



OPEN

## Overcoming the limitations of patch-based learning to detect cancer in whole slide images

Ozan Ciga<sup>1✉</sup>, Tony Xu<sup>5</sup>, Sharon Nofech-Mozes<sup>3,4</sup>, Shawna Noy<sup>2</sup>, Fang-I Lu<sup>3,4</sup> & Anne L. Martel<sup>1,2</sup>

Whole slide images (WSIs) pose unique challenges when training deep learning models. They are very large which makes it necessary to break each image down into smaller patches for analysis, image features have to be extracted at multiple scales in order to capture both detail and context, and extreme class imbalances may exist. Significant progress has been made in the analysis of these images, thanks largely due to the availability of public annotated datasets. We postulate, however, that even if a method scores well on a challenge task, this success may not translate to good performance in a more clinically relevant workflow. Many datasets consist of image patches which may suffer from data curation bias; other datasets are only labelled at the whole slide level and the lack of annotations across an image may mask erroneous local predictions so long as the final decision is correct. In this paper, we outline the differences between patch or slide-level classification versus methods that need to localize or segment cancer accurately across the whole slide, and we experimentally verify that best practices differ in both cases. We apply a binary cancer detection network on post neoadjuvant therapy breast cancer WSIs to find the tumor bed outlining the extent of cancer, a task which requires sensitivity and precision across the whole slide. We extensively study multiple design choices and their effects on the outcome, including architectures and augmentations. We propose a negative data sampling strategy, which drastically reduces the false positive rate (25% of false positives versus 62.5%) and improves each metric pertinent to our problem, with a 53% reduction in the error of tumor extent. Our results indicate classification performances of image patches versus WSIs are inversely related when the same negative data sampling strategy is used. Specifically, injection of negatives into training data for image patch classification degrades the performance, whereas the performance is improved for slide and pixel-level WSI classification tasks. Furthermore, we find applying extensive augmentations helps more in WSI-based tasks compared to patch-level image classification.

Convolutional neural networks (CNNs) are able to extract features from raw images, which has made them attractive in many visual tasks, including medical image analysis. In computational pathology, current approaches mostly focus on isolated patches extracted from WSIs, or tasks that are tailored around a slide or patient-level predictions, as opposed to requiring consistently high precision and sensitivity across a WSI<sup>1</sup>. We argue that applying CNNs on WSIs might not achieve similar performances due to several biases outlined below, in addition to the different characteristics of WSIs and patches extracted from WSIs.

Much of the early work in digital pathology involved the extraction of a small fraction of patches from each WSI, which were then annotated or labelled by pathologists. This approach is vulnerable to a data curation bias which refers to the implicit assumptions made by curators that may affect the task outcome. In patch-level problems, training and validation datasets are generally collected by the same experts or with the same instructions. For example, both training and validation sets contain the same number of unique classes with similar ratios. Data curation is also generally biased towards the positive class since these images are usually collected from patients undergoing treatment. Curation bias is exploited for tuning the network hyperparameters (e.g.,

<sup>1</sup>Department of Medical Biophysics, University of Toronto, Toronto, Canada. <sup>2</sup>Sunnybrook Research Institute, Toronto, Canada. <sup>3</sup>Department of Laboratory Medicine and Molecular Diagnostics, University of Toronto, Toronto, Canada. <sup>4</sup>Division of Anatomic Pathology, Sunnybrook Health Science Center, Toronto, ON, Canada. <sup>5</sup>Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, Canada. <sup>6</sup>Tony Xu: This work was conducted while the author was doing an internship at Sunnybrook Research Institute. ✉email: ozan.ciga@mail.utoronto.ca

the confidence threshold) on validation sets to maximize performance on the held-out test set. While these are appropriate in the context of a challenge, they do not necessarily hold for tasks with WSIs, where it is not possible to determine a universally best threshold or to comprehensively annotate data (e.g., representing all negatives in cancer detection, such as ink, creases, out-of-focus regions or non-cancerous tissue).

Although CNNs can only be applied to individual patches within a WSI, slide level tasks can be performed by aggregating the individual patch level results. It is, however, important to take into account how this aggregation is carried out when evaluating and comparing algorithms. Many tasks that involve WSIs are binary decision problems or allow a number of errors without negatively impacting the measured metric. For instance<sup>2</sup>, aim to determine the severity of breast cancer by measuring the extent of metastasis via a five-class problem. However, if a single macro-metastasis is identified, the exact number of metastases become irrelevant for certain classes by definition. Similarly<sup>3</sup>, aim to determine whether cancer is present or not present in a WSI. On slides labelled as containing cancer, false positive are not penalized because the location of the detected cancer is not important; conversely, in slides labelled as negative, there is no differentiation between one or many false positive detections. In contrast, the number of errors at the patch level will be related directly to task performance in slide level tasks that involve quantification or segmentation.

