

## Deep learning-based classification of breast cancer molecular subtypes from H&E whole-slide images

Masoud Tafavvoghi <sup>a,\*</sup>, Anders Sildnes <sup>b</sup>, Mehrdad Rakaee <sup>c,e,f</sup>, Nikita Shvetsov <sup>b</sup>, Lars Ailo Bongo <sup>b</sup>, Lill-Tove Rasmussen Busund <sup>c,d</sup>, Kajsa Møllersen <sup>a</sup>

<sup>a</sup> Department of Community Medicine, Uit The Arctic University of Norway, Tromsø, Norway

<sup>b</sup> Department of Computer Science, Uit The Arctic University of Norway, Tromsø, Norway

<sup>c</sup> Department of Medical Biology, Uit The Arctic University of Norway, Tromsø, Norway

<sup>d</sup> Department of Clinical Pathology, University Hospital of North Norway, Tromsø, Norway

<sup>e</sup> Department of Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, USA

<sup>f</sup> Department of Cancer Genetics, Oslo University Hospital, Oslo, Norway

### ARTICLE INFO

#### Keywords:

Breast cancer  
Classification  
Deep learning  
Molecular subtypes  
Whole-slide images

### ABSTRACT

Classifying breast cancer molecular subtypes is crucial for tailoring treatment strategies. While immunohistochemistry (IHC) and gene expression profiling are standard methods for molecular subtyping, IHC can be subjective, and gene profiling is costly and not widely accessible in many regions. Previous approaches have highlighted the potential application of deep learning models on hematoxylin and eosin (H&E)-stained whole-slide images (WSIs) for molecular subtyping, but these efforts vary in their methods, datasets, and reported performance. In this work, we investigated whether H&E-stained WSIs could be solely leveraged to predict breast cancer molecular subtypes (luminal A, B, HER2-enriched, and Basal). We used 1433 WSIs of breast cancer in a two-step pipeline: first, classifying tumor and non-tumor tiles to use only the tumor regions for molecular subtyping; and second, employing a One-vs-Rest (OvR) strategy to train four binary OvR classifiers and aggregating their results using an eXtreme Gradient Boosting model. The pipeline was tested on 221 hold-out WSIs, achieving an F1 score of 0.95 for tumor vs non-tumor classification and a macro F1 score of 0.73 for molecular subtyping. Our findings suggest that, with further validation, supervised deep learning models could serve as supportive tools for molecular subtyping in breast cancer. Our codes are made available to facilitate ongoing research and development.

### Introduction

Breast cancer accounts for 12.5% of all diagnosed cancer types globally, with around 2.3 million new cases and 685,000 fatalities annually, and is expected to grow to 3 million newly diagnosed cases and 1 million deaths by 2040.<sup>1</sup> Breast cancer is a heterogeneous disease, and its outcome depends on patients' demographic factors and tumor characteristics, including the crucial distinction among molecular subtypes, which play a significant role in determining treatment strategies. Broadly, breast cancer has four molecular subtypes: luminal A (LumA), luminal B (LumB), HER2-enriched (HER2), and basal-like (BL). Normally, BL tumors exhibit higher rates of recurrence during the initial 5 years following detection and treatment, but they show higher response to chemotherapy. On the other hand, luminal cancers, accounting for 60%–70% of all breast cancers,<sup>2</sup> respond poorly to chemotherapy, and LumA tumors have lower early recurrence compared to other breast cancer molecular subtypes.<sup>3,4</sup> Therefore, identifying the molecular subtypes of breast cancer is crucial for treatment decisions.

Currently, gene expression profiling serves as a new technology for breast cancer molecular subtyping, which is substantially more expensive and not available in all healthcare systems.<sup>5</sup> As a result, immunohistochemistry (IHC) staining is still widely used to classify the subtypes in clinical practice. IHC staining involves using specific antibodies to detect and visualize specific proteins' presence, localization, and abundance within breast cancer tissue samples. The IHC staining is typically performed for four key biomarkers: estrogen receptor (ER), progesterone receptor (PR), human epidermal growth factor receptor 2 (HER2), and antigen Ki-67. Based on the results of these four stainings, breast cancer can be classified into four main molecular subtypes: LumA, LumB, HER2, and BL.<sup>6</sup>

In recent years, deep learning has emerged as a transformative technology in many fields, notably in medical image analysis.<sup>7,8</sup> As a branch of machine learning, deep learning uses neural networks with multiple layers to identify complex patterns in raw data without manual feature extraction, a major improvement over traditional methods. This capability is especially valuable in medical imaging, where it processes large volumes of data to

\* Corresponding author.

E-mail address: [masoud.tafavvoghi@uit.no](mailto:masoud.tafavvoghi@uit.no) (M. Tafavvoghi).

deliver precise and automated analyses. Deep learning has demonstrated exceptional performance in histopathology analysis, a critical aspect of cancer diagnosis and research. It can accurately identify subtle features in histopathological slides, such as cell morphology, tissue structures, and biomarker expressions, with high precision.<sup>9–11</sup> Such breakthroughs not only streamline the diagnostic process but also hold the potential to improve the accuracy and reproducibility of results, ultimately benefiting patient care and advancing our understanding of complex diseases like breast cancer.

Traditional methods for classifying breast cancer molecular subtypes rely heavily on histopathological examination, which is often time-consuming, subjective, and sometimes inconsistent in interpretation.<sup>12</sup> Additionally, new technologies, while promising, are indeed expensive and may not be readily available in many countries and healthcare systems. In contrast, hematoxylin and eosin (H&E) staining, as the gold-standard in histopathology,<sup>13</sup> offers a more accessible and cost-effective approach. This technique is widely recognized for its reliability in tissue characterization and pathology diagnosis.

In the field of digital pathology, several studies have utilized H&E-stained histopathological images for classifying breast cancer molecular subtypes, each facing certain limitations. Couture et al.<sup>14</sup> and Jaber et al.<sup>15</sup> demonstrated the potential of deep learning for predicting BL and non-BL subtypes, highlighting its significant clinical implications, especially in resource-limited settings. However, their models faced challenges such as misclassification risks due to subtype heterogeneity and the presence of non-cancerous tiles (also known as patches) in cancer-rich clusters. In a similar vein, Abbasi et al.<sup>16</sup> noted the variability in model performance across different scanners, underscoring the necessity for equipment-specific model tuning. Liu et al.<sup>17</sup> explored the use of weakly supervised learning models on a private image dataset to classify breast cancer subtypes, concluding that while AI can aid preliminary screening, it cannot yet fully replace traditional human analysis.

Expanding on these aforementioned studies, we hypothesize that H&E-stained histopathology images contain sufficient information to classify breast cancer molecular subtypes by exhibiting different morphological patterns in the breast tissue. To evaluate this hypothesis, we have developed a multi-stage model that labels H&E whole-slide images (WSIs) after performing non-relevant tile exclusion, color normalization, and tile classification. Additionally, we enriched our analysis with an expanded dataset combining slides from publicly available datasets, enhancing the reproducibility of our results.

Building on this hypothesis, this study aims to investigate whether H&E-stained histopathology images contain sufficient information to classify molecular subtypes without carrying out additional analyses like IHC staining or gene expression profiling. This approach can offer cost and time efficiency, simpler diagnostics, wider accessibility, reduced high-tech reliance, better resource use, and suitability for large-scale research.

## Methods

The procedure of classification of breast cancer molecular subtypes in this article consists of two main parts. In the first part, we trained a deep learning model for classifying tumor and non-tumor tiles in a WSI, intending to utilize only tumor regions for the classification of molecular subtypes. In the second part, we trained separate classifiers for breast cancer molecular subtyping. In the following subsections, we delve into descriptions of utilized datasets, preprocessing steps, and training of our models. The CNN classifier codes used in this study were adapted from the repository by Foersch et al.<sup>18</sup> and modified as needed. The codes are available at [https://github.com/uit-hdl/BC\\_MolSubtyping](https://github.com/uit-hdl/BC_MolSubtyping).

