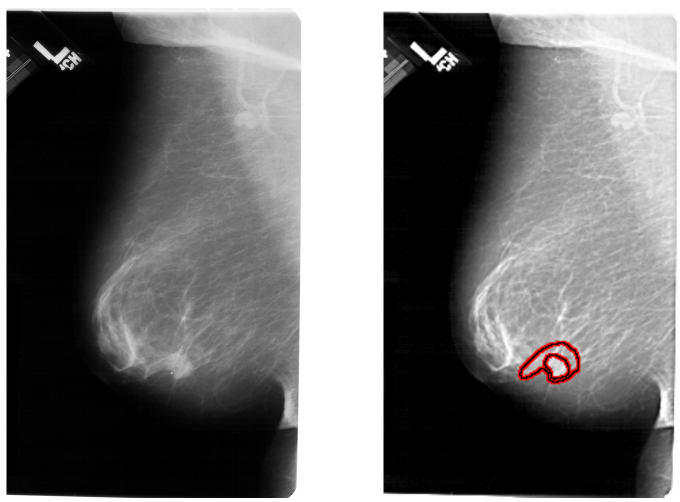

(a) Original image (b) Labeled cancerous regions

Fig. 1: Example of a mammographic image from the DDSM dataset (Case ID: A_1171). The left image (a) shows the original mammogram, while the right image (b) displays the same image with annotated cancerous regions.

3.2 Preprocessing methods

Unsupervised GrowCut preprocessing algorithm The GrowCut algorithm, first introduced by Vezhnevets and Konouchine [21], is designed for multi-label segmentation using a Cellular Automaton framework, where the image is represented as a grid of cells, and individual pixels act as these cells. Each pixel p within an image P is defined by a triplet consisting of its label l_p, a confidence score θ_p, which indicates the certainty that the pixel belongs to class l_p, and a feature vector C_p. The process begins with a set of user-labeled pixels, each assigned a confidence of 1. The algorithm then iteratively updates the labels and confidence values of neighboring cells until no further changes occur, indicating convergence to a stable segmentation. The rule governing cell evolution is expressed in Equation 1, where q represents a neighboring pixel, defined using either the von Neumann or Moore neighborhood system, while the function g is formulated in Equation 2.

$$g(||\overrightarrow{C_p} - \overrightarrow{C_q}||_2) \cdot \theta_q^t > \theta_p^t \tag{1}$$

$$g(x) = 1 - \frac{x}{max||\overrightarrow{C}||_2} \tag{2}$$

An unsupervised adaptation of GrowCut, proposed by Ghosh et al. [7], eliminates the need for user input by initializing with a randomly distributed set of seed points that gradually expand to form segmented regions. This variant, referred to as Unsupervised GrowCut (UGC), addresses two key limitations of the original algorithm: the requirement for manual input and the restriction on the number of segment classes. Initially, random seed points are assigned a confidence value of 1, and their growth follows a modified transition rule, which compares the output of the monotonic function g (from Equation 2) to a predefined threshold. This threshold is also used for merging regions during segmentation. Additionally, UGC establishes equivalence classes that group similar labels, which are dynamically updated when regions merge. If the similarity condition defined in Equation 3 is met, pixels p and q are considered part of the same equivalence class, and their respective regions are fused.

The algorithm was tested on 30 MRI scans of the lungs and brain, using a threshold value of 0.95 and an initial set of 100 randomly placed labels. The segmentation results were visually assessed and compared to those obtained with the MeanShift and NCut methods. Although no ground truth data was available, qualitative analysis suggested that UGC performed comparably to the other two techniques.

$$g(||\overrightarrow{C_p} - \overrightarrow{C_q}||_2) \cdot \theta_q^t > threshold \tag{3}$$

Marginean et al. [13] further advanced the GrowCut method by introducing a more autonomous variation known as Competitive Unsupervised GrowCut (Competitive UGC). This approach retains the Cellular Automaton structure while integrating the label-merging process from UGC [7] and the soft label propagation mechanism from the classical GrowCut algorithm [21].

One key improvement addresses a limitation in UGC, where some segmented regions may expand into adjacent areas without merging correctly. To resolve this, Competitive UGC reinstates the original GrowCut evolution rule (Equation 1) while preserving the region-merging mechanism. However, merging occurs only when a neighboring pixel q actively propagates its label.

