



HAL
open science

Natural vs Balanced Distribution in Deep Learning on Whole Slide Images for Cancer Detection

Ismat Ara Reshma, Sylvain Cussat-Blanc, Radu Tudor Ionescu, Hervé Luga,
Josiane Mothe

► To cite this version:

Ismat Ara Reshma, Sylvain Cussat-Blanc, Radu Tudor Ionescu, Hervé Luga, Josiane Mothe. Natural vs Balanced Distribution in Deep Learning on Whole Slide Images for Cancer Detection. 36th ACM/SIGAPP Symposium on Applied Computing (SAC 2021), Association for Computing Machinery (ACM) - Special Interest Group on Applied Computing (SIGAPP), Mar 2021, Virtual Event, South Korea. <hal-03057209>

HAL Id: hal-03057209

<https://hal.science/hal-03057209v1>

Submitted on 11 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Natural vs Balanced Distribution in Deep Learning on Whole Slide Images for Cancer Detection

Ismat Ara Reshma*
IRIT, UMR5505 CNRS, Univ. de
Toulouse
Toulouse, France
ismat-ara.reshma@irit.fr

Sylvain Cussat-Blanc
IRIT, ANITI, UMR5505 CNRS, Univ.
de Toulouse
Toulouse, France
sylvain.cussat-blanc@irit.fr

Radu Tudor Ionescu
Univ. of Bucharest
Bucharest, Romania
raducu.ionescu@gmail.com

Hervé Luga
IRIT, UMR5505 CNRS, Univ. de
Toulouse
Toulouse, France
herve.luga@irit.fr

Josiane Mothe
IRIT, UMR5505 CNRS, Univ. de
Toulouse
Toulouse, France
josiane.mothe@irit.fr

ABSTRACT

The class distribution of data is one of the factors that regulates the performance of machine learning models. However, investigations on the impact of different distributions available in the literature are very few, sometimes absent for domain-specific tasks. In this paper, we analyze the impact of natural and balanced distributions of the training set in deep learning (DL) models applied on histological images, also known as whole slide images (WSIs). WSIs are considered as the gold standard for cancer diagnosis. In recent years, researchers have turned their attention to DL models to automate and accelerate the diagnosis process. In the training of such DL models, filtering out the non-regions-of-interest from the WSIs and adopting an artificial distribution—usually a balanced distribution—is a common trend. In our analysis, we show that keeping the WSIs data in their usual distribution—which we call natural distribution—for DL training is better than the artificially obtained balanced distribution. We conduct an empirical comparative study with 10 random folds for each distribution, comparing the resulting average performance levels in terms of five different evaluation metrics. Experimental results show the effectiveness of the natural distribution over the balanced one across all the evaluation metrics.

CCS CONCEPTS

• **Computing methodologies** → *Supervised learning; Image processing; Image segmentation*; • **Applied computing** → *Health informatics*.

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '21, March 22–26, 2021, Virtual Event, Republic of Korea

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8104-8/21/03...\$15.00

<https://doi.org/10.1145/3412841.3441884>

KEYWORDS

Deep learning, class-biased training, class distribution analysis, whole slide images, cancer detection.

ACM Reference Format:

Ismat Ara Reshma, Sylvain Cussat-Blanc, Radu Tudor Ionescu, Hervé Luga, and Josiane Mothe. 2021. Natural vs Balanced Distribution in Deep Learning on Whole Slide Images for Cancer Detection. In *The 36th ACM/SIGAPP Symposium on Applied Computing (SAC '21), March 22–26, 2021, Virtual Event, Republic of Korea*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3412841.3441884>

1 INTRODUCTION

Histological images or whole slide images (WSIs) are the digital conversions of conventional histological glass slides containing tissue samples [7, 17]. During the last two decades, the analysis of WSIs has become common practice for the clinical diagnosis of cancers by pathologists. However, the process of diagnosis has some well-known limitations, such as being very time consuming, laborious, and subjective to pathologists' expertise and fatigue. To address these limitations, researchers are investigating automatic cancer detection using machine learning models. Specifically, the huge success of deep learning models, such as convolutional neural networks (CNNs), in visual recognition [15, 18, 25] motivated researchers to explore their use in cancer detection in WSIs [1, 22, 32].

However, the success of CNN models depends on different hyper-parameters settings. One of the most important hyper-parameters is the class distribution of the training data. In machine learning, the imbalanced data distribution has been shown to lead to inferior models, compared to the balanced distribution [3, 12]. This is likely to be the reason why the balanced distribution is the most adopted in deep learning state-of-the-art methods [1, 8, 22], although Prati et al. [24], for example, showed it is not optimal in all cases. Unfortunately, very few comparative studies exist in the literature [24, 34, 35], and these are mainly conducted on toy data sets. However, real data sets can be far different and more complex than the toy ones. Thus, there is no evidence that the balanced distribution will work on cancer WSIs. Consequently, being a special kind of image type, domain-specific studies are required for WSIs data.