CNNs usually work best with small images such as 224×224 pixels, whereas a WSI can be 100,000×100,000 pixels. Even a drastic down-sampling (e.g., 16×) will necessitate tiling and working on a portion of the WSI. Modifying the architecture, or simply feeding a larger image to overcome this issue is not always possible due to optimization challenges and graphical memory constraints. While the latter is less prohibitive and may be mitigated by technological advancements, training models with more parameters to accommodate WSIs with larger convolutional kernels are infeasible due to  $O((\text{kernel side length})^2)$  scaling. This upper bounds the performance on WSIs since CNNs heavily rely on context and edges, which may not always be present, especially when the WSI is viewed under high resolution (e.g., the outline of a gland or duct may only be partially available if the images are sampled at high resolution over a limited field of view).

In this paper, we extensively study the differences between WSI analysis and patch level analysis. We use WSIs of breast tissue which has been surgically resected following neoadjuvant therapy (NAT) to explore the differences between the following three tasks:

1. Patch-level classification: determine whether cancer cells are present or not in a curated set of patches from WSIs<sup>4</sup>
2. Slide-level classification: determine whether a WSI contains cancer cells or not without localizing the cancer within the WSI.
3. Slide-level segmentation: Find the convex hull that contains all cancer cells in a WSI (i.e., the tumor bed estimation).

Importantly, a binary classification network is used in each task. For the segmentation task, the WSI is tiled into patches, and each patch is classified as cancer or no cancer prior to computing the convex hull. We observe that the best practices, including the most suitable augmentations and the appropriate network complexity, are different for each task. We find a network achieving > 99% F1 score on patch-level classification fails to identify the tumor bed accurately, with a Dice score of ~ 50%. We then study failure points of the patch-level classifier and propose a novel negative sampling method that improves the task performance beyond random or no sampling. We also observe a large gap between slide-level classification and segmentation.

## Problem definition

To exemplify the differences between patches and WSIs, we tackle the automatic detection of breast cancer. Breast cancer is the second most common invasive cancer in women, where around 12% of women having a lifetime risk of getting breast cancer<sup>5</sup>. New therapies have improved outcomes and it is increasingly important to assess the stage and subtype of cancer more precisely in order to select the most appropriate treatment. While tools that can identify and characterize cancer in a specific region of the tissue specimen selected by a pathologist have value, there is a need for automated tools capable of processing WSIs since cancer may be anywhere within a slide, or even spread across multiple slides.

An evaluation of the residual tumour in surgically resected breast tissue after neoadjuvant therapy for breast cancer is clinically relevant in determining treatment response and estimating prognosis. It has been shown that the extent of the residual tumor bed, the cellularity within the tumor bed, and the presence or absence of lymph node metastases can be used to estimate the five-year recurrence risk<sup>6,7</sup>. While automated analysis of WSIs for breast cancer tasks is common<sup>2,8,9</sup>, there have been few studies that focus on post-NAT images or slide-wide analysis<sup>10</sup>. Recently<sup>3</sup>, conducted a large scale study with 9894 slides from axillary lymph nodes for breast cancer metastasis detection, and achieved 0.989 area under the curve (AUC). 52% of the false negatives were from slides that showed signs of neoadjuvant chemotherapy, indicating the challenges associated with identifying cancer in post neoadjuvant therapy images, even when the dataset size is large. In addition, 8.6% of false negatives were due to isolated single tumor cells, which play an important role in determining the tumor bed outline as they may change the extent significantly.

Due to its non-trivial nature, clinical relevance and sensitivity towards both false positives and negatives, we conduct experiments for the tumor bed estimation task on WSIs of resected breast tissue. The tumor bed is defined as the convex hull that contains all residual cancer cells and the task requires that all residual cancer cells are detected; after therapy this is more difficult as isolated cancer cells may be surrounded by stromal tissue. The task also requires a very low false positive detection rate as the erroneous detection of cancer cells can lead to large errors in the boundaries of the tumor bed.

## Methods

**Data.** The BreastPathQ training dataset<sup>10</sup> consists of 2394 patches (of size  $512 \times 512$ ) extracted from 69 WSIs of Post-NAT-BRCA specimens<sup>11</sup>. BreastPathQ images are given labels according to the percentage of cellularity in each patch. The training set is imbalanced, where 71% of the images are tumor-positive. The corresponding WSIs<sup>4</sup> also include tumor bed outlines for each slide which we use for sampling additional negative patches from tumor-free regions of the WSIs.

We use the testing images from the BreastPathQ challenge for the patch-level classification task. This set includes 1106 images, and is annotated by two independent pathologists. The test set for the slide-level classification and segmentation tasks is composed of a private set of 50 WSIs, evenly and randomly sampled from 10 independent cancer patients in different treatment stages. Three pathologists have independently annotated tumor bed outlines in each WSI and later reviewed each WSI together to create a consensus ground truth; this allows us to assess the inter-expert disagreement rate. In our experiments, we use the consensus ground truth unless otherwise stated. According to the consensus dataset, 30% of the WSIs are tumor-free. For augmentation, architecture and stride experiments, we allocate 40% of the WSIs as the validation set, and use the remaining 60% for testing. Both sets contain roughly the same number of tumor-free slides, split into each set at random. Using this data, we compare pathologists to each other, to the consensus, as well as to our method.