The method was tested on 20 MRI scans of the human heart and compared against UGC and AdaPri UGC [14]. Results showed that segmentation performance was largely dependent on the thresholding strategy, with each method demonstrating optimal effectiveness under specific application conditions. Figure 2a shows the result after applying this preprocessing technique.

Erode and dilate The second preprocessing technique used in this work is inspired by Omer et al. [12]. Its main goal is to remove unwanted elements, such as scan marks and artifacts, from mammograms during tumor identification and segmentation. This method relies on erosion and dilation, two fundamental morphological operations in image processing.

Erosion is a transformation that shrinks foreground regions by removing pixels from object boundaries. Mathematically, given an input image I and a structuring element S, erosion, denoted as $I \ominus S$, consists of all points where S is entirely contained within the foreground of I. This operation reduces object size, eliminates small noise, and disconnects thin structures.

In contrast, dilation expands foreground regions by adding pixels to object boundaries. Formally, the dilation of I by S, denoted as $I \oplus S$, consists of all points where the reflection of S intersects with the foreground of I. Dilation enlarges object boundaries, bridges gaps between close structures, and restores eroded features. When applied sequentially, erosion and dilation form morphological operators such as opening and closing, which are widely used in segmentation and noise reduction.

This preprocessing technique follows a three-step process:

1. **Erode and dilate:** In this first phase, the image is processed using erosion and dilate operations. The image undergoes erosion and dilation to smooth surfaces and distinguish breast tissue from unwanted artifacts.
2. **Generate a Boolean mask:** The next step is to build a mask that differentiates the breast areas –the areas that are desired to keep in the image–, from the unwanted areas of the image. A contour detection algorithm [20] identifies segmented areas. The largest region is assumed to represent the breast, while other areas are discarded.
3. **Apply the mask to the original image:** In the final step, the binary mask is applied to the original image, isolating the breast region while preserving fine details.

Figure 2b includes an example of the generated image after applying this preprocessing method.

Combination of two processing methods The experimentation also included an evaluation of the performance when combining the two preprocessing techniques presented in the previous sections. In this scenario, the method based on [12] was applied first, using the erode and dilate operations to generate a mask that segments the original image. The resulting image was then used as a reference for applying the Unsupervised GrowCut technique from [7]. Figure 2c illustrates the outcome of applying both preprocessing techniques together.

(a) Unsupervised GrowCut (b) Erode and dilate based (c) Both methods

Fig. 2: Preprocessing results applied to the original mammogram from Case A_1171 (Fig. 1a). The differences are visible compared to the original version of the images. The GrowCut-based method (a) removes the pectoralis muscle, the erode and dilate-based method (b) removes background artifacts such as labeling, and the combined method (c) removes both.

3.3 Deep Learning methods

This study compares two model families: the Ultralitics' YOLO family [9] and the Faster R-CNN [19].

You Only Look Once YOLO [9] is a family of real-time object detection models that follow a single-stage approach, predicting object classes and bounding boxes in a single pass through the network. This design prioritizes speed while maintaining competitive accuracy. Over time, multiple YOLO versions have been developed, each improving efficiency, accuracy, and feature extraction.

In particular, **YOLOv8** [9] introduces a more efficient backbone and anchor-free detection, improving speed and flexibility for multiple tasks like object detection, segmentation, and pose estimation. On the other hand, **YOLOv9** [22] enhances detection precision using transformer-based modules and improved spatial attention, optimizing computational performance. The most recent release of YOLO is **YOLOv11** [10], which refines efficiency further with better feature extraction and reduced computational cost, making it adaptable for various deployment scenarios.

Faster Region-based Convolutional Neural Network Faster R-CNN [19] is a state-of-the-art object detection algorithm based on Convolutional Neural Networks (CNNs). It enhances previous detection methods by incorporating a Region Proposal Network (RPN), which efficiently generates candidate object regions. These proposals are then refined and classified by a detection network, allowing for precise localization and recognition. As a two-stage detector, Faster R-CNN offers high accuracy and robustness in object detection tasks but requires substantial computational resources.