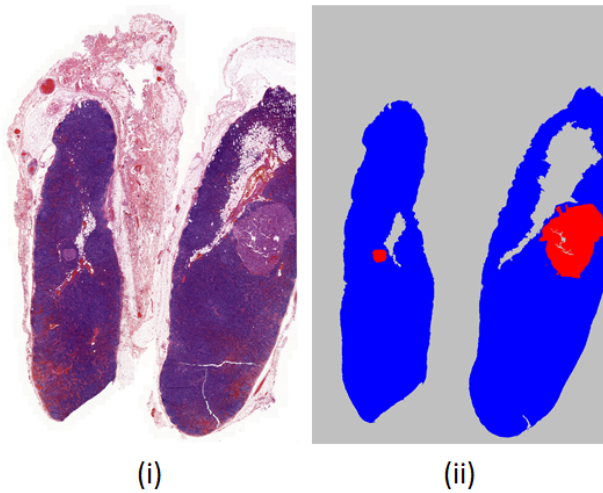


Figure 1: A WSI of a metastatic lymph node (i) with its annotated mask (ii), where each color represents a specific part. The red regions represent metastasis (cancer) in the lymph node (ROI), the blue regions represent lymph node without metastasis (also ROI), while the gray regions represent background and other histological structures (non-ROI). Best viewed in color.

In a whole slide image, some particular regions are considered as regions-of-interest (ROIs), drawing the interest of pathologists to check for abnormalities, and are in turn divided into several classes, depending on the use purpose. In contrast, the remaining regions are considered as non-ROIs, which are mainly background. For example, in a WSI of a metastatic lymph node, the regions comprising the lymph node tissue are ROIs, while the remaining ones are non-ROIs. In this case, ROIs are further divided into cancer, unhealthy versus non-cancer, healthy parts. Non-ROI parts are background as well as some other histological structures such as blood, adipose, or fibrous tissue. Figure 1 presents an example WSI of metastatic lymph node with all the mentioned parts.

One of the important features in WSI data sets is that the pixel distribution is usually biased towards the non-ROI class, which we call the *natural* distribution, that is to say, the distribution provided in the data set. Typically, a WSI contains 70 to 80% of non-ROI pixels on average [20, 32]. It is common practice to remove the non-ROIs from the computation pipeline by applying some pixel filtering algorithms [20, 32]. However, to the best of our knowledge, the usability of non-ROIs during training has not been investigated so far.

Regarding the ROI classes, the natural distribution is either balanced or imbalanced depending on the included WSIs in the data set. It is common practice to artificially balance the ROI classes when they are not in the natural distribution of the data set. In the case of WSI, related work does not analyze the impact of these distributions on the results.

While related work concludes that it is usually more appropriate to balance the examples in each class in the training data set, we believe the cancer WSI case can be very different from the used

toy data sets. Indeed, cancer WSI data have very specific inter-class similarities and intra-class differences. Wang et al. [33] considers it can lead to false predictions. Moreover, according to our own manual observations, ROIs and non-ROIs contain some common histological structures, which we think could be useful in resolving confusion regarding class variability and similarity. Thus filtering out non-ROI may not be the best practice in such cases.

In this paper, we present an empirical comparative study between the popular artificially-balanced and imbalanced distribution (e.g. original distribution of real data) of WSI data. Through this study, we also found useful insights regarding the usefulness of non-ROIs during training. We believe that the results of our study are specifically important when building new data sets, in order to know which class examples are worth annotating.

The rest of this paper is structured as follows. We first present the related work in Section 2. We describe our methodological framework along with the data set, hyper-parameter settings, and evaluation metrics in Section 3, while the results are presented and discussed in Section 4. In Section 5, we draw some conclusions and mention future directions.

2 RELATED WORK

Most of the real-world data is skewed towards certain classes [27]. This skewness is usually considered as a problem that makes the classification task challenging for a classifier. Many studies have been carried out to measure the impact of artificially-balanced training sets, some suggesting straightforward data-level modifications [3, 10, 16, 19], such as oversampling, undersampling, or using synthetic data, and others suggesting algorithm-level modifications [12, 30], such as applying a weighted cost function. There are even pieces of work that tried to combine data-level and algorithm-level changes [4, 36, 37]. Prati et al. [24] listed some of the most popular data and model balancing methods, focusing on classical machine learning problems, yet, only a few studies discuss the deep learning perspective [2, 11]. Buda et al. [2] compared different methods to address the class-biased problem. According to the authors, imbalanced data has an adversarial effect on the classification accuracy of CNNs, as for classical machine learning techniques. Johnson and Khoshgoftaar [11] conducted a detailed survey on recent techniques to deal with the imbalanced data problem in deep learning, concluding that no method has enough evidence of superiority in dealing with class imbalance using deep learning.

All the aforementioned articles consider the imbalanced data as inauspicious for learning, suggesting methods to make the class distribution of the training data balanced, although the generalization of such methods to specific domains is yet to be proven. In particular, comparative studies on different distributions are required to make a conclusive decision on each data category. However, such studies are few. Weiss and Provost [34, 35] showed that neither the naturally-occurring class distribution nor the balanced distribution is the best for learning, and often, substantially better performance can be obtained by using a different class distribution. In their analysis, they employed 26 data sets from the UCI repository [6]. According to the analysis of Prati et al. [24] on 20 data sets from UCI and a few other private data sets, the most recommended distribution for seven different learning algorithms, including neural

networks, is the balanced distribution. They conclude that only Support Vector Machines are less affected by the class imbalance. However, the study was done on toy data sets. Thus, there is no evidence that the commonly-prescribed balanced distribution is a generalizable solution for more complex real data sets. Consequently, a separate dedicated analysis is required for every special kind of data, like WSIs.

To our knowledge, there is no such comparative study for the cancer detection task from WSIs, while this is beneficial to be known for an optimal outcome of a cancer detection system. Moreover, in state-of-the-art cancer detection systems for WSIs [21–23, 31, 32], the common practice is to filter out non-ROIs and to create a training set with an artificially-balanced or a slightly-skewed distribution to the negative ROI class. Thus, the original data distribution and the use of non-ROIs are yet to be explored.