**Preprocessing.** Regions without nuclei are not relevant in identifying tumors. Therefore, we threshold each WSI in HSV color space to remove such regions, with  $0.65 > \text{hue} > 0.5$ ,  $\text{saturation} > 0.1$ , and  $0.9 > \text{value} > 0.5$ . These bounds are experimentally determined using a separate public dataset for breast cancer segmentation<sup>9</sup>. We tile the WSI and pass patches with foreground ratios  $\geq 25\%$  through the model during the evaluation stage.

**Patch-level classification task.** We use the BreastPathQ dataset for patch-level classification. The continuous cellularity percentage values per patch are mapped to two classes, with cellularity percentages above zero labelled as cancer.

**Slide-level tasks.** In addition to the same dataset used for patch-level classification, we incorporate additional negative (no cancer) samples for the slide level tasks (see Section 3.5 for details).

Our method achieves slide-level segmentation (tumor bed estimation) by tiling the WSI into distinct patches and generating a probability value (of being tumor-positive) per patch. One may slide a window in strides smaller than the patch length to improve the classification of each pixel through a voting scheme. For example, given a patch size of 512, a stride of 256 will generate four separate probability values per pixel, representing if the patch contains any tumor. Then, these values may be averaged to obtain more accurate probability values per pixel. Since each value is obtained using a different patch, striding helps to incorporate context better.


The predicted label for the slide-level classification can be obtained from the segmentation output. Specifically, if the tumor bed exists, or there is one or more positive patch on the WSI, the WSI is labeled as tumor-positive.

**Negative sampling.** A single WSI may contain artifacts (e.g., creases, random blurring of regions, ink or staining residues), as well as structures that are considered as background (e.g., histocytes, adenosis, or red blood cells in tumor bed estimation) which may not have been filtered out in the preprocessing step. While all datasets provide negative samples, it is not feasible to label all forms of negatives. Therefore in digital histopathology, it is common for the negative (e.g., healthy or benign tissue in breast cancer detection) class to be underrepresented.


BreastPathQ dataset consists of 69 WSIs, where each WSI is annotated for the tumor bed outline. By definition, patches extracted from outside of the tumor bed region are considered as completely cancer-free. We sample patches from outer extents of each WSI and collect a total of 1.4 million images. Then, we sample 21 thousand images randomly or by clustering. For clustering, we use features (CR<sup>1280</sup>) generated by EfficientNet-B0 model trained on ImageNet<sup>12</sup>. Specifically, we use features corresponding to each image obtained by 2d-average pooling of the final layer of the network prior to the fully-connected classification layer (also called as the pre-activation). We then cluster 1.4 million feature vectors into three thousand clusters using the mini-batch K-means algorithm and select seven instances from each cluster. To find the optimal cluster size, we use the *Elbow method* principle. We run K-means with cluster sizes {1000, 1500, ..., 10000}. 2500 clusters achieve an explained variance of 67.1%, 3000 clusters is 83.3%, and 5000 is 91%. Examples of sampled images from both random and K-means sampling strategies can be viewed in Fig. 1.

**Assessing the prediction accuracy of tumor bed outline.** Tumor beds do not have obvious boundaries, unlike structures such as breast ducts. Therefore, a metric which emphasizes the difference in extent rather than the overlap is desired. As a quantitative measure of the tumor extent<sup>6</sup>, use the longest diagonal line  $d_1$  inside the tumor bed, and the longest line  $d_2$  within the tumor bed that is perpendicular to  $d_1$  (see Fig. 2), and combine them in a single quantity  $d_{prim} = \sqrt{d_1 d_2}$ . This measure was found to be predictive of long-term survival post neoadjuvant therapy, allows for a more refined evaluation of response among tumors with similar sizes with different shapes, and is being adopted by a growing number of clinical trials as a primary outcome measure<sup>13</sup>.

In the following, we demonstrate the advantages of  $d_{prim}$  over Dice in representing the spread of cancer.  $d_{prim}$  is reported in millimeters as opposed to percentages, which preserves the scale information. For instance, a false negative for a very large tumor bed will result in the same Dice score (0%) as of a single cell tumor, whereas it is possible to quantify the error using  $d_{prim}$ . The same holds for nonzero overlap: In Fig. 3a, the overlap is 86% with a  $d_{prim}$  error of 2.4 mm, whereas in Fig. 3b the error in  $d_{prim}$  is 0.5 mm with a 89% Dice, since the tumor bed is much smaller. The convex hull is robust to the deformities within the tumor bed that are insignificant for prognosis,



**Figure 1.** Negative mining with random sampling versus feature based K-means clustering from whole-slide images. Both figures are extracted from actual experiments. (a) Random sampling (b) K-means sampling. In (b), we show the three samples closest to the centres of six randomly selected clusters for illustration.