### Datasets

#### Classification of tumor and non-tumor regions

In this study, we used two publicly available sets of data, one for classifying tumor and non-tumor tiles in H&E-stained WSIs and the other set for the classification of breast cancer molecular subtypes (Table 1). To develop

**Table 1**  
Characteristics of breast H&E WSIs used in this study.

Dataset	# WSIs	Source	Pixel size [ $\mu\text{m}/\text{pixel}$ ]	Scanner type	Image format
<i>Classification of breast tumor and non-tumor image tiles</i>					
TCGA-BRCA	195	USA	0.25, 0.50	Variant	svs
BRACS	129	Italy	0.25, 0.50	Aperio AT2	svs
<i>Classification of breast cancer molecular subtypes</i>					
TCGA-BRCA	980	USA	0.25, 0.50	Variant	svs
CPTAC-BRCA	382	USA	0.25, 0.50	Variant	svs
HER2-Warwick	71	UK	0.23	NanoZoomer C9600	ndpi

and test a binary deep learning model for tumor and non-tumor classification of image tiles, we used 195 WSIs of TCGA-BRCA<sup>19</sup> and 129 WSIs of BRACS<sup>20</sup> datasets. The BRACS dataset has annotated regions of interest that can be used to extract tumor tiles from WSIs. The TCGA-BRCA dataset has no officially published annotations of regions in WSIs. However, there are 195 WSIs from TCGA-BRCA in DRYAD<sup>21</sup> dataset that have annotated tumor regions. Because both images and annotations in the DRYAD dataset are downsized to a 1/10 scale, we replicated the annotated regions for the full-size WSIs and then extracted image tiles from those areas.

#### Classification of breast cancer molecular subtypes

The selection of datasets for classifying breast cancer molecular subtypes was based on the availability of pertinent labels for WSIs. According to a review article we published earlier on publicly available datasets of breast cancer histopathology images,<sup>22</sup> there were only two datasets featuring these specific labels: TCGA-BRCA and CPTAC-BRCA.<sup>23</sup> From 1139 available WSIs in the TCGA-BRCA dataset, we acquired 980 WSIs labeled with molecular subtypes and excluded 159 WSIs that either lacked labels or were categorized as *Normal-like* tumors. This was the only selection criterion we applied to the datasets used for the classification of breast cancer molecular subtypes. In the CPTAC-BRCA dataset, there are 653 WSIs of breast tissue. Of these, 382 WSIs are labeled with molecular subtypes, and the remaining 271—either unlabeled or labeled as *Normal-like*—were excluded from this study. In both TCGA-BRCA and CPTAC-BRCA datasets, the major and minor classes were LumA and HER2, with an overall share of 50.1% and 8.6%, respectively, which can cause significant dataset imbalance. To mitigate this imbalanced distribution of classes, we added the HER2-Warwick dataset,<sup>24</sup> which has 86 H&E-stained WSIs of invasive breast carcinomas from 86 patients, 71 of those with positive HER2 expression scores that were used in our study (Table 1). Detailed distribution of breast molecular subtypes in each dataset is presented in Table 2.

### Preprocessing

#### Classification of tumor and non-tumor regions

Typically, WSIs are extremely large, containing billions of pixels, and cannot be directly fed into any deep learning model. Therefore, regions of interest within the WSIs are divided into tiles to be compatible with these algorithms. For the classification of tumor and non-tumor tiles, we used QuPath<sup>25</sup> software to make and extract non-overlapping tiles with size  $512 \times 512$  at  $0.5 \mu\text{m}/\text{pixel}$  magnification from the annotated tumor regions in TCGA-BRCA and BRACS WSIs. This yielded 38,392 tumor tiles from the TCGA-BRCA and BRACS datasets, sufficient to serve as the tumor class for training a binary classifier. Additionally, we extracted 37,407 tiles from the non-tumor areas in the same WSIs to have roughly balanced tumor and non-tumor classes. The non-tumor class comprised normal tissues (including normal epithelium), folded tissue areas, marker signs, and white areas on the slides to prevent extra steps to remove low-quality and white tiles for the classification of breast molecular subtypes in the subsequent steps (Fig. S1). All 75,799 tiles were then split into 70%, 15%, and 15% for training, validation, and testing, respectively, ensuring that tiles from each WSI were assigned exclusively to one of these sets. This split also maintained the balance of classes in each set. Fig. 1 illustrates the workflow of

**Table 2**

Distribution of WSIs in the four classes used for the classification of breast molecular subtypes.

Dataset	Labels source	LumA	LumB	HER2	BL	Excluded
TCGA-BRCA	PAM50	507	219	78	176	122 unlabeled and 37 normal-like WSIs
CPTAC-BRCA	PAM50	176	44	39	123	258 unlabeled and 13 normal-like WSIs
HER2-Warwick	IHC	0	0	71	0	15 WSIs without positive HER2 expression scores
Total	–	683	263	188	299	445

preprocessing steps for classifying breast H&E images into tumor and non-tumor classes.

#### Classification of breast cancer molecular subtypes

Data preprocessing for the classification of breast cancer molecular subtypes was more comprehensive. In the first step, we used QuPath to detect the tissue areas in 1433 WSIs. Detected areas were then divided into tiles of size  $512 \times 512$  pixels, all extracted at a consistent magnification of  $20 \times$ , regardless of their highest available magnification level. All extracted tiles were stored in TIFF format without data compression. For LumA, LumB, and BL classes, we extracted the tiles without overlapping. However, to increase the number of extracted tiles (instances) for the minority class, we set vertical and horizontal overlaps of 64 pixels for the HER2 WSIs. We employed this approach to maximize the amount of training data, as existing literature suggests that the performance of deep learning models improves with larger datasets.<sup>26–28</sup> Another reason for this choice was that WSIs in the HER2-Warwick dataset originate from biopsies rather than tissue resections, resulting in fewer tiles compared to surgical resections (Fig. S2).

Because normal areas and artifacts in the image do not contribute to our classification task, we chose to use only tumor tiles. However, most of the WSIs in the TCGA-BRCA dataset and all WSIs of the CPTAC-BRCA and HER2-Warwick datasets lacked annotations of tumor areas. Therefore, to take only tumor tiles, we fed all 3,571,651 extracted tiles to the earlier trained binary tumor/non-tumor classifier to determine the likelihood of each tile belonging to the tumor class. Following that, to create a balanced dataset with a nearly equal number of tiles in each of four classes of breast cancer molecular subtypes, namely LumA, LumB, HER2, and BL, we used the minor class (HER2) as the reference class with 278,675 tumor tiles and balanced the four breast cancer classes based on that. Because the number of WSIs in each class was different, we took 441, 1180, and 1410 random tumor tiles per WSI from LumA, LumB, and BL classes, respectively, to have a roughly equal distribution of tiles among each class (Fig. 2). For a detailed illustration of data partitioning, see Fig. S3. It is important to note that the actual counts of selected tumor tiles for model development differ from the expected values. This discrepancy arose because many WSIs contained small tumor regions, resulting in fewer tumor tiles than specified for each class.

The selected tumor tiles were then color normalized using the Macenko<sup>29</sup> method to ensure consistency in color representation across all images and minimize potential variations caused by differences in image acquisition conditions. To normalize the color of tumor tiles, we generated a reference image using Python by picking one random tile out of the tumor tiles per 256 randomly selected WSIs (Fig. S4). Following that, to mitigate the risk of overfitting, the color-normalized tumor tiles were

divided into three distinct sets: 70% for training and validation of convolutional neural networks (CNNs), 15% for training and validation of the eXtreme Gradient Boosting (XGBoost)<sup>30</sup> model, and the remaining 15% for testing the entire pipeline on the classification of BC molecular subtypes. Each set maintained a stratified split to ensure a proportional representation of images per class, and WSIs from each patient were restricted to a single set, ensuring that different images from the same patient were not included in multiple sets.