3 METHODOLOGICAL FRAMEWORK

3.1 Different Class Distributions

To evaluate the importance of the class distribution in the training set, we run a series of experiments that use two different distributions. Specifically, we consider the *natural distribution*, which refers to the initial distribution of the data set, i.e., the distribution a WSI data set naturally has or as provided by stakeholders or pathologists. On the other hand, we consider the *balanced distribution*, which is artificially generated such that classes become uniformly distributed.

WSIs are a special kind of image. They are produced by a whole slide scanner which is a microscope under robotic and computer control, equipped with highly specialized cameras containing advanced optical sensors [7]. Unlike a usual image, a WSI contains information of multiple levels (usually, 8 to 10) of resolution in the same file that allows the extraction of images with low to high magnification levels. The extracted WSIs are usually very large in size, considering the level of extraction. For example, at 40× magnification or level 0, a WSI could be several gigapixels in size. It is impractical to use gigapixels WSIs as direct input to CNNs. To tackle the memory issue, the usual practice is to consider patches rather than whole WSIs for CNN training. We also follow this principle in our study. Thus, we extract overlapping patches and corresponding ground-truth annotation masks from the annotated training data of metastatic lymph node WSIs.

The extracted patches are categorized into different categories to facilitate the creation of different class distributions for the analysis. These categories of patches are defined according to the pixel classes of the patches. In particular, patches that contain more than 99.99% non-ROI pixels are labeled as the *other* (\circ) category. The remaining patches belong to one of the following categories based on the presence of corresponding class pixels: *cancer* (\mathbb{C}), *lymph node* ($\neg\mathbb{C}$), and *mixed* ($\mathbb{C}\&\neg\mathbb{C}$) with an optional presence of class \circ pixels. Among all the patch categories, we consider the category $\mathbb{C}\&\neg\mathbb{C}$ as a multiclass patch category (since it contains two ROI classes at the same time).

We generated the two distributions of classes in the training set by selecting patches from these different categories, as shown in Table 1. To train a model with a balanced distribution, we design experiment E.a by considering the same number of patches from

Table 1: Experiments: balanced distribution (E.a) and natural over-representation of class \circ (E.b) with a total of 9 units of patches in the training set of each experiment.

Experiment ID	Distribution	Patch ratio ($\circ : \mathbb{C} : \neg\mathbb{C}$)
E.a	Balanced	3 : 3 : 3
E.b	Natural	7 : 1 : 1

each of the three classes. On the other hand, to train a model with the natural distribution, we design E.b, where, the training set is highly biased (7 times) toward the class \circ similar to the usual natural distribution in the WSI. In other words, in the training set of E.b, almost 78% of examples are from class \circ . We keep the total number of patches equal in both experiment settings for the sake of fair comparison. Hence, it is not possible to utilize all the extracted patches at the same time in an experiment. Rather, a subset from each category is selected randomly by maintaining the class ratio of a distribution in a particular experiment.

3.2 Network and Hyper-Parameter Setting

We selected U-net [26] as our CNN architecture because it has been proven to be effective even when using very few training images and its training time is manageable. We implemented the U-net architecture using Keras [5] with the TensorFlow backend.

In all the experiments, 80% of the training data samples are used to train the model, the remaining 20% being kept for validation. The validation set is selected randomly from the training set. The models were trained from scratch, i.e., without using transfer learning. We evaluated the trained model on the designated test data set, which was completely unseen during training.

After preliminary empirical evaluation, we set the number of epochs to 35 and the mini-batch size to 5. We opted for the categorical cross-entropy as loss function, similar to the original U-net. To optimize the objective function, we used the Adam optimizer [13]. The initial weights were drawn randomly from the zero-mean Gaussian distribution, as recommended by Krizhevsky et al. [15]. We used a standard deviation of 0.05, which is the default setting in Keras.

We ran the same experiments for different learning rates (10^{-5} , 10^{-4}) and different seeds, observing consistent results across the various configurations. Hence, in Section 4, we report the results using one setting only.

3.3 Training / Testing Principle

In the training step, we generate the natural and the balanced class distributions in the training set by selecting patches from the different categories as defined in the previous sub-section and considering the distribution in Table 1.

The random patch selection is followed by the shuffling of the whole training set to prevent all patches in a mini-batch being from the same category. This makes the convergence faster during training and provides better accuracy [14]. Once the training set with a particular class distribution is created, the CNN model is trained on top of it. Since the CNN model is influenced by the random

Table 2: The number of patches that belong to each category in the training set of the MLNTO data set, when WSIs are downsampled by a factor of $8\times$ and the stride is $d/2$ (i.e., 192).

Patch category	#patches
Other (O)	90,374
Cancer (C)	15,328
Lymph node (\neg C)	17,274
Mixed (C& \neg C)	4,922

weight initialization and the stochastic training process, training the model once might not be enough to validate our hypothesis. Thus, we perform 10 runs for each class distribution, averaging the corresponding results. For two experiments presented in Table 1, we have a total of $2 \times 10 = 20$ runs for a particular parameter setting.

During inference, the trained model is employed to predict the annotation mask of the unseen test WSIs. For this purpose, equally-sized overlapping patches from the test WSIs are considered for prediction. We empirically set half of the pixels of a test patch in the middle as a central region (CR) and the remaining part as a border for each patch. Although the patches are overlapping, the CRs themselves are non-overlapping in a test WSI. The trained model predicts the pixel-level confidence score for each class in the corresponding patch. We annotate a pixel with the class corresponding to the highest confidence score. However, the predicted score for the CR of a patch is taken into account during evaluation.