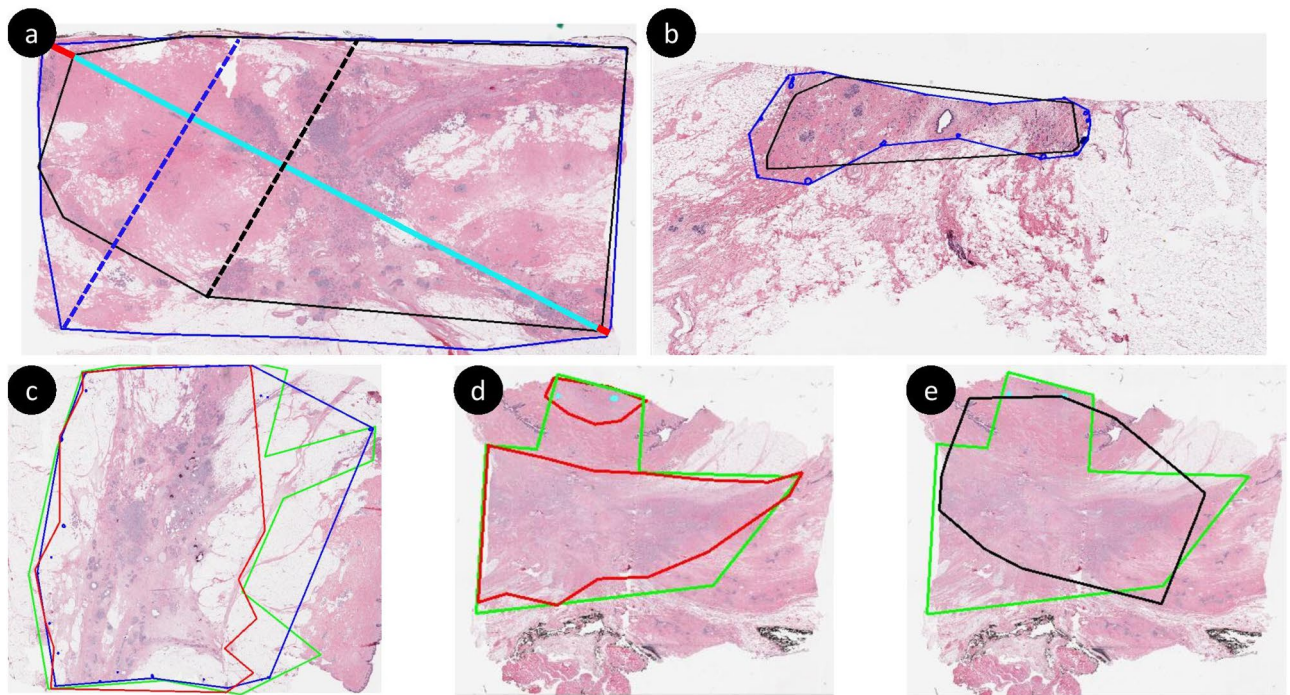


**Figure 2.**  $d_1$  (turquoise) and  $d_2$  (blue) for the tumor bed convex hull.

which alleviates the burden on the annotator by allowing less refined boundaries so long as the distant tumors are identified. For instance, in Fig. 3c, the overlap between red and both green and blue lines is  $> 85\%$ , whereas the difference in  $d_{prim}$  is around 7 mm. Conversely, the overlap for green and blue are 93%, with  $\sim 2$  mm difference. Finally, given a large tumor bed, we observed significant overlap is still possible even if multiple regions are missed. The overlap in Fig. 3d is 95% ( $d_{prim}$  error is 1.1 mm) whereas in Fig. 3e it is 85% ( $d_{prim}$  error is 2.5 mm).

## Experiments

**Experimental setup.** For all experiments, we use the EfficientNet framework<sup>12</sup>, and train the network for predicting if a patch of size  $224 \times 224$  (training patches are of size  $512 \times 512$ , which are randomly cropped to  $448 \times 448$  before resizing each dimension in half) contains cancer cells. EfficientNet is based on AutoML framework and compound scaling, and achieves better performance than standard architectures such as residual networks with a fewer number of trainable parameters. We use Adam optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , learning rate of 0.0001, batch size of 20, and weighted cross-entropy loss function, and train for 250 epochs for each experiment. We split 15% of the training data for validation, and select the model with minimum validation loss in testing. Unless otherwise stated, we use the EfficientNet-B0 model.



**Figure 3.** Comparison between Dice and  $d_{prim}$ . Each annotation is converted to its convex hull to extract  $d_{prim}$  between the compared masks. To avoid crowding, we only show the tumor bed extents for (a), where the longest diagonal difference (shown in red lines) is  $\sim 1$  mm, whereas the perpendicular diagonal difference is  $\sim 2$  mm.