#### Model training


##### Deep convolutional neural networks

**Classification of tumor and non-tumor regions.** To train the tumor/non-tumor classifier, we used the Inception\_V3<sup>31</sup> architecture with pre-trained weights, implemented in PyTorch (version 1.7.1 + cu110). Inception\_V3 is known for its efficiency in capturing complex hierarchical features through the use of Inception modules, which perform convolutions of various sizes ( $1 \times 1$ ,  $3 \times 3$ ,  $5 \times 5$ ) within the same layer, allowing the network to capture multi-scale features effectively. The architecture also incorporates auxiliary classifiers at intermediate layers to help propagate gradients and improve convergence during training.

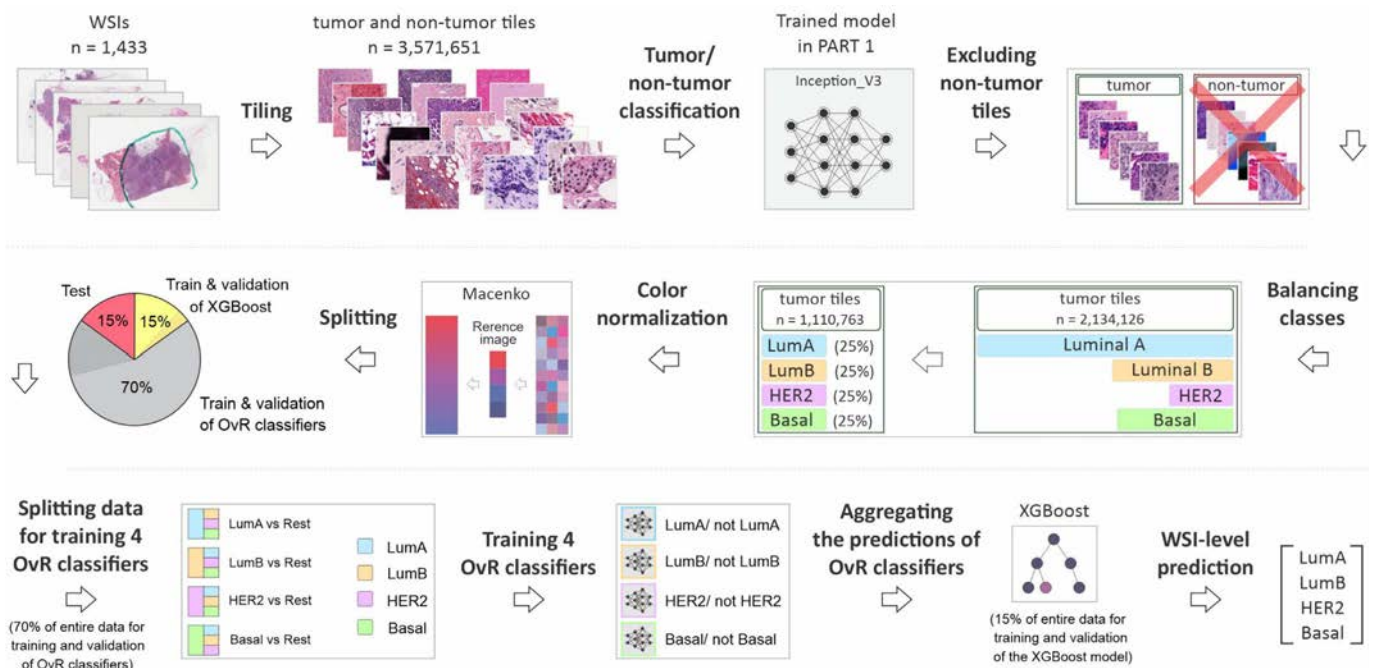
Inception\_V3's modular approach allows it to capture intricate patterns and features within histopathology images,<sup>32,33</sup> making it a suitable choice for this classification task. The auxiliary classifiers embedded within the network also aid in preventing gradient vanishing issues, which can be prevalent in deep networks.

The training was performed using an RTX-3090 GPU with 24 GB of VRAM. The hyperparameters were set as follows: a batch size of 64, a learning rate of  $1e-5$ , and a dropout rate of 0.33 to prevent overfitting. We used the ADAM optimizer, known for its robustness to sparse gradients, and the cross-entropy loss function, which is suitable for binary classification.

**Classification of breast cancer molecular subtypes.** For classifying breast cancer molecular subtypes from H&E-stained images, we trained four separate binary classifiers using the One-vs-Rest (OvR) strategy to simplify a complex multi-class classification into binary tasks. For training and validation, 70% of the selected tumor tiles were allocated (the gray part of the pie chart in Fig. 2), with 80% (56% of the total data in light gray) for training and 20% (14% of the total data in dark gray) for validation (model fine-tuning). Each classifier was trained using the ResNet-18 architecture,<sup>34</sup> with pre-trained weights, incorporating all tiles from the target subtype and one-third from each of the other three subtypes as the *rest* class, ensuring balanced binary classes (Fig. 2 and Fig.S3) for training the models.



**Fig. 1.** The workflow of training a deep binary classifier for tumor and non-tumor breast tiles.



**Fig. 2.** The workflow of classifying breast H&E WSIs into four molecular subtypes, where the results of four binary OvR classifiers were aggregated by an XGBoost model to predict the subtype in WSIs.

This architecture was chosen due to its proven effectiveness in various image classification tasks, offering a good balance between depth and computational efficiency.<sup>35,36</sup> Moreover, ResNet-18's relatively shallow depth compared to deeper variants ensures faster training and inference times without significant loss in accuracy, which is crucial when training four deep classifiers with large-scale datasets. The training was performed using the same hardware with the following hyperparameters: a batch size of 128, a learning rate of  $5e-6$ , a dropout rate of 0.33, the ADAM optimizer, and the cross-entropy loss function. Additionally, the training data underwent shuffling, along with random horizontal and vertical flips and rotations as image augmentation techniques, to enhance model training and mitigate potential bias arising from overlapping pixels in the extracted tiles.

**Thresholding.** The output of a binary classifier for the target class is a single score between 0 and 1 for each image tile. To aggregate these scores into a definitive classification for each WSI, we set a threshold for each classifier, which determines the predicted class of each tile. Using a fixed threshold of 0.5 for assigning classes to images can be sub-optimal.<sup>37,38</sup> Therefore, we employed precision-recall (PR) analysis to establish the optimum decision thresholds for each classifier. PR analysis aids in identifying a threshold that balances precision (positive-predictive value) and recall (sensitivity) effectively. The PR curve plots precision against recall, enabling the selection of a threshold that maximizes the classifier's performance. Although our classes are balanced, this approach is particularly relevant as we train four separate binary OvR classifiers, where the target class in each classifier is treated as the positive class. This allows us to fine-tune each OvR classifier for optimal performance in distinguishing the target class from all others. We used only the validation set, which constitutes 20% of the classification data (equivalent to 14% of the total tumor tiles), to adjust the threshold of OvR classifiers. The optimum thresholds for LumA, LumB, HER2, and BL classifiers were determined to be 0.434, 0.415, 0.481, and 0.424, respectively. These thresholds were applied in subsequent stages of our study.