3.4 Data Set

We utilized the Metastatic Lymph Node data set, named MLNTO, from the University Cancer Institute Toulouse Oncopole. MLNTO is a private data set containing metastases (cancer that spread from its origin to other parts of the body) of 16 primary cancer types and organs, for example, melanoma, adenocarcinoma and squamous cell carcinoma of various anatomical sites.

The data set contains WSIs of lymph nodes stained with hematoxylin and eosin (H&E). Two expert pathologists have provided the ground-truth segmentation masks for all of the WSIs. The three classes as mentioned previously were considered during the annotation (see Figure 1): *metastasis/cancer* (C), *lymph node* (\neg C), and so-called *other* (O), the latter being either background or histological structures, such as adipose or fibrous tissue. The ground-truth masks are given for WSIs that are downsampled by a factor of 8 from the highest resolution, i.e., at level 3 resolution.

We first extract patches from the WSIs. Although random patch sampling is usually applied for patch extraction [21, 22, 32], we extract overlapping patches by maintaining a fixed stride, since we empirically found that using a regular grid performs better than random sampling. Furthermore, there is a lack of sufficient contextual information at the edges of a patch, and CNNs, like U-net, mostly focus on the center region of an input image. Hence, it is preferable to learn from overlapping patches to capture all regions of a WSI.

A total of 127,898 overlapping patches of dimension ($d \times d = 384 \times 384$ pixels) have been extracted with a stride of $d/2$ from the training set. The value of d has been empirically determined for

Table 3: A pixel statistics reporting the average number (in millions and percentage) of pixels of each class per WSI in the training and test sets of MLNTO data set. Considered WSIs are downsampled by a factor of from the original resolution.

		C	\neg C	O
Train	Mean in millions	15.2	14.6	107.4
	Mean in %	11.1	10.6	78.3
Test	Mean in millions	20.0	9.8	114.8
	Mean in %	13.8	6.8	79.4

the segmentation task. The number of patches that fall into each category are given in Table 2.

In our study, we exclude the patches from the mixed (C& \neg C) categories, since these might be problematic in loss calculation and gradient optimization for CNNs, according to Halicek et al. [9]. Among the remaining categories, C has the minimum number of patches (15k), thus, for a balanced distribution, we consider 15k patches from each category in E.a. Likewise, in E.b, to respect the natural data imbalance without increasing the total number of patches, we considered 35k patches from the category O and 5k from each of the other two categories, C and \neg C. In summary, the training data for each experiment consists of a total of 45k patches.

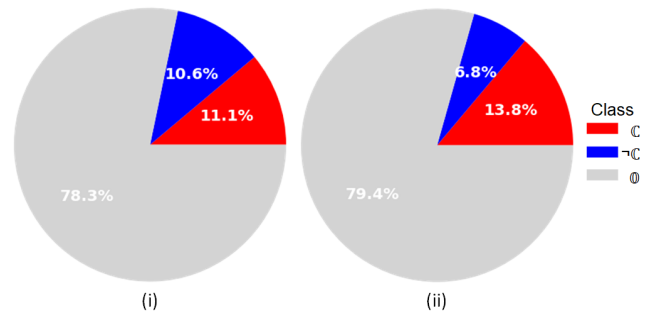


Figure 2: Class distributions of the pixels in the training (i) and test set (ii) of the MLNTO data set.

Table 3 and Figure 2 show the class distributions of the pixels in the training and test data sets. The statistics show that the natural data distribution is imbalanced with an over-representation of the class O in comparison to classes C and \neg C. In the training set, the percentage of class O, C, and \neg C in a WSI are on average 78.3%, 11.1%, and 10.6%, respectively, while in the test set they are 79.4%, 13.8%, and 6.8%, respectively. The imbalanced nature of the data convinced us to better understand the impact of balanced versus imbalanced distribution of the classes on the trained model.

3.5 Evaluation Metrics

For a clinical system, false negative (FN) and false positive (FP) values are both important. Therefore, we considered evaluation metrics that capture both features.

In particular, we utilized the following well-known evaluation metrics to evaluate our trained models: accuracy (Acc), precision

(P), a.k.a positive predictive value (PPV), recall (R), a.k.a sensitivity or true positive rate (TPR), F1-measure, receiver operating characteristic (ROC) curve, and precision-recall (PR) curve. Here, the former four metrics were computed based on the argmax decision (based on a single threshold or cutoff point) of the model’s prediction scores, while the latter two metrics were computed based on all possible threshold values. The precision and recall are responsive to the FP and FN, respectively, while the PR curve expresses the trade-off between precision and recall. Another widely-used metric, the ROC curve, plots the recall/TPR at every false positive rate (FPR) point. Here, FPR is equivalent to 1-specificity.

Furthermore, we applied F1-measure, accuracy, and area under curve (AUC) to summarize the performance of the models. Among all the metrics, the precision, F1-measure, and PR curve can capture the poor performance of the classifier for the imbalanced test set [28].

All aforementioned metrics were applied in our patch-level evaluation since it is impractical to consider each pixel in a WSI. In this regard, we employed some post-processing on the predicted scores to compute the curve-based evaluation metrics. More precisely, we took the maximum of the predicted scores of all CR pixels in a test patch as the predicted score of that patch for a particular class, in this case, for class C. To calculate the mean curve of the 10 runs, we normalized the predicted scores for each run by applying min-max normalization then we used interpolation to compute the mean curve.