**Metrics.** In binary decision problems, such as determining *any* presence of tumor in WSIs, finding a single positive region suffices for identifying the positive WSIs, and any additional false positives or false negatives for the same WSI do not incur any penalty. In contrast, our task requires us to be consistently precise throughout the WSI in order to estimate the tumor bed accurately. Therefore, we report both the confusion matrix to quantify if our system is able to detect tumor accurately or discard slides without tumor, as well as the error in  $d_{prim}$ , defined as the absolute difference between the prediction and ground truth averaged over the test set, and the overlap (measured by the Dice coefficient) between the ground truth and the prediction. Note that  $d_{prim}$  is especially important in assessing the quality of predictions since, unlike the Dice coefficient, this metric is very sensitive to distant false positives and false negatives towards the WSI's outer edges.

**Tasks.** *Patch-level cancer detection.* We use models trained with minority oversampling weighing (see Table 3), and assess their performance on patch-level BreastPathQ challenge test set with 1106 samples. We also conduct experiments to assess the effects of augmentation, size of the architecture and number of test time augmentations.

*Whole-slide image cancer detection.* After separately annotating the test set, pathologists reviewed the dataset together and discarded four slides due to high disagreement rates. In these slides, pathologists disagreed as to the presence of tumor, annotated non-overlapping regions as the tumor bed, or disagreed on highly contentious regions where the presence of tumor could not be determined by only using H&E stained slides. For the remaining slides, a consensus tumor bed was selected using a combination of expert annotations. Any additional regions were not examined during this step.

**Results.** *Patch image analysis.* We compare three different negative sampling methods in Table 1. The negative instances are sampled from the WSIs as described in Section 3.5. Minority oversampling performed the best for all sampling strategies. We compare multiple augmentation strategies, number of test time augmentations (TTA), and the effect of model complexity on the classification performance in Table 2.

*Whole slide image analysis.* The results of our experiments are presented in Tables 3, 4, 5 and 6. We use the minimum error in  $d_{prim}$  to select the best performing model in our experiments. Note that the average ground truth  $d_{prim}$  is 14.5 millimeters, with a standard deviation of 6.1. We compare the discrepancy between references (i.e., the pathologists) to each other, as well as to the network in Table 7 after the consensus set (with 46 slides) has been established. We also make the same comparison prior to consensus, accounting for all 50 slides in Table 8.

We first determine the best negative sampling method, along with options for weighing different samples depending on their class (Table 3). We compare three different sampling strategies, including the clustering approach defined in Section 3.5, randomly selecting negative instances as well as no additional negative sampling beyond what the training data already has. Negative mining creates an imbalance in favor of the negatives (85%

Negative sampling	Acc (%)	Cfs. mtx
K-means	92.6	$\begin{pmatrix} 173 & 58 \\ 23 & 852 \end{pmatrix}$
Random	96.3	$\begin{pmatrix} 200 & 31 \\ 9 & 866 \end{pmatrix}$
None	98.7	$\begin{pmatrix} 224 & 7 \\ 7 & 868 \end{pmatrix}$

**Table 1.** Binary (tumor vs. no tumor) classification accuracies on BreastPathQ test set with different negative sampling schemes using minority oversampling.

Setting	Acc (%)	Cfs. mtx
No augmentation	93.5	$\begin{pmatrix} 186 & 45 \\ 27 & 848 \end{pmatrix}$
Stain norm	92.4	$\begin{pmatrix} 180 & 51 \\ 32 & 843 \end{pmatrix}$
Color jitter	93.5	$\begin{pmatrix} 205 & 26 \\ 46 & 829 \end{pmatrix}$
Affine trs	94.4	$\begin{pmatrix} 207 & 24 \\ 38 & 837 \end{pmatrix}$
TTA=1	92.9	$\begin{pmatrix} 198 & 33 \\ 45 & 830 \end{pmatrix}$
TTA=2	93.5	$\begin{pmatrix} 186 & 45 \\ 27 & 848 \end{pmatrix}$
TTA=3	93.3	$\begin{pmatrix} 183 & 48 \\ 26 & 849 \end{pmatrix}$
TTA=4	92.9	$\begin{pmatrix} 178 & 53 \\ 25 & 850 \end{pmatrix}$
TTA=16	92.2	$\begin{pmatrix} 167 & 64 \\ 22 & 853 \end{pmatrix}$
B0	93.5	$\begin{pmatrix} 186 & 45 \\ 27 & 848 \end{pmatrix}$
B3	93.1	$\begin{pmatrix} 189 & 42 \\ 34 & 841 \end{pmatrix}$
B5	95.1	$\begin{pmatrix} 201 & 30 \\ 24 & 851 \end{pmatrix}$
B7	93.5	$\begin{pmatrix} 193 & 38 \\ 34 & 841 \end{pmatrix}$

**Table 2.** Binary (tumor vs. no tumor) classification accuracies on BreastPathQ test set under different settings. Unless otherwise stated, test-time augmentations (TTA) are set to two.

of the training data). Therefore, we use multiple balancing strategies, including multiplying the incurred loss with a coefficient that is directly ( $\propto$ ) or inversely ( $\frac{1}{\propto}$ ) proportional to its corresponding class. We also conduct experiments using no balancing (=) and balancing by oversampling. For our training data, each negative sample is passed through the network two times, whereas each positive sample is passed 11 times per epoch during the training.