3.4 Comparative analysis procedure

The comparative analysis conducted in this work involves the three Deep Learning models presented –YOLO, and Faster R-CNN—along with four different versions of the DDSM dataset. The dataset versions used for experimentation are as follows:

1. The original dataset, without any preprocessing applied.
2. A version where preprocessing is applied using only the GrowCut algorithm.
3. A version where preprocessing is applied using only the Erode and Dilate algorithm.
4. A version combining both preprocessing algorithms.

During the experimentation, all three models were trained using each of the four dataset versions, and their Average Precision (AP) [6, 17] was recorded for comparison. To determine the best parameter configurations, an initial study was conducted with various settings, leading to the selection of the values detailed in Table 1. The training process was set to a maximum of 100 epochs, with early stopping applied in cases where no further improvement in model performance was observed.

Table 1: Hyperparameters used for training the models during experimentation.

Parameter	**YOLO**	**Faster R-CNN**
Optimizer	AdamW (momentum: 0.9)	SGD (momentum: 0.9)
Learning Rate	1.667×10^{-3}	2.5×10^{-4}
Batch Size	16	128

4 Results and discussion

This section analyzes the impact of different preprocessing techniques on the performance of object detection models for breast tumor detection. By comparing YOLO-based models and Faster R-CNN, we examine how preprocessing influences detection accuracy and whether it provides significant advantages in medical imaging.

Table 2 presents the results obtained from the experimentation, as described in Section 3. This table compares the Average Precision (AP) at confidence thresholds of 50% and 75% for each model, evaluated using the four different versions of the dataset processed with various preprocessing techniques.

Table 2 presents the results obtained from the experimentation, as described in Section 3. This table compares the Average Precision (AP) at confidence thresholds of 50% and 75% for each model, evaluated using the four different versions of the dataset processed with various preprocessing techniques.

The most obvious finding from these results is the performance gap between the YOLO-based models and Faster R-CNN. Across all experimental scenarios, Faster R-CNN consistently achieves superior accuracy, significantly outperforming the YOLO models.

Table 2: Average Precision (AP) results for each model and preprocessing strategy. AP is reported at confidence thresholds of 50% (AP50) and 75% (AP75). The highest values in each section are highlighted in **bold**.

Preprocessing Method	**AP50 (%)**			
	YOLOv8	**YOLOv9**	**YOLOv11**	**Faster R-CNN**
No Preprocessing (Original)	22.1	14.2	12.7	0.4
Erode and Dilate	12.8	12.0	12.6	**91.1**
GrowCut	18.7	14.7	12.4	77.9
Combined (Erode + GrowCut)	16.8	11.8	9.7	57.0
Preprocessing Method	**AP75 (%)**			
	YOLOv8	**YOLOv9**	**YOLOv11**	**Faster R-CNN**
No Preprocessing (Original)	10.1	6.3	5.2	0.0
Erode and Dilate	6.5	5.7	6.1	27.4
GrowCut	9.6	5.7	6.0	**33.0**
Combined (Erode + GrowCut)	7.8	5.5	4.8	8.3

On the other hand, when analyzing the evolution of the metrics with the different preprocessing algorithms, it can be observed that, in the case of YOLO, there are no major differences in the AP, so it cannot be concluded that the preprocessing represents an important modification for the models trained with YOLO.

In analyzing the results with Fast R-CNN, there are significant differences between the different preprocessing techniques. With this model, it can be seen how the preprocessing algorithm based on Dilate and Erode operations achieves a superior improvement compared to the algorithm based on Grow Cut. This difference between the two methods may be because the grow-cut algorithm removes a larger number of pixels from the image, which in some cases may lead to the elimination of part of the area needed to detect and classify tumors.

Furthermore, if the AP of the preprocessing including only the Erode and Dilate-based algorithm is compared with the preprocessing using the two techniques, it can be observed that the Erode and Dilate-based algorithm can achieve better results alone and is more effective without combining with the other technique. This also points out that the Grow Cut-based technique can be inefficient, with some images of the model, causing crucial information to be lost for tumor detection and classification.

Additionally, it is worth noting that the results indicate a strong correlation between the preprocessing method and the complexity of the object detection model used. While YOLO-based architectures exhibit resilience to different preprocessing methods—showing relatively small variations in AP—the Faster R-CNN model experiences a drastic improvement when high-quality preprocessing is applied. This suggests that preprocessing has a greater impact on two-stage detectors like Faster R-CNN, where precise region proposals depend heavily on the clarity and consistency of input images.