4 RESULTS AND DISCUSSIONS

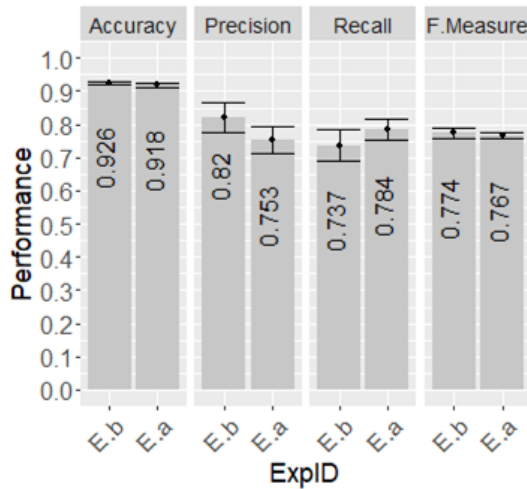


Figure 3: Comparison between experiments, E.a (for balanced distribution) and E.b (for natural distribution) for class C. Here, the top of each bar represents mean performance and the black line segment, which is known as the error bar, represents \pm standard deviation. The experiments are arranged according to the ascending order of the precision.

Figure 3 presents the results of our experiments based on the argmax decision on the predicted scores of the models. We present both the mean and standard deviation considering the 10 runs.

On Figure 3, the results show that the naturally imbalanced distribution (E.b) is better than the artificially balanced distribution (E.a) in terms of all evaluation metrics, except for recall. The precision of E.b (second sub-part) is 7% higher than that of E.a., which means that the distribution with over-represented class C (E.b) produces less false positives than the balanced distribution (E.a).

On the other hand, according to the recall (third part of Figure 3), the result is opposite, namely the balanced distribution (E.a) gives higher recall than the other distribution. One of the reasons could be the number of positive examples in the training set. To keep the same total number of training examples in both experiments (for a fair comparison), we did not use all the positive examples (class C) in E.b, while in E.a, we utilized almost all positive examples.

Since recall and precision are inversely proportional, we also plot the F1-measure to compare the overall result by considering both precision and recall at the same time. According to the F1-measure, E.b is better than the E.a.

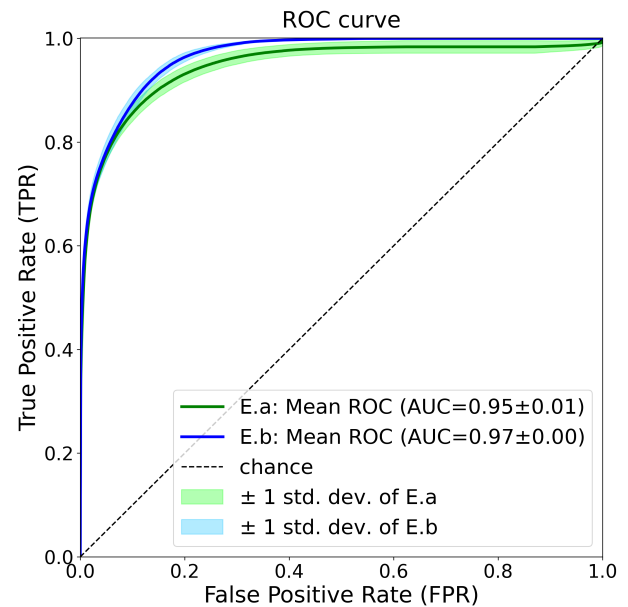


Figure 4: Mean ROC curves with standard deviation (std. dev.) of E.a (balanced distribution) and E.b (natural distribution) for class C.

While Figure 3 presents results for a single threshold point (cutoff), Figures 4 and 5 illustrate the results for all possible threshold points. In these figures, we present the mean curves resulting from the 10 runs as well as the standard deviation.

According to both ROC and PR curves, the natural distribution (E.b) is better than the balanced distribution (E.a). In Figure 4, we can observe that E.b has higher TPR (i.e., sensitivity or recall) than E.a, at almost every FPR points. In other words, models trained under

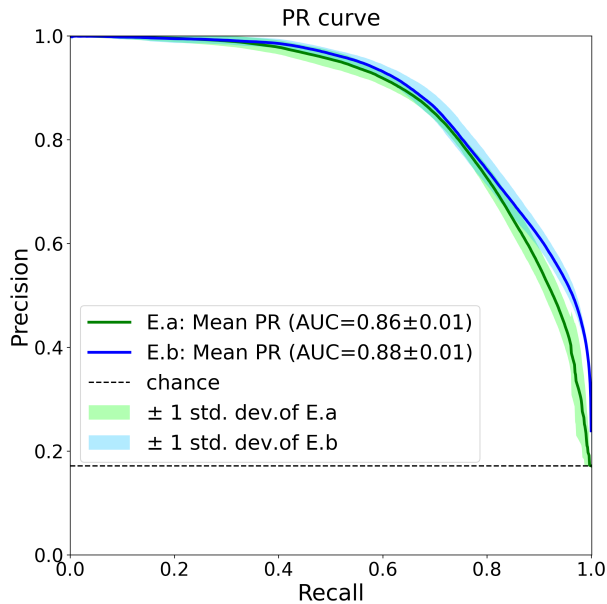


Figure 5: Mean PR curves with standard deviation (std. dev.) of E.a (balanced distribution) and E.b (natural distribution) for class C.