#### XGBoost

By feeding the tumor tiles into the four binary OvR classifiers, we obtained eight scores for each tile, two from each classifier. Each pair of scores

represents the classifier's confidence in the images belonging to either the target class or the rest class. To aggregate the tile-wise predictions of four OvR CNNs to classify the molecular subtype of WSIs, we quantified the number of tiles within a WSI that had scores exceeding the optimal threshold for each class. This process involved counting the tiles classified as LumA,  $\neg$  LumA (not-LumA), LumB,  $\neg$  LumB, HER2,  $\neg$  HER2, BL, and  $\neg$  BL. Subsequently, these eight values were used as new features to train an XGBoost model, which predicts the molecular subtype of the WSIs (Fig. 3). XGBoost, an implementation of gradient-boosted decision trees, is a highly flexible and versatile machine learning tool, proven effective in a wide range of supervised learning tasks, demonstrating its ability to learn complex patterns in large volumes of data.<sup>30,39</sup>

To train and fine-tune the XGBoost model, we allocated 15% of the entire dataset (the yellow part of the pie chart in Fig. 2), which was further divided randomly into 80% for training and 20% for validation. We used the thresholds established in the earlier classification step to count the tiles predicted as either the target class or the rest class for all four breast cancer molecular subtypes in each WSI.

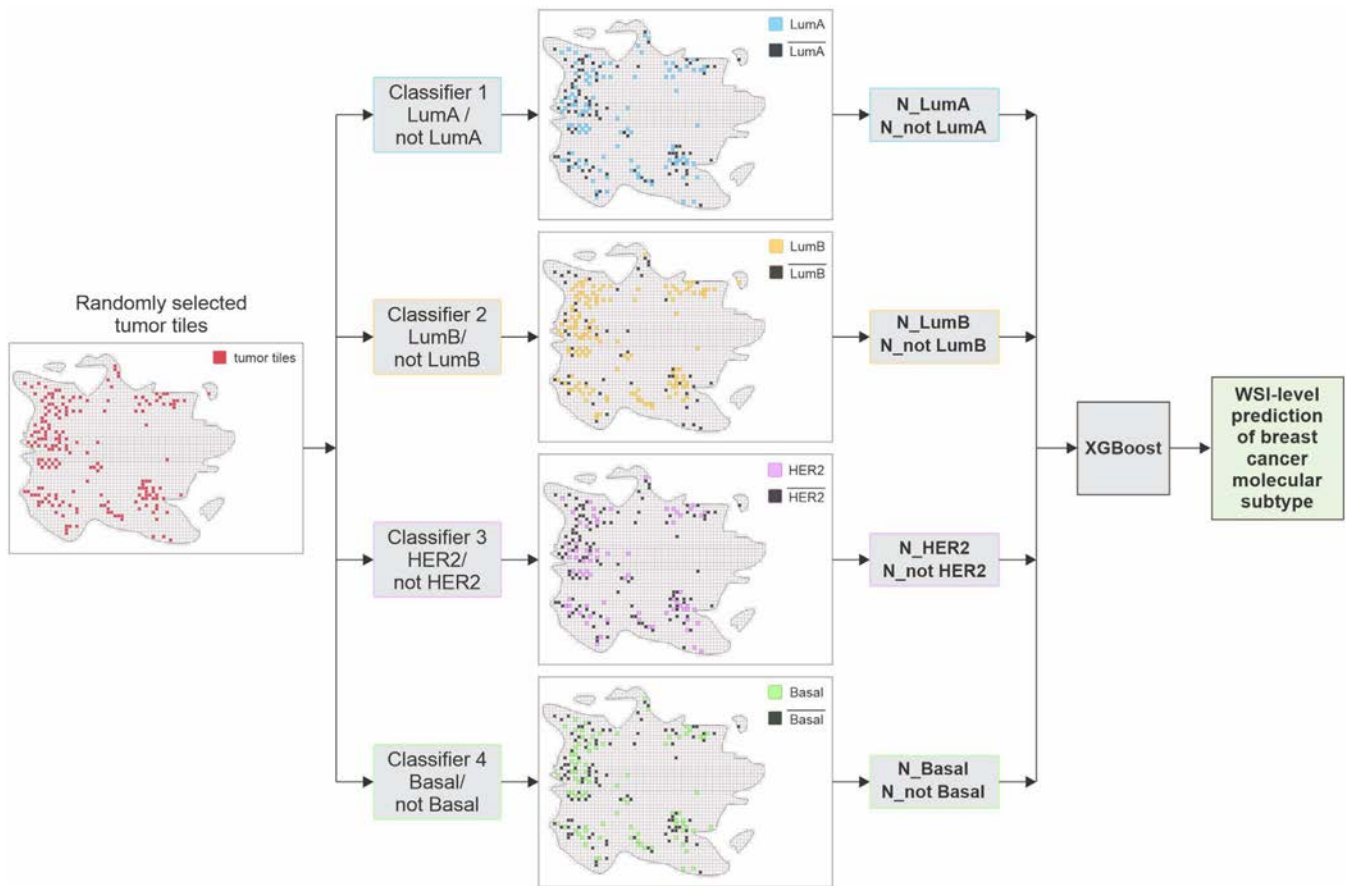
## Results

### Classification of tumor and non-tumor tiles

The model trained on 75,799 tiles labeled as 'tumor' or 'non-tumor' from TCGA-BRCA and BRACS datasets was evaluated using the F1 score, achieving a value of 0.954. The F1 score provides a balance between precision and recall (sensitivity), making it a suitable metric for our classification task. Such a high F1 score indicates that the model is reliable and performs very well in distinguishing between tumor and non-tumor tiles. Table 3 demonstrates additional performance metrics of our model on the test set. In addition, the model's overall accuracy was 0.955, indicating a high level of overall correct predictions across both classes (Fig. 4).

### Classification of breast cancer molecular subtypes

We evaluated the entire pipeline on a hold-out test set (15% of the entire data, shown as the red part in the pie chart of Fig. 2) consisting of 221 H&E WSIs that had not been used in any part of the workflow.



**Fig. 3.** Aggregation of four OvR binary classifiers for predicting breast cancer molecular subtypes in a WSI. Randomly selected tumor tiles are independently processed by each classifier. The counts of tiles classified into target and non-target classes are used as features in an XGBoost model, which integrates the results from all four classifiers to determine the molecular subtype.

Table 4 presents the classification results for breast cancer molecular subtypes within the test set at the WSI level, obtained using an XGBoost model that aggregates predictions from four binary OvR models.

The performance of our model varies significantly across different breast cancer subtypes. We present the F1 score as our primary metric, commonly used in multi-class classification problems to balance precision and sensitivity, thereby providing a more comprehensive view of model performance. For the LumA subtype, the model achieved an F1 score of 0.922, indicating strong performance in correctly identifying cases of this subtype. Conversely, the HER2 subtype presented more of a challenge, with an F1 score of 0.545. This lower score suggests that whereas the model is particularly effective at ruling out false positives (high specificity), it frequently misses actual cases of HER2 (low sensitivity). For the LumB and BL subtypes, the F1 scores were 0.742 and 0.698, respectively, suggesting a moderate capability of the model to differentiate these subtypes. However, these scores also highlight potential limitations in distinguishing these subtypes from others or each other.

Given the imbalance in class distribution, with LumA comprising a significant majority of the WSIs in the test set (46%), to ensure a balanced evaluation, we employed a macro-averaging approach to report the accuracy. Macro-averaging calculates metrics independently for each class and then computes their average, thereby treating all classes with equal importance, irrespective of their frequency. The individual accuracy rates for

**Table 3**  
Classifier's performance metrics on breast histopathology tumor/non-tumor tiles.

Class	F1 score	Precision	Sensitivity	Specificity	Accuracy
Tumor	0.954	0.963	0.945	0.965	0.965

LumA, LumB, HER2, and BL were 0.931, 0.837, 0.469, and 0.667, respectively, resulting in a macro-average accuracy of 0.726.