We also compare the effect of different types of image augmentation on slide-level classification and segmentation tasks.

(Table 4). We compare randomly cropping 448×448 rectangular boxes from the training instances versus no cropping both during training and evaluation. We modify the color properties of each instance by applying color jittering, which amounts to randomly changing brightness, contrast, saturation and hue values by 5% from their original values, as well as by stain normalization using a reference staining matrix<sup>14</sup>. We experiment applying affine image transformations by scaling the cropped patch down or up by 10%, and shearing by 60° in both x and y axes.

We compare four different models with varying widths and depths to understand the effect of model complexity on slide-level task outcomes (Table 5). We also compare multiple sliding window approaches for finding the

Negative sampling	Weighing	Dice (% overlap)	Error in $\sqrt{d_1 d_2}$	Confusion matrix
K-means	$\alpha$	59	2.58	$\begin{pmatrix} 5 & 3 \\ 0 & 12 \end{pmatrix}$
	$\frac{1}{\alpha}$	58	3.29	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$
	=	58	2.56	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$
	Minority oversam-pling	75	1.69	$\begin{pmatrix} 6 & 2 \\ 0 & 12 \end{pmatrix}$
Random	$\alpha$	64	1.90	$\begin{pmatrix} 4 & 4 \\ 1 & 11 \end{pmatrix}$
	$\frac{1}{\alpha}$	61	3.55	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$
	=	58	2.21	$\begin{pmatrix} 2 & 6 \\ 0 & 12 \end{pmatrix}$
	Minority oversam-pling	57	3.18	$\begin{pmatrix} 3 & 5 \\ 1 & 11 \end{pmatrix}$
None	$\alpha$	30	15.37	$\begin{pmatrix} 0 & 8 \\ 0 & 12 \end{pmatrix}$
	$\frac{1}{\alpha}$	30	14.58	$\begin{pmatrix} 0 & 8 \\ 0 & 12 \end{pmatrix}$
	=	26	16.33	$\begin{pmatrix} 0 & 8 \\ 0 & 12 \end{pmatrix}$
	Minority oversam-pling	30	15.25	$\begin{pmatrix} 0 & 8 \\ 0 & 12 \end{pmatrix}$

**Table 3.** The effect of sampling negative instances randomly versus a clustering approach. Under each sampling strategy, we compare various methods to handle class imbalance.  $\alpha$  indicates using class weights in cross-entropy loss directly proportional to class distributions, = uses weights 1 for all classes. Each setting uses the EfficientNet-B0 framework, 256 as the tile stride, and only random cropping as data augmentation.

Random crop	Stain norm	Color jitter	Affi-ne trs	Dice (% over-lap)	Error in $\sqrt{d_1 d_2}$	Cfs. mtx
				63	3.63	$\begin{pmatrix} 4 & 4 \\ 0 & 12 \end{pmatrix}$
✓	✓			70	2.01	$\begin{pmatrix} 4 & 4 \\ 0 & 12 \end{pmatrix}$
✓		✓		73	1.67	$\begin{pmatrix} 5 & 3 \\ 0 & 12 \end{pmatrix}$
✓			✓	67	1.56	$\begin{pmatrix} 4 & 4 \\ 1 & 11 \end{pmatrix}$
✓	✓	✓	✓	75	1.63	$\begin{pmatrix} 6 & 2 \\ 0 & 12 \end{pmatrix}$

**Table 4.** The effect of data augmentation on slide-level tasks using the K-means negative sampling strategy, 256 as the tile stride, and EfficientNet-B0 architecture.

most suitable tile stride (Table 6). Finally, we compare our method to each of the three pathologists, along with comparing pathologists to each other in tables 7 and 8.

## Discussion and conclusion

8% of the test set had to be discarded as the three pathologists were unable to reach consensus, and one slide was relabeled as non-tumor; this demonstrates the challenge associated with estimating the tumor bed, even for experts. In clinical practice, such challenging cases can be resolved by using an additional stain (e.g., immunostaining). Interestingly, two false positives and two false negative WSIs that our system initially detected were discarded in the consensus set, indicating errors made by the network are not random, and the failure points can be attributed to the uncertainty in the slide.

We find using randomly sampled negative instances to be particularly ineffective in identifying tumor-free slides. In effect, randomly sampling negative regions undersamples rare regions that exist in WSIs. In contrast,

Architecture (# params)	Dice (% overlap)	Error in $\sqrt{d_1 d_2}$	Cfs. mtx
B0 (5 M)	75	1.69	$\begin{pmatrix} 6 & 2 \\ 0 & 12 \end{pmatrix}$
B3 (12 M)	71	1.52	$\begin{pmatrix} 6 & 2 \\ 0 & 12 \end{pmatrix}$
B5 (30 M)	64	1.79	$\begin{pmatrix} 4 & 4 \\ 0 & 12 \end{pmatrix}$
B7 (58 M)	55	2.93	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$

**Table 5.** The effect of architecture size on slide-level tasks, using the K-means negative sampling, 256 as the tile stride, and only random cropping as augmentation.