One explanation for this phenomenon is the difference in feature extraction strategies between YOLO and Faster R-CNN. YOLO, as a single-stage detector, directly predicts object locations and classifications, making it inherently robust to minor distortions introduced by preprocessing. Conversely, Faster R-CNN relies on a Region Proposal Network (RPN) that requires well-defined edges and object boundaries to generate accurate bounding boxes. As seen in the results, preprocessing that enhances contrast and removes noise significantly benefits RPN-based models like Faster R-CNN.

Moreover, the results indicate that preprocessing not only affects detection accuracy but also has implications for computational efficiency. Faster R-CNN, despite achieving the highest AP scores, is known for its higher computational cost compared to YOLO models. However, when preprocessing techniques like Erode and Dilate are applied, the improvement in Faster R-CNN's accuracy could justify its computational expense, particularly in clinical applications where precise tumor detection is critical.

Focusing on the main objective of this article, which is the role of preprocessing algorithms in the segmentation and detection of breast cancer, we can extract the main argument that preprocessing in this medical image is of special importance. The results of the best model evidence this –Fast R-CNN–, that in the case of AP50, we can go from 0.4% with the original model to 91.1% in the best of the scenarios proposed in this experimentation. Another important discussion to keep in mind is that in addition to the importance of including a preprocessing algorithm, choosing which algorithm is the best is also extremely important. In our experimentation, the best-case scenario achieves an efficiency of 91.1% by applying the best algorithm, but a poor choice of preprocessing algorithm leaves the AP50 value at 77.9% or even 57% using the same data set and the same model.

Finally, these results highlight a crucial consideration for the integration of AI-based mammography analysis into clinical workflows. Given that preprocessing significantly impacts detection accuracy, it is imperative to standardize preprocessing pipelines when training AI models for medical applications. Inconsistent preprocessing across datasets could lead to significant performance variations when models are deployed in real-world settings.

5 Conclusion

This paper addresses the problem of breast cancer detection, segmentation, and classification to study the importance of model selection and preprocessing algorithms in developing techniques to automate this task. There were two main objectives: i) to determine the importance of the use of quality preprocessing algorithms and ii) to determine which Deep Learning model architectures are capable of achieving the best results. Different versions of the YOLO model and the Fast R-CNN model have been evaluated to address the objectives. Each model has been trained in different scenarios corresponding to applying different preprocessing algorithms to evaluate the differences between the different techniques.

Experimental results have shown that Fast R-CNN represents an architecture that obtains better results than YOLO. As for the preprocessing algorithms, the experimental results show that using a good preprocessing technique can improve up to 90% of AP. Furthermore, the importance of which technique is used becomes apparent as the difference in AP between different techniques can be up to 40%.

Future work in this line of development will be to explore and propose more dataset preprocessing algorithms that follow the motivation of the Erode and Dilate algorithm, as it is the one that has given us the best results in this comparison. The motivation of this future study is to try to increase the AP values obtained, which currently stand at 91.1%. Additionally, future research should investigate the integration of data augmentation techniques alongside preprocessing. Since deep learning models often struggle with limited training data, augmentation strategies such as geometric transformations, contrast adjustments, and synthetic image generation could further enhance model robustness and generalization in breast cancer detection.