Fig. 5 shows our classifier's confusion matrix and PR curves. The confusion matrix presents a detailed insight into the model's ability to classify breast cancer molecular subtypes at the WSI level. For LumA, the model demonstrates a high accuracy, correctly identifying 94 out of 101 instances, reflecting a solid capability to distinguish this subtype from others with minimal confusion, as indicated by the misclassification of only seven LumA cases as LumB ( $n = 3$ ) and HER2 ( $n = 4$ ). LumB WSIs were classified with lower accuracy; out of 43 samples, 36 were correctly identified. Most misclassifications occurred with the BL subtype ( $n = 5$ ), suggesting possible similarities in the features recognized by the model between these two subtypes. The classification of the HER2 subtype poses significant challenges, with only 13 correct predictions out of 32 samples, highlighting a critical area of weakness in the model's performance. Misclassifications are broadly distributed across all other subtypes, implying a fundamental difficulty in isolating defining features of HER2 within the model's current framework. BL subtype classification shows a better outcome, with 30 out of 45 cases correctly identified. However, misclassifications into LumB and HER2 suggest overlapping characteristics or insufficient specificity in the model's learning parameters.

Due to class imbalance in our dataset, we opted to use PR curves instead of receiver operating characteristic (ROC) curves. PR curves are more informative than ROC curves for imbalanced datasets as they focus on the classifier's performance with respect to the minority class, offering a clearer picture of the trade-offs between precision and recall.

The PR curves with 95% confidence intervals for each breast cancer molecular subtype, presented in Fig. 5, provide additional insights into the model's performance variability. The PR curve for the LumA subtype

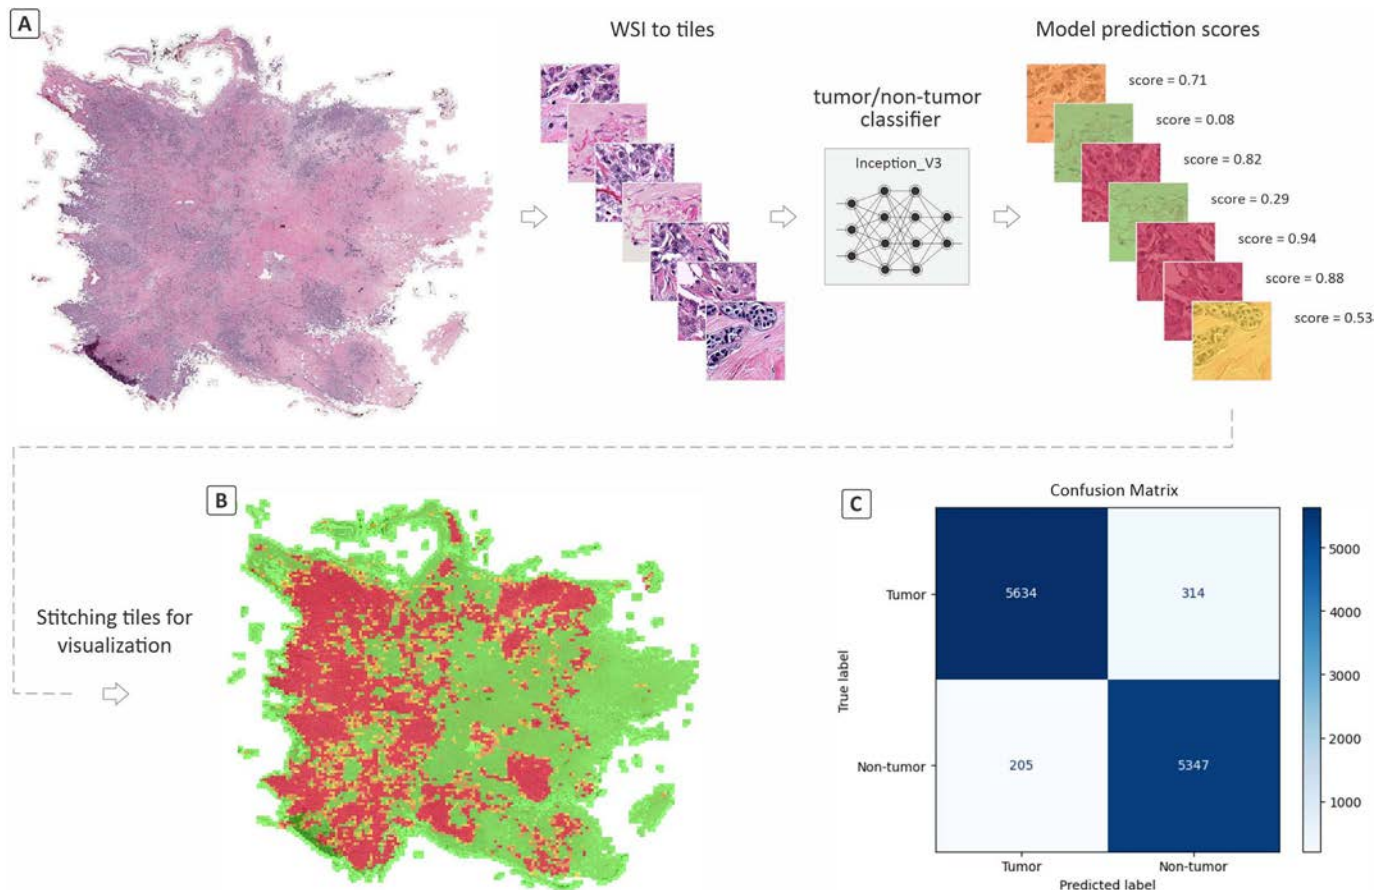


Fig. 4. Performance of the classification model in distinguishing between tumor and non-tumor tiles. (A) Original WSI from the TCGA-BRCA dataset. (B) Recreated WSI by stitching the tiles with an overlay heatmap to illustrate tumor areas in the WSI. Red and yellow areas in the stitched image show tumor areas with high- and low confidence scores, while green areas show non-tumor areas. (C) Confusion matrix showing the actual labels versus the model's predictions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

**Table 4**  
WSI level classification metrics of breast cancer molecular subtypes with 95% confidence intervals.

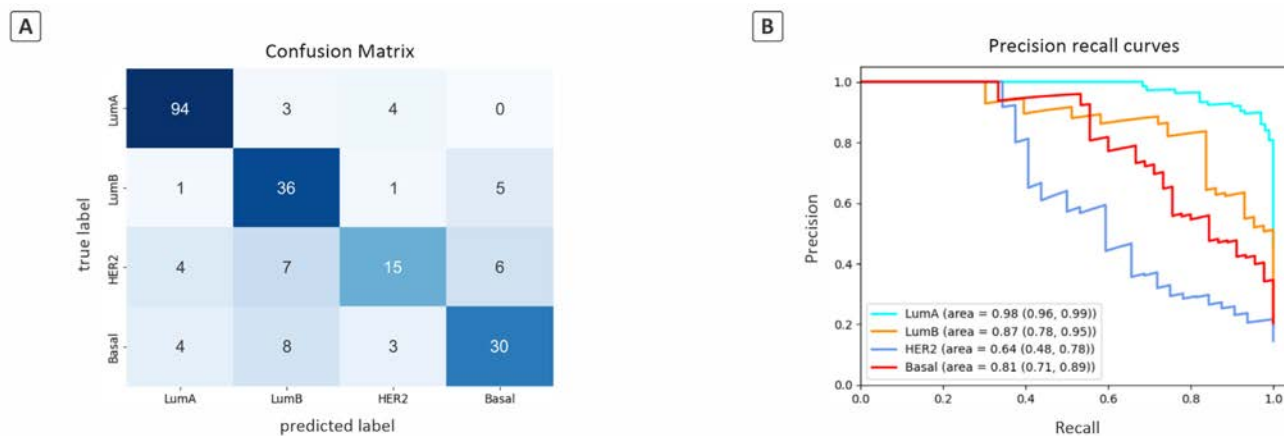
Class	#WSIs	F1 score	Precision	Recall	Specificity
Luminal A	101	0.922 (0.880, 0.956)	0.913 (0.856, 0.963)	0.931 (0.873, 0.973)	0.925 (0.876, 0.969)
Luminal B	43	0.742 (0.629, 0.833)	0.667 (0.537, 0.788)	0.837 (0.714, 0.941)	0.899 (0.852, 0.941)
HER2-enriched	32	0.545 (0.364, 0.690)	0.652 (0.435, 0.840)	0.469 (0.286, 0.647)	0.958 (0.927, 0.984)
Basal-like	45	0.698 (0.571, 0.800)	0.732 (0.587, 0.868)	0.667 (0.532, 0.792)	0.938 (0.895, 0.972)
Macro-average	-	0.727	0.741	0.726	0.930