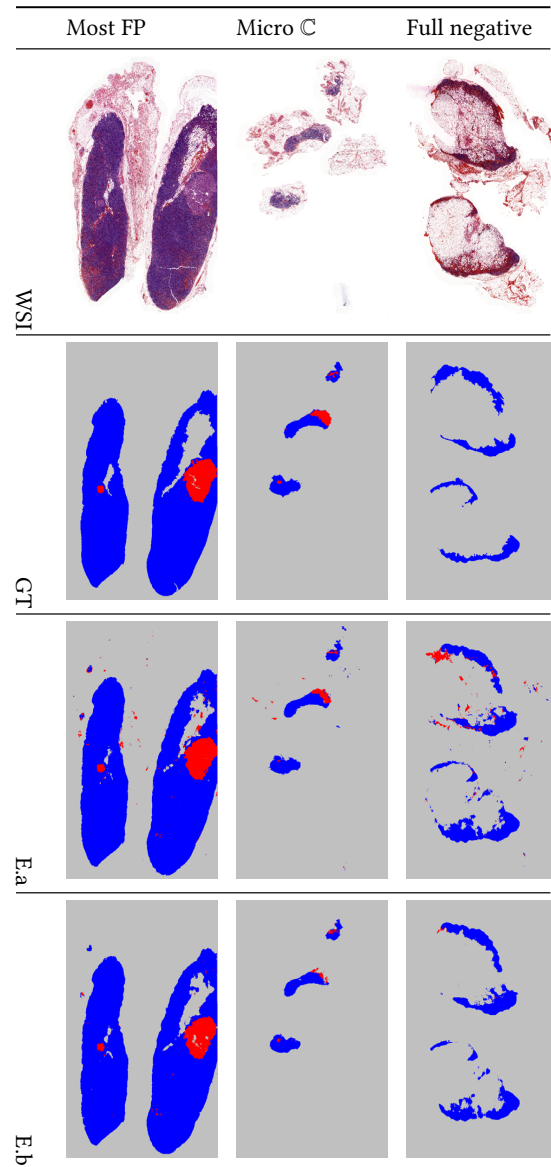
the natural distribution give higher recall, while also producing less false positives (FP) than the models trained under the balanced distribution.

Although the ROC curve is a popular metric to evaluate biomedical models, neither TPR nor FPR can capture the poor performance of the models for the imbalanced test data [28]. Thus, as an alternative, we report the results in terms of the PR curve as well (see Figure 5). According to the PR curve, the precision (i.e., PPV) of E.b is higher at almost every recall (i.e., TPR/sensitivity) point than E.a. Even at the recall point 1.0 (i.e., at 100% recall), the precision of E.b is better than that of the random chance baseline, while the precision of E.a is as low as the random chance.

In Table 4, we present some example predictions of E.a (3rd row) and E.b (4th row) with the corresponding test WSIs (1st row) and ground-truth (GT) masks (2nd row). Here, we choose the WSIs with three characteristics, most FP generative (1st column), WSI with micro metastasis/C (2nd column), and full negative (3rd column). The predictions show that E.b produces less FP than E.a.

We further investigate the reasons why using the natural distribution turned out to be superior, by manually looking at the predicted masks. According to our thorough observation, we figured out that the miss-classification by E.a (balanced distribution) is due to the inter-class similar regions: the regions containing common histological structures, e.g., blood, fibrous tissue, etc. These histological structures are common in both ROI (C and \neg C) and non-ROI (\emptyset) regions (see Figure 6). When an ROI contains such a histological structure by comprising a small area in a gigapixel image, it is practical to overlook that small area and annotate it

Table 4: Examples with ground-truth annotations and predictions taken from both E.a and E.b experiments. Best viewed in color.



as a corresponding ROI (C or \neg C) instead of annotating it as non-ROI, although the annotation is actually wrong. During training, we thus need enough examples of these inter-class similar regions (i.e., over-representation of class \emptyset) with their actual annotation to compensate the unavoidable pitfall of the annotation, otherwise it might cause false positives during prediction [29].

5 CONCLUSION

In this research, we performed a data-level comparison between naturally imbalanced and artificially balanced distributions, thus