Stride	Dice (% overlap)	Error in $\sqrt{d_1 d_2}$	Cfs. mtx
512	48	9.38	$\begin{pmatrix} 1 & 7 \\ 0 & 12 \end{pmatrix}$
256	71	1.52	$\begin{pmatrix} 6 & 2 \\ 0 & 12 \end{pmatrix}$
128	62	2.45	$\begin{pmatrix} 5 & 3 \\ 0 & 12 \end{pmatrix}$
64	66	1.65	$\begin{pmatrix} 4 & 4 \\ 0 & 12 \end{pmatrix}$
32	64	1.63	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$
16	61	1.66	$\begin{pmatrix} 3 & 5 \\ 0 & 12 \end{pmatrix}$

**Table 6.** The effect of sliding window tile stride on slide-level tasks using the EfficientNet-B3 architecture. In all experiments, we use K-means negative sampling and only random cropping as data augmentation.

clustering negative samples based on a salient feature representation allows us to group visually similar instances. Therefore, clustering lets us sample the negative space more evenly, which improves precision. Furthermore, relying only on annotator labeled negative data (*None* setting) is incapable of identifying true negatives and has a much lower performance compared to either K-means or random sampling. Upon examination of the predictions, we find that the network is incapable of identifying structures such as red blood cells or scanner imperfections (e.g., shadow effects), which are not present in the original BreastPathQ dataset, however negative sampling from Post-NAT BRCA WSIs provide us with this information. While it might not always be possible to obtain a boundary beyond which there is no cancer present, many datasets contain negative WSIs that can be sampled in a similar fashion.

After incorporating the negative samples, our dataset is highly imbalanced in favor of negative samples (85%). A common strategy in handling class imbalance is to increase the cost of mislabeling the minority class by using weighted cross-entropy.

where the weights per class are inversely proportional to the number of training samples with that label. We found that this strategy artificially skews the network to predict more samples from the minority class, which increases the number of false negatives. A counter-intuitive approach, where we use directly proportional weights, performs better, and not accounting for the imbalance (equal weights) gives the best performance out of these three approaches. Finally, we observe that oversampling the minority class to rebalance class distributions gives the best overall performance. Therefore, we use K-means negative samples with minority oversampling in slide-level experiments presented from Table 5 onward. In contrast, the results for the patch-level classification task shown in Table 1 suggest that any form of negative sampling adversely impacts prediction accuracy. Injecting information that is not present in the test set has a negative effect since the network has learned a more complex decision boundary that is not utilized for testing. Interestingly, random sampling is not as detrimental, possibly because most of the negative training instances look significantly different from positives.

We find that randomly cropping the instances improve results in both slide-level segmentation and classification tasks. We find that light color jittering performs better than stain normalization. Affine transformations improve results for the WSI task. In patch level analysis, we observe that augmentations are less significant, and that more than two test time augmentations (including random cropping by  $448 \times 448$ , flips and rotations) negatively impact the prediction accuracy. This is in contrast to WSI level predictions, where a smaller tile stride (corresponding to test-time augmentations in patches) is usually correlated with better performance.



Reference	Target	Dice (% over-lap)	Error in $\sqrt{d_1 d_2}$	Cfs. mtx
P#1	CNN	64	2.12	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
	P#2	74	2.20	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
	P#3	79	1.75	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
P#2	CNN	66	3.06	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
	P#3	81	1.63	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
P#3	CNN	67	1.86	$\begin{pmatrix} 6 & 0 \\ 0 & 20 \end{pmatrix}$
Consensus	CNN	66	2.33	$\begin{pmatrix} 6 & 1 \\ 0 & 19 \end{pmatrix}$
	P#1	73	1.65	$\begin{pmatrix} 6 & 1 \\ 0 & 19 \end{pmatrix}$
	P#2	88	0.79	$\begin{pmatrix} 6 & 1 \\ 0 & 19 \end{pmatrix}$
	P#3	79	1.48	$\begin{pmatrix} 6 & 1 \\ 0 & 19 \end{pmatrix}$

**Table 7.** Comparison between references (expert pathologists), and our method on slide-level tasks using the EfficientNet-B3 architecture, including the consensus ground truth (with 8% of the dataset removed). P#\* indicates the expert ID, and CNN is the trained network to segment the tumor bed automatically. In all experiments, we use K-means negative sampling, 256 as the tile stride, and color jittering and affine transformations as data augmentation.

Reference	Target	Dice (% over-lap)	Error in $\sqrt{d_1 d_2}$	Cfs. mtx
P#1	CNN	56	2.61	$\begin{pmatrix} 6 & 1 \\ 1 & 22 \end{pmatrix}$
	P#2	64	3.42	$\begin{pmatrix} 6 & 1 \\ 0 & 23 \end{pmatrix}$
	P#3	68	2.53	$\begin{pmatrix} 6 & 1 \\ 0 & 23 \end{pmatrix}$
P#2	CNN	60	3.38	$\begin{pmatrix} 6 & 0 \\ 2 & 22 \end{pmatrix}$
	P#3	77	1.96	$\begin{pmatrix} 6 & 0 \\ 0 & 24 \end{pmatrix}$
P#3	CNN	61	2.15	$\begin{pmatrix} 6 & 0 \\ 1 & 23 \end{pmatrix}$

**Table 8.** Comparison between references (expert pathologists), and our method on slide-level tasks using the EfficientNet-B3 architecture, prior to including consensus ground truth. P#\* indicates the expert ID, and CNN is the trained network to segment the tumor bed automatically. In all experiments, we use K-means negative sampling, 256 as the tile stride, and color jittering and affine transformations as data augmentation.