displayed an exemplary area under the curve (AUPRC) of 0.98, supported by a tight confidence interval ranging from 0.96 to 0.99. This result underscores the model's consistent and precise capability in identifying the LumA subtype. Conversely, the LumB subtype recorded an AUPRC of 0.87 with a broader confidence interval between 0.78 and 0.95, reflecting reliable yet slightly variable performance, which suggests minor challenges in accurate classification. The model exhibited more pronounced difficulties with the HER2 subtype, where the AUPRC was significantly lower at 0.64, and the confidence interval widened from 0.48 to 0.78, highlighting substantial inconsistency and overlap with other subtypes in its classification. Lastly, the Basal subtype attained an AUPRC of 0.81 and a confidence interval from 0.71 to 0.89, indicating a moderate level of performance with some degree of variability. This thorough analysis of the PR curves and their respective confidence intervals builds upon the earlier discussed F1 scores and accuracy metrics to provide a detailed assessment of the model's classification effectiveness and pinpoint areas needing further improvement.

**Discussion**

In this study, we developed a supervised deep learning model to investigate whether H&E-stained histopathological images contain sufficient information for classifying breast cancer molecular subtypes.

IHC staining is fundamental for molecular subtyping of breast cancer, offering greater precision than H&E staining but at a higher cost and longer processing time. However, it is susceptible to inter-observer variability, which can lead to diagnostic discrepancies. IHC-based classification may not always correspond with gene expression profiles, with discrepancies up to 31%.<sup>40</sup> Additionally, about one out of five LumB cases are identified as HER2 positive through IHC analysis, highlighting the complexity of differentiating between these subtypes.<sup>6</sup> Therefore, achieving diagnostic information from H&E WSIs comparable to that obtained through IHC would significantly reduce the time and costs associated with the diagnostic process, including both laboratory work and pathologists' evaluation.



**Fig. 5.** Performance of the XGBoost model in classifying breast cancer molecular subtypes. (A) Confusion matrix of model predictions. (B) PR curves of the subtypes with 95% confidence intervals.

To address this, we developed a two-step process. First, we trained a tile-wise CNN classifier to use only tumor regions for the classification of breast molecular subtypes. Our model achieved an overall F1 score and accuracy of 0.954 and 0.955, respectively, indicating a high effectiveness in identifying tumor-containing tiles. This high level of performance is critical for ensuring that subsequent analyses for subtype classification are conducted only on relevant tumor tissues. The decision to use Inception\_V3 for this task was based on its robust feature extraction capabilities, which are well-suited for a smaller dataset.

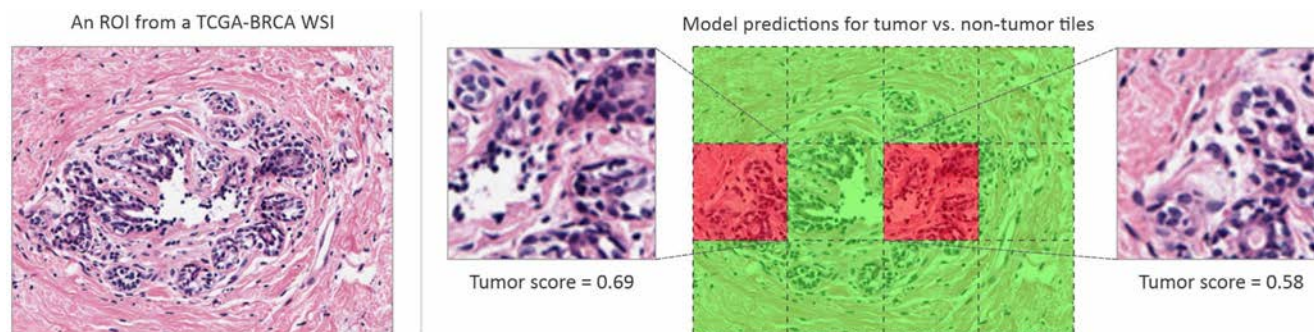
One of the limitations in the classification of tumor and non-tumor tiles is the potential presence of both classes within a single tile. This issue is inherent in histology images, especially when extracting large tiles from WSIs. Therefore, during inference, a given tile may be classified as tumor while also containing small non-tumor regions, which can introduce noise into the training data for subsequent molecular subtype classification. Additionally, class imbalance within the non-tumor classes could affect the performance of the classifier. For instance, tiles with folded tissue artifacts or those containing normal epithelium were substantially fewer than those with normal stroma, which may limit the model's performance on such cases (Fig. 6).

The main part of the workflow focused on classifying breast cancer molecular subtypes, where we used the OvR strategy to simplify the problem of multi-class classification into a series of binary classification tasks, particularly given the complexity and similarity among the histology images of different breast cancer molecular subtypes. In this approach, the goal is to distinguish a specific class from all others. One benefit of this method is its interpretability; because each class is represented by a single classifier, it allows for a better understanding of the target class by analyzing its corresponding classifier. Moreover, the aggregation of predictions from four CNN binary classifiers with an XGBoost classifier model allowed us to

leverage the strengths of binary classification while effectively addressing the complexities inherent in the classification of breast cancer molecular subtypes. The use of ResNet-18 for the CNNs in this part of the workflow was driven by its balance between computational efficiency and performance. With fewer parameters than deeper models like Inception\_V3, ResNet-18 was a practical choice for handling the large dataset of 1.1 million images of size  $512 \times 512$  pixels, ensuring reasonable training times and memory usage while still providing strong performance.

Our results revealed that the model performed differently across various subtypes with a macro F1 score of 0.72. It exhibited strong discrimination for LumA and LumB subtypes, but there were areas for improvement, especially for the HER2 class. The challenges in achieving high precision and sensitivity may be attributed to the underlying complexity of histopathological images, the variability in tissue preparation, staining procedures and scanners, and the subtle differences that separate the molecular subtypes of breast cancer in practice. The impact of factors such as tumor cellularity and heterogeneity, as well as variations in tissue preparation and WSI scanners, on model performance, could be an important area for future research, depending on the availability of data.

While a comparison of several similar studies follows, it is important to acknowledge that using deep learning models for analyzing histopathological WSIs is a relatively young field of research. In addition, the existing studies have used different datasets, which limits the direct comparability of the findings. Despite this, we present performance metrics from various studies below to offer readers a general impression of how our results stack up against others in the field. Couture et al.<sup>14</sup> trained a VGG-16 binary classifier specifically for BL and non-BL tumors, achieving an overall accuracy of 0.77, with a sensitivity of 0.78 and specificity of 0.73, using a private dataset. In contrast, our model, which is not explicitly trained to distinguish between BL and non-BL tumors, can still be evaluated in this



**Fig. 6.** An illustration of two non-tumor tiles containing normal breast epithelium that were classified as tumor. Tiles with red and green overlays indicate classification as tumor and non-tumor, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

binary context by grouping LumA, LumB, and HER2 as non-BL. In this adjusted setting, our model achieves an overall accuracy of 0.75 for BL and non-BL WSIs with a sensitivity of 0.67 and a specificity of 0.94.