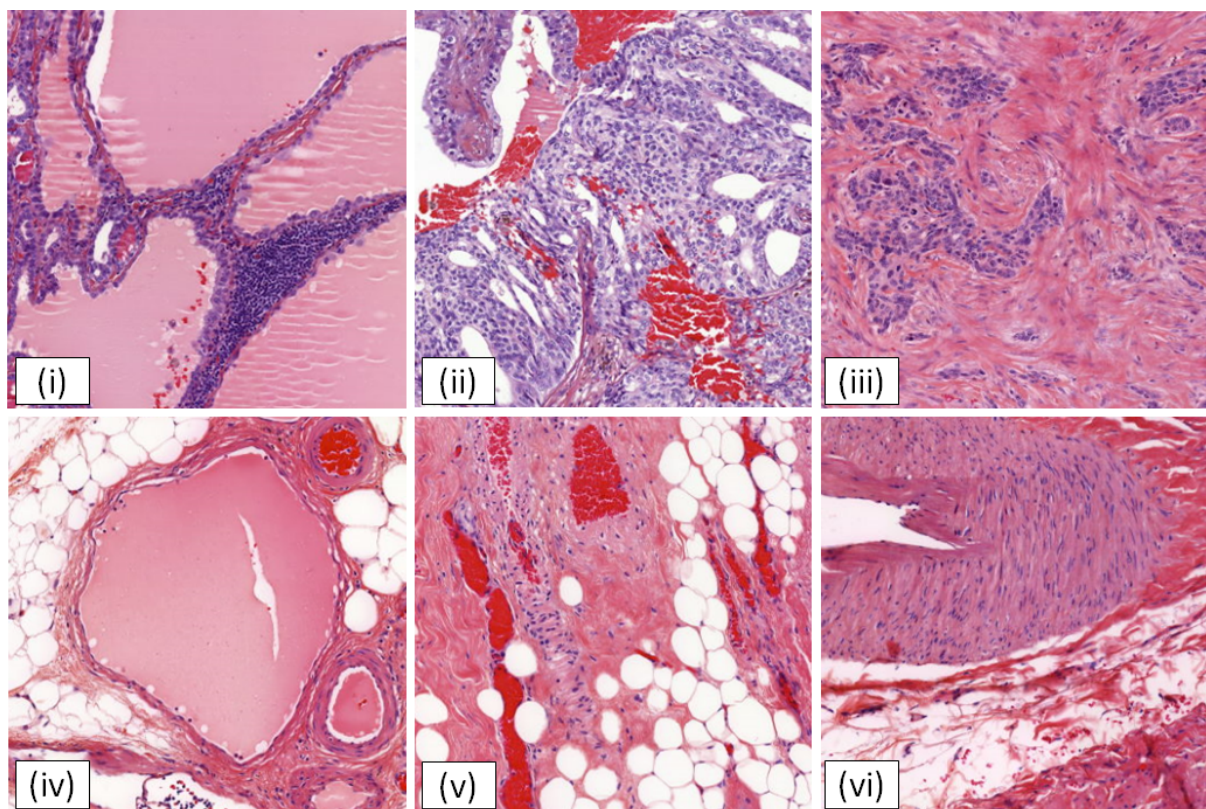


Figure 6: Some examples of confusion between classes \mathcal{O} and \mathcal{C} caused by inter-class similar region with common histological structures. Images (i), (ii) and (iii) represent patches from the training set that belong to category \mathcal{C} , which also contain fluid, blood and fibrosis with lymphocytes, respectively. Images (iv), (v) and (vi) represent patches from the test set that belong to category \mathcal{O} , containing the same histological structures as the above images. During training, if there are not enough examples of class \mathcal{O} containing the mentioned structures, the respective histological structures in the test patches are predicted as class \mathcal{C} (false positive). Best viewed in color.

determining the usability of the non-ROIs in WSIs for CNN training. According to our analysis, the commonly-prescribed balanced distribution for CNNs does not provide the best performance. Instead, the original training set, which is naturally biased towards the non-ROI class, produces the best performance for cancer detection. Moreover, the results illustrate the usefulness of the usually filtered out non-ROI data. Specifically, these examples are useful to be added during CNN training to solve the inter-class similarity between ROIs and non-ROIs, thus reducing false positives during prediction. Notably, prediction with less false positives reduces pathologists' workloads during cancer detection. We believe that the findings of this study can be useful for researchers in cancer detection from WSIs, providing a new viewpoint that can make them reconsider using the natural distribution and the non-ROIs in building training data. In some regard, our work is limited because we only considered the natural and the balanced distribution. In future work, it would be interesting to observe the impact of other distributions, e.g., lymph-node biased or cancer biased distributions.

ACKNOWLEDGEMENTS

Authors would like to thank Professor Pierre Brousset, Dr. Camille Franchet, and Dr. Margot Gaspard for their fruitful discussions and providing the data set.

The research leading to these results has received funding from the NO Grants 2014-2021, under project ELO-Hyp contract no. 24/2020.

REFERENCES

- [1] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermesen, Quirine F Manson, Maschenka Balkenhol, et al. 2017. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama* 318, 22 (2017), 2199–2210.
- [2] Mateusz Buda, Atsuto Maki, and Maciej A. Mazurowski. 2017. A systematic study of the class imbalance problem in convolutional neural networks. *CoRR* abs/1710.05381 (2017).
- [3] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [4] Nitesh V Chawla, Aleksandar Lazarevic, Lawrence O Hall, and Kevin W Bowyer. 2003. SMOTEBoost: Improving prediction of the minority class in boosting. In *European conference on principles of data mining and knowledge discovery*. Springer, 107–119.
- [5] François Chollet et al. 2015. Keras. <https://keras.io>.