We find that EfficientNet-B3, which is comparable to ResNet34 and ResNet50 in terms of model complexity, provides the best bias and variance trade-off for our data and slide-level tasks. For the patch-level task, we find larger architectures to be more suitable, where the best performing network (B5) is 2.5× the size of best performing model (B3) for tumor bed estimation.

We find that stride lengths smaller than the patch length primarily help identify negative regions, and marginally improve the tumor bed outlines. We also find that stride lengths drastically smaller than the patch length do not significantly improve the results, and are computationally prohibitive (e.g., a stride length of 16 processes each pixel 1024 times, and each WSI in this setting can have more than a million patches to be evaluated).

In conclusion, we find that best practices in patches versus WSI analysis vary significantly, despite employing the same networks with the same training data. This discrepancy also highlights that the advancements in either setting cannot be directly applied to the other, and independent research is required to further both fields.

Received: 9 December 2020; Accepted: 6 April 2021

Published online: 26 April 2021

## References

1. Srinidhi, C. L., Ciga, O. & Martel, A. L. Deep neural network models for computational histopathology: A survey. *arXiv preprint arXiv:1912.12378* (2019).
2. Bandi, P. *et al.* From detection of individual metastases to classification of lymph node status at the patient level: the CAMELYON17 challenge. *IEEE Trans. Med. Imaging* **38**, 550–560 (2018).
3. Campanella, G. *et al.* Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Med.* **25**, 1. <https://doi.org/10.1038/s41591-019-0508-1> (2019).
4. Martel, A. L., Nofech-Mozes, S., Salama, S., Akbar, S. & Peikari, M. Assessment of residual breast cancer cellularity after neoadjuvant chemotherapy using digital pathology. <https://doi.org/10.7937/TCIA.2019.4YIBTJNO> (2019).
5. DeSantis, C., Ma, J., Bryan, L. & Jemal, A. Breast cancer statistics, 2013. *CA: a cancer journal for clinicians* **64**, 52–62 (2014).
6. Symmans, W. F. *et al.* Measurement of residual breast cancer burden to predict survival after neoadjuvant chemotherapy. *J. Clin. Oncol.* **25**, 4414–4422 (2007).
7. Symmans, W. F. *et al.* Long-term prognostic risk after neoadjuvant chemotherapy associated with residual cancer burden and breast cancer subtype. *J. Clin. Oncol.* **35**, 1049 (2017).
8. Ehteshami Bejnordi, B. *et al.* Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. *JAMA* **318**, 2199–2210 (2017).
9. Aresta, G. *et al.* Bach: Grand challenge on breast cancer histology images. *Med. image analysis* (2019).
10. Akbar, S. *et al.* Automated and manual quantification of tumour cellularity in digital slides for tumour burden assessment. *Sci. Reports* **9**, 14099 (2019).
11. Peikari, M., Salama, S., Nofech-Mozes, S. & Martel, A. L. Automatic cellularity assessment from post-treated breast surgical specimens. *Cytom. Part A* **91**, 1078–1087 (2017).
12. Tan, M. & Le, Q. EfficientNet: Rethinking model scaling for convolutional neural networks. In Chaudhuri, K. & Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning*, vol. 97 of *Proceedings of Machine Learning Research*, 6105–6114 (PMLR, Long Beach, California, USA, 2019).
13. Peintinger, F. *et al.* Reproducibility of residual cancer burden for prognostic assessment of breast cancer after neoadjuvant chemotherapy. *Mod. Pathol.* **28**, 913 (2015).
14. Macenko, M. *et al.* A method for normalizing histology slides for quantitative analysis. In *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 1107–1110 (IEEE, 2009).

## Acknowledgements

This research was enabled in part by support provided by Compute Canada ([www.computecanada.ca](http://www.computecanada.ca)). This work was funded by Canadian Cancer Society (grant #705772) and NSERC RGPIN-2016-06283.

## Author contributions

O.C. contributed to conceptualization, methodology, software, data curation, writing - original draft, visualization, and validation. T.X. contributed to methodology, software. S.N.-M., S.N., and F.-I.L. contributed to problem definition, annotation of data sets, analysis of results. A.L.M. contributed to supervision, investigation, writing - review and editing. All authors contributed to the preparation and review of manuscript.

## Competing interests

ALM is co-founder and CSO of Pathcore. Other authors have no conflict to declare.

## Additional information

**Correspondence** and requests for materials should be addressed to O.C.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021