Jaber et al.<sup>15</sup> applied a deep learning model to extract features from multiscale tiles of WSIs in the TCGA-BRCA dataset. Their approach involved using principal component analysis alongside gene expression data to exclude patients with tumors exhibiting heterogeneous subtype characteristics from the training data. Using a multiclass support vector machine algorithm, they achieved an accuracy of 0.66 for WSIs in the test set, which is lower than the macro-average of our multi-class model (0.73). In addition, they trained a binary classifier for BL/non-BL tumors with an AUC-ROC of 0.86 and an accuracy of 0.87 for a balanced dataset. Despite its advantages, their model still faced challenges, such as the risk of misclassification due to subtype heterogeneity and the non-cancer richness of some tiles within cancer-rich clusters.

Abbasi et al.<sup>16</sup> investigated the predictiveness of morphological features in H&E-stained tissues for the classification of breast molecular subtypes. They leveraged a specialized backbone pre-trained on H&E WSIs in a self-supervised setting for large unlabeled data. Their network showcased enhanced generalizability to unseen data from new scanner types despite the dataset's relatively modest size used for backbone pre-training. Their model achieved an average AUC-ROC of 0.73. Their results showed the inconsistent performance of two models with different backbones when tested on unseen data from different scanners, which points to the potential of model tuning specific to the imaging source.

Liu et al.<sup>17</sup> proposed a weakly supervised learning model, where they used co-teaching to reduce the effect of noise patches and multi-instance learning for the classification of WSIs. Their top multi-class classifier demonstrated an accuracy of 0.58 and a macro F1 score of 0.65; however, aggregating the results of four binary classifiers and using the weighted fusion method improved their classification metrics to an accuracy and F1 of 0.64 and 0.69, respectively. They concluded that such AI models can be used for preliminary screening, making molecular subtyping more accessible and cost-effective, although it cannot yet replace traditional pathologist analysis entirely.

The XGBoost model in our study was trained on a dataset with imbalanced classes at the WSI level. Training on imbalanced data, however, can lead to several issues, particularly affecting the model's performance for the less-represented classes. Although we balanced the data at the tile level, this approach did not compensate for the imbalances at the WSI level, where the XGBoost model was trained. This imbalance was reflected in our results, with the model performing best for the major class (LumA) and underperforming for the minor class (HER2), which experienced the most misclassifications. Recognizing these challenges, future research could explore the impact of balancing the dataset on performance metrics. Specifically, adjusting the training set to better represent minority classes like HER2 could potentially improve the model's overall accuracy.

Our evaluation of minimum tumor content for reliable classification of breast cancer molecular subtypes demonstrates an overall positive correlation between the number of tumor tiles and model accuracy, as shown in Fig. S5. However, determining a cut-off for the minimum number of tumor tiles is complex due to several factors. The number of tumor tiles per WSI in each class was adjusted to balance the number of images used for training the four binary OvR classifiers. Additionally, aggregation of OvR classifiers predictions occurs at the WSI level, where data imbalance is significant. As a result, defining a clear cut-off for minimum tumor content at the WSI level remains difficult, as it is influenced by both the tumor content and the inherent variabilities in the datasets.

While our model demonstrated potential during internal validation (the hold-out test set), external validation is critical to ensure its generalizability and robustness across different datasets from diverse institutes and scanners. Our study lacked external test set validation due to the unavailability of publicly accessible datasets with the necessary annotations for breast cancer molecular subtypes. This limitation reflects a common challenge in

this field, where high-quality, annotated data are scarce and often proprietary. External validation ensures that the model is not overfitted and performs reliably in practical applications. It also helps identify potential weaknesses when exposed to unseen data, highlighting areas that require further refinement. Demonstrating effectiveness across multiple institutions enhances credibility and facilitates broader acceptance within the research community. Additionally, the explainability of AI models remains an important consideration, as it provides insights into the decision-making process of the models. Although our work did not focus on explainability due to moderate model performance, future research should incorporate methods to enhance transparency and trust in the model's predictions.

## Conclusion

Our findings contribute to the evidence that H&E-stained WSIs contain pertinent information for classifying breast cancer molecular subtypes. Nonetheless, as this research is still emerging, our model, while promising, needs further enhancements and validation to ensure its generalisability. Future efforts could focus on refining these models to better manage the variability in histopathological images and extend their validation through external datasets.

## Declaration of competing interest

All authors declare they have no conflicts of interest.

## Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jpi.2024.100410>.

## References

- Arnold M, Morgan E, Rungay H, et al. Current and future burden of breast cancer: Global statistics for 2020 and 2040. *The Breast* 2022;66:15–23. <https://doi.org/10.1016/j.breast.2022.08.010>. ISSN 0960-9776. URL: <https://www.sciencedirect.com/science/article/pii/S0960977622001448>.
- Mir MA, Ud' Din IM. *Molecular Subtypes of Breast Cancer and Cdk Dysregulation*. Singapore: Springer Nature Singapore. 2023:133–148. [https://doi.org/10.1007/978-981-19-8911-7\\_6](https://doi.org/10.1007/978-981-19-8911-7_6). ISBN 978-981-19-8911-7.
- Sarkar DK, ed. *Breast Diseases, Guidelines for Management*. CRC Press; 2024. <https://doi.org/10.1201/9780367821982>. ISBN 9780367421281 (hbk).
- Değerli E, Öztaş NŞ, Alkan G, et al. Relationship between pathological response and molecular subtypes in locally advanced breast cancer patients receiving neoadjuvant chemotherapy. *J Chemother* 2023;35(1):29–38. <https://doi.org/10.1080/1120009X.2022.2043514>. [PMID: 35220928].
- Nguyen TMT, Le RD, Nguyen CV. Breast cancer molecular subtype and relationship with clinicopathological profiles among vietnamese women: a retrospective study. *Pathol Res Pract* 2023;250:154819. <https://doi.org/10.1016/j.prp.2023.154819>. ISSN 0344-0338. URL: <https://www.sciencedirect.com/science/article/pii/S0344033823005198>.
- Johnson KS, Conant EF, Soo MS. Molecular subtypes of breast cancer: a review for breast radiologists. *J Breast Imaging* 12 2020;3(1):12–24. <https://doi.org/10.1093/jbi/wbaa110>. ISSN 2631-6110.
- Suk H-I, Liu M, Cao X, Kim J. Editorial: advances in deep learning methods for medical image analysis. *Front Radiol* 2023;2. <https://doi.org/10.3389/fradi.2022.1097533>. ISSN 2673-8740.
- Szilágyi L, Kovács L. Special issue: artificial intelligence technology in medical image analysis. *Appl Sci* 2024;14(5). <https://doi.org/10.3390/app14052180>. ISSN 2076-3417. URL: <https://www.mdpi.com/2076-3417/14/5/2180>.
- Andrew S, Lee H, Tan X, et al. A deep learning model for molecular label transfer that enables cancer cell identification from histopathology images. *npj Precis Oncol* 2022;6(1):14. <https://doi.org/10.1038/s41698-022-00252-0>. ISSN 2397-768X.
- Farahani H, Boschman J, Farnell D, et al. Deep learning-based histotype diagnosis of ovarian carcinoma whole-slide path ology images. *Modern Pathol* 2022;35(12):1983–1990. <https://doi.org/10.1038/s41379-022-01146-z>. ISSN 0893-3952. URL: <https://www.sciencedirect.com/science/article/pii/S0893395222055107>.
- Sandarenu P, Millar EKA, Song Y, et al. Survival prediction in triple negative breast cancer using multiple instance learning of histopathological images. *Sci Rep* 2022;12(1):14527. <https://doi.org/10.1038/s41598-022-18647-1>. ISSN 2045-2322.
- Rashmi R, Prasad K, Chethana Babu KU. Breast histopathological image analysis using image process ing techniques for diagnostic purposes: a methodological review. *J Med Syst* 2021;46(1):7. <https://doi.org/10.1007/s10916-021-01786-9>. ISSN 1573-689X (Electronic), 0148-5598 (Print), 0148-5598 (Linking). URL: <https://pubmed.ncbi.nlm.nih.gov/34860316/>.