- [6] Dua Dheeru and Efi Karra Taniskidou. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [7] Navid Farahani. 2015. Whole slide imaging in pathology: advantages, limitations, and emerging perspectives. (2015).
- [8] Martin Halicek, Maysam Shahedi, James V Little, Amy Y Chen, Larry L Myers, Baran D Sumer, and Baowei Fei. 2019. Head and neck cancer Detection in Digitized Whole-Slide Histology Using convolutional neural networks. *Scientific reports* 9, 1 (2019), 1–11.
- [9] Martin Halicek, Maysam Shahedi, James V Little, Amy Y Chen, Larry L Myers, Baran D Sumer, and Baowei Fei. 2019. Head and neck cancer Detection in Digitized Whole-Slide Histology Using convolutional neural networks. *Scientific reports* 9, 1 (2019), 1–11.
- [10] Nicolas Jaccard, Thomas W Rogers, Edward J Morton, and Lewis D Griffin. 2017. Detection of concealed cars in complex cargo X-ray imagery using deep learning. *Journal of X-ray Science and Technology* 25, 3 (2017), 323–339.
- [11] Justin M Johnson and Taghi M Khoshgoftaar. 2019. Survey on deep learning with class imbalance. *Journal of Big Data* 6, 1 (2019), 27.
- [12] Salman H Khan, Munawar Hayat, Mohammed Benmamoun, Ferdous A Sohel, and Roberto Togneri. 2017. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems* (2017).
- [13] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- [14] Oscar Koller, Hermann Ney, and Richard Bowden. 2015. Deep learning of mouth shapes for sign language. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*. 85–91.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105.
- [16] Miroslav Kubat, Stan Matwin, et al. 1997. Addressing the curse of imbalanced training sets: one-sided selection. In *Icml*, Vol. 97. Nashville, USA, 179–186.
- [17] Neeta Kumar, Ruchika Gupta, and Sanjay Gupta. 2020. Whole Slide Imaging (WSI) in Pathology: Current Perspectives and Future Directions. *Journal of Digital Imaging* (2020).
- [18] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (May 2015), 436–444. <https://doi.org/10.1038/nature14539>
- [19] Gil Levi and Tal Hassner. 2015. Age and gender classification using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 34–42.
- [20] Huangjing Lin, Hao Chen, Qi Dou, Liansheng Wang, Jing Qin, and Pheng-Ann Heng. 2018. Scannet: A fast and dense scanning framework for metastatic breast cancer detection from whole-slide image. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 539–546.
- [21] Huangjing Lin, Hao Chen, Qi Dou, Liansheng Wang, Jing Qin, and Pheng-Ann Heng. 2018. Scannet: A fast and dense scanning framework for metastatic breast cancer detection from whole-slide image. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 539–546.
- [22] Yun Liu, Krishna Kumar Gadepalli, Mohammad Norouzi, George Dahl, Timo Kohlberger, Subhashini Venugopalan, Aleksey S Boyko, Aleksei Timofeev, Philip Q Nelson, Greg Corrado, Jason Hipp, Lily Peng, and Martin Stumpe. 2017. Detecting Cancer Metastases on Gigapixel Pathology Images. <https://arxiv.org/abs/1703.02442> Initial publication on arxiv, then submit to MICCAI.
- [23] Yun Liu, Timo Kohlberger, Mohammad Norouzi, George E Dahl, Jenny L Smith, Arash Mohtashamian, Niels Olson, Lily H Peng, Jason D Hipp, and Martin C Stumpe. 2018. Artificial Intelligence–Based Breast Cancer Nodal Metastasis Detection: Insights Into the Black Box for Pathologists. *Archives of pathology & lab. medicine* (2018).
- [24] Ronaldo C Prati, Gustavo EAPA Batista, and Diego F Silva. 2015. Class imbalance revisited: a new experimental setup to assess the performance of treatment methods. *Knowledge and Information Systems* 45, 1 (2015), 247–270.
- [25] Waseem Rawat and Zenghui Wang. 2017. Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review. *Neural Comput.* 29, 9 (Sept. 2017), 2352–2449. <https://doi.org/10.1162/neco-a-00990>
- [26] O. Ronneberger, P. Fischer, and T. Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI) (LNCS)*, Vol. 9351. Springer, 234–241. <http://lmb.informatik.uni-freiburg.de/Publications/2015/RF15a>
- [27] Anandarup Roy, Rafael MO Cruz, Robert Sabourin, and George DC Cavalcanti. 2018. A study on combining dynamic selection and data preprocessing for imbalance learning. *Neurocomputing* 286 (2018), 179–192.
- [28] Takaya Saito and Marc Rehmsmeier. 2015. The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS one* 10, 3 (2015), e0118432.
- [29] Leslie Shanks, Derryck Klarkowski, and Daniel P O'Brien. 2013. False positive HIV diagnoses in resource limited settings: operational lessons learned for HIV programmes. *PLoS One* 8, 3 (2013), e59906.
- [30] Yanmin Sun, Mohamed S Kamel, Andrew KC Wong, and Yang Wang. 2007. Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition* 40, 12 (2007), 3358–3378.
- [31] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. 2018. Rotation equivariant CNNs for digital pathology. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 210–218.
- [32] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H. Beck. 2016. Deep Learning for Identifying Metastatic Breast Cancer. *CoRR* abs/1606.05718 (2016). <http://dblp.uni-trier.de/db/journals/corr/corr1606.html#WangKGB16>
- [33] Jiazhao Wang, John D MacKenzie, Rageshree Ramachandran, Y Zhang, H Wang, and Danny Ziyi Chen. 2015. Segmenting subcellular structures in histology tissue images. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 556–559.
- [34] Gary M Weiss and Foster Provost. 2001. The effect of class distribution on classifier learning: an empirical study. *Rutgers Univ* (2001).
- [35] Gary M Weiss and Foster Provost. 2003. Learning when training data are costly: The effect of class distribution on tree induction. *Journal of Artificial Intelligence Research* 19 (2003), 315–354.
- [36] Yirui Wu, Yukai Ding, and Jun Feng. 2020. SMOTE-Boost-based sparse Bayesian model for flood prediction. *EURASIP Journal on Wireless Communications and Networking* 2020 (2020), 1–12.
- [37] Xiaohui Yuan, Lijun Xie, and Mohamed Abouelenien. 2018. A regularized ensemble framework of deep learning for cancer detection from multi-class, imbalanced training data. *Pattern Recognition* 77 (2018), 160–172.