References

1. American Cancer Society: Breast cancer facts & figures 2024-2025. https://www.cancer.org/ (2024), accessed: October 25, 2024
2. Avcı, H., Karakaya, J.: A novel medical image enhancement algorithm for breast cancer detection on mammography images using machine learning. Diagnostics **13**(3), 348 (2023)
3. Bajcsi, A., Chira, C.: Textural and shape features for lesion classification in mammogram analysis. In: Hybrid Artificial Intelligent Systems. pp. 755–767. Springer Nature Switzerland, Cham (2023)
4. Beeravolu, A.R., Azam, S., Jonkman, M., Shanmugam, B., Kannoorpatti, K., Anwar, A.: Preprocessing of breast cancer images to create datasets for deep-cnn. IEEE Access **9**, 33438–33463 (2021)
5. Bhatt, U., Shams, Z.: Trust in Artificial Intelligence: Clinicians Are Essential, pp. 127–141. Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-81030-6_10
6. Everingham, M., Eslami, S.M.A., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. International Journal of Computer Vision **98**(1), 111–136 (Jun 2015). https://doi.org/10.1007/s11263-014-0733-5
7. Ghosh, P., Antani, S.K., Long, L.R., Thoma, G.R.: Unsupervised grow-cut: Cellular automata-based medical image segmentation. In: 2011 IEEE First International Conference on Healthcare Informatics, Imaging and Systems Biology. pp. 40–47 (2011). https://doi.org/10.1109/HISB.2011.44
8. Heath, M., Bowyer, K., Kopans, D., Moore, R., Kegelmeyer, W.P.: The digital database for screening mammography. In: Yaffe, M.J. (ed.) Proceedings of the Fifth International Workshop on Digital Mammography. pp. 212–218. Medical Physics Publishing (2001)

9. Jocher, G., Chaurasia, A., Qiu, J.: YOLO by Ultralytics. https://github.com/ultralytics/ultralytics (2023), accessed: June 20, 2024
10. Jocher, G., Qiu, J.: Ultralytics yolo11. https://github.com/ultralytics/ultralytics (2024), accessed: January 20, 2025
11. Lu, H.C., Loh, E.W., Huang, S.C.: The classification of mammogram using convolutional neural network with specific image preprocessing for breast cancer detection. In: 2019 2nd international conference on artificial intelligence and big data (ICAIBD). pp. 9–12. IEEE (2019)
12. M. Omer, A., Elfadil, M.: Preprocessing of digital mammogram image based on otsu's threshold. American Scientific Research Journal for Engineering, Technology, and Sciences **37**(1), 220–229 (Nov 2017), https://asrjetsjournal.org/index.php/American_Scientific_Journal/article/view/3476
13. Marginean, R., Andreica, A., Diosan, L., Bálint, Z.: Autonomous image segmentation by competitive unsupervised growcut. In: 2019 21st International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC). pp. 313–319 (09 2019). https://doi.org/10.1109/SYNASC49474.2019.00051
14. Marinescu, A.I., Zoltán, B., Diosan, L., Andreica, A.: Dynamic autonomous image segmentation based on grow cut. In: European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESAAN 2018). pp. 67–72 (2018)
15. Murphy, A.: Craniocaudal view. https://radiopaedia.org/articles/craniocaudal-view (2023)
16. Murphy, A.: Mediolateral oblique view. https://radiopaedia.org/articles/mediolateral-oblique-view (2023)
17. Padilla, R., Passos, W.L., Dias, T.L.B., Netto, S.L., da Silva, E.A.B.: A comparative analysis of object detection metrics with a companion open-source toolkit. Electronics **10**(3) (2021). https://doi.org/10.3390/electronics10030279
18. Ramani, R., Vanitha, N.S., Valarmathy, S.: The pre-processing techniques for breast cancer detection in mammography images. International Journal of Image, Graphics and Signal Processing **5**(5), 47 (2013)
19. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. Advances in neural information processing systems **28** (2015)
20. Suzuki, S., be, K.: Topological structural analysis of digitized binary images by border following. Computer Vision, Graphics, and Image Processing **30**(1), 32–46 (1985). https://doi.org/https://doi.org/10.1016/0734-189X(85)90016-7
21. Vezhnevets, V., Konushin, V.: "growcut" - interactive multi-label nd image segmentation by cellular automata. Graphicon **1** (11 2004)
22. Wang, C.Y., Yeh, I.H., Liao, H.Y.M.: Yolov9: Learning what you want to learn using programmable gradient information. arXiv preprint arXiv:2402.13616 (2024)
23. Warren, L.M., Given-Wilson, R.M., Wallis, M.G., Cooke, J., Halling-Brown, M.D., Mackenzie, A., Chakraborty, D.P., Bosmans, H., Dance, D.R., Young, K.C.: The effect of image processing on the detection of cancers in digital mammography. American journal of roentgenology **203**(2), 387–393 (2014)
24. World Health Organization: Breast cancer: Key facts. https://www.who.int/news-room/fact-sheets/detail/breast-cancer (2024), accessed: October 25, 2024