13. Tehrani KF, Park J, Chaney EJ, Tu H, Boppart SA. Nonlinear imaging histopathology: a pipeline to correlate gold-standard hematoxylin and eosin staining with modern nonlinear microscopy. *IEEE J Select Top Quant Elect* 2023;29(4 Biophotonics):6800608. <https://doi.org/10.1109/jstqe.2022.3233523>. ISSN 1077-260X (Print), 1558-4542 (Electronic), 1077-260X (Linking). URL: <https://pubmed.ncbi.nlm.nih.gov/37193134/>.
14. Couture HD, Williams LA, Geradts J, et al. Image analysis with deep learning to predict breast cancer grade, er status, histologic subtype, and intrinsic subtype. *NPJ Breast Cancer* 2018;4:30. <https://doi.org/10.1038/s41523-018-0079-1>. ISSN 2374-4677.
15. Jaber MI, Song B, Taylor C, et al. A deep learning image-based intrinsic molecular subtype classifier of breast tumors reveals tumor heterogeneity that may affect survival. *Breast Cancer Res* 2020;22(1):12. <https://doi.org/10.1186/s13058-020-1248-3>. ISSN 1465-542X.
16. Abbasi-Sureshjani S, Yüce A, Schönenberger S, et al. Molecular subtype prediction for breast cancer using H&E specialized backbone. *MICCAI Workshop on Computational Pathology*. PMLR; 2021. p. 1–9.
17. Liu H, Xu W-D, Shang Z-H, et al. Breast cancer molecular subtype prediction on pathological images with discriminative patch selection and multi-instance learning. *Front Oncol* 2022;12, 858453. <https://doi.org/10.3389/fonc.2022.858453>.
18. Foersch S, Glasner C, Woerl A-C, et al. Multistain deep learning for prediction of prognosis and therapy response in colorectal cancer. *Nat Med* 2023;29:430–439. <https://doi.org/10.1038/s41591-022-02134-1>.
19. The Cancer Genome Atlas (TCGA). Genomic Data Commons Data Portal (GDC). <https://portal.gdc.cancer.gov/projects/TCGA-BRCA>. Accessed 07 Jul. 2023.
20. BRACS: BReAst Carcinoma Subtyping. Institute of High-Performance Computing and Networking. 2020. <https://www.bracs.icar.cnr.it/>. Accessed 07 Jul. 2023.
21. Cruz-Roa A, Gilmore H, Basavanahally A, et al. High-throughput adaptive sampling for whole-slide histopathology image analysis (hashi) via convolutional neural networks: application to invasive breast cancer detection. *PLoS One* 2018;13(5):e0196828. <https://doi.org/10.5061/dryad.1g2nt41>. Accessed 07 Jul. 2023 via.
22. Tafavoghi M, Bongo LA, Shvetsov N, Busund L-TR, Møllersen K. Publicly available datasets of breast histopathology H&E whole-slide images: a scoping review. *J Pathol Inform* 2024;15, 100363. <https://doi.org/10.1016/j.jpi.2024.100363>.
23. National Cancer Institute Clinical Proteomic Tumor Analysis Consortium. *The Clinical Proteomic Tumor Analysis Consortium Breast Invasive Carcinoma Collection (CPTAC-BRCA)*. The Cancer Imaging Archive. 2020 <https://wiki.cancerimagingarchive.net/pages/viewpage.action?pageId=70227748>. [Accessed 07 Jul. 2023].
24. Her2 Scoring Contest. Tissue Image Analytics (TIA) Centre. 2016. [https://warwick.ac.uk/fac/cross\\_fac/tia/data/her2contest/](https://warwick.ac.uk/fac/cross_fac/tia/data/her2contest/). Accessed 07 Jul. 2023.
25. Bankhead P, Loughrey MB, Fernández JA, et al. QuPath: Open source software for digital pathology image analysis. *Sci Rep* 2017;7:16878. <https://doi.org/10.1038/s41598-017-17204-5>.
26. Richter ML, Byttner W, Krumnack U, Wiedenroth A, Schallner L, Shenk J. Input size matters for cnn classifiers. In: *Farkaš I, Masulli P, Otte S, Wermter S, eds. Artificial Neural Networks and Machine Learning – ICANN 2021*. Lecture Notes in Computer Science Cham: Springer; 2021. p. 11–18. [https://doi.org/10.1007/978-3-030-86340-1\\_11](https://doi.org/10.1007/978-3-030-86340-1_11).
27. Luo C, Li X, Wang L, He J, Li D, Zhou J. How does the data set affect cnn-based image classification performance?. 2018 5th International Conference on Systems and Informatics (ICSAI); 2018. p. 361–366. Nanjing, China. <https://doi.org/10.1109/ICSAI.2018.8599448>.
28. Uchida S, Ide S, Iwana BK, Zhu A. A further step to perfect accuracy by training cnn with larger data. 2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR); 2016. p. 405–410. Shenzhen, China. <https://doi.org/10.1109/ICFHR.2016.0082>.
29. Macenko M, Niethammer M, Marron JS, et al. A method for normalizing histology slides for quantitative analysis. 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro; 2009. p. 1107–1110. <https://doi.org/10.1109/ISBI.2009.5193250>.
30. Chen T, Guestrin C. Xgboost: a scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. URL: <https://api.semanticscholar.org/CorpusID:4650265>.
31. Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z. Rethinking the inception architecture for computer vision. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2016. p. 2818–2826.
32. Xie J, Liu R, Luttrell J, Zhang C. Deep learning based analysis of histopathological images of breast cancer. *Front Genet* 2019;10. <https://doi.org/10.3389/fge.ne.2019.00080>. ISSN 1664-8021.
33. Diao S, Luo W, Hou J, et al. Deep multi-magnification similarity learning for histopathological image classification. *IEEE J Biomed Health Inform* 2023;27(3):1535–1545. <https://doi.org/10.1109/JBHI.2023.3237137>.
34. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. p. 770–778. <https://doi.org/10.1109/CVPR.2016.90>.
35. Tan Y, Feng Lj, Huang Yh, Xue Jw, Long Ll, Feng Z-B. A comprehensive radiopathological nomogram for the prediction of pathological staging in gastric cancer using ct-derived and wsi-based features. *Translat Oncol* 2024;40:101864. <https://doi.org/10.1016/j.tranon.2023.101864>. ISSN 1936-5233. URL: <https://www.sciencedirect.com/science/article/pii/S1936523323002504>.
36. Sun Q, Zhong W, Zhou J, Lai C, Teng X, Lai M. *Rcdpia: A Renal Carcinoma Digital Pathology Image Annotation Dataset Based on Pathologists*. 2024.
37. Saito T, Rehmsmeier M. The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS One* 2015;10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>.
38. Brownlee J. ROC Curves and Precision-Recall Curves for Imbalanced Classification. URL: <https://machinelearningmastery.com/roc-curves-and-precision-recall-curves-for-imbalanced-classification/>; 2018. Accessed: 2024-05-16.
39. Liew XY, Hameed N, Clos J. An investigation of xgboost-based algorithm for breast cancer classification. *Mach Learn Appl* 2021;6:100154. <https://doi.org/10.1016/j.mlwa.2021.100154>. ISSN 2666-8270. URL: <https://www.sciencedirect.com/science/article/pii/S2666827021000773>.
40. Prat A, Pineda E, Adamo B, et al. Clinical implications of the intrinsic molecular subtypes of breast cancer. *The Breast* 2015;24:S26–S35. <https://doi.org/10.1016/j.breast.2015.07.008>. ISSN 0960-9776. URL: <https://www.sciencedirect.com/science/article/pii/S0960977615001460>. 14th St.Gallen International Breast Cancer Conference – Proceedings Book.