

H&E-adversarial network: a convolutional neural network to learn stain-invariant features through Hematoxylin & Eosin regression

Niccolò Marini^{*1,2}, Manfredo Atzori^{1,3}, Sebastian Otálora^{1,4}, Stephane Marchand-Maillet², and Henning Müller^{1,5}

¹Institute of Information Systems, HES-SO (University of Applied Sciences and Arts Western Switzerland)

²Department of Computer Science, University of Geneva

³Department of Neurosciences, University of Padua

⁴Inselspital, Bern Hospital

⁵Medical Faculty, University of Geneva

Abstract

Computational pathology is a domain that aims to develop algorithms to automatically analyze large digitized histopathology images, called whole slide images (WSI). WSIs are produced scanning thin tissue samples that are stained to make specific structures visible. They show stain colour heterogeneity due to different preparation and scanning settings applied across medical centers. Stain colour heterogeneity is a problem to train convolutional neural networks (CNN), the state-of-the-art algorithms for most computational pathology tasks, since CNNs usually underperform when tested on images including different stain variations than those within data used to train the CNN. Despite several methods that were developed, stain colour heterogeneity is still an unsolved challenge that limits the development of CNNs that can generalize on data from several medical centers. This paper aims to present a novel method to train CNNs that better generalize on data including several colour variations. The method, called H&E-adversarial CNN, exploits H&E matrix information to learn stain-invariant features during the training. The method is evaluated on the classification of colon and prostate histopathology images, involving eleven heterogeneous datasets, and compared with five other techniques used to handle stain colour heterogeneity. H&E-adversarial CNNs show an improvement in performance compared to the other algorithms, demonstrating that it can help to better deal with stain colour heterogeneous images.

*niccolo.marini@hevs.ch

1. Introduction

Stain colour heterogeneity remains a critical challenge in computational pathology [39], despite the increasing number of methods developed to tackle the problem.

Stain colour heterogeneity is related to the colour variations of digitized histopathology images [40, 39, 21], called whole slide images (WSI). WSIs include slices of tissue samples collected through biopsies or surgical resections to identify findings that may lead to diseases such as cancer [14]. Colour variations are due to the inconsistencies in sample thickness, tissue preparation and tissue scanning across different medical centers [28, 36, 19, 31, 39]. Tissue preparation involves several procedures, including tissue staining. The staining involves the application of chemical reagents on the sample to increase contrast among structures within the tissue. Frequently, tissues are stained with the combination of Hematoxylin and Eosin (H&E). The Hematoxylin stains the cellular nuclei with shades of blue, while the Eosin stains the cytoplasm and extracellular structures with shades of pink [40]. The concentrations of H&E are not standardized, leading to variations of staining across medical centers. Tissue scanning is the high-resolution digital acquisition of the image through whole slide scanners. Several vendors produce whole slide scanners, such as Aperio, 3DHistech, Hamamatsu. Each type of scanner presents different properties and characteristics, such as the physical temperature [26] (that influence the stain reagents) or the light acquired by the scanner during the acquisition, that influence the colour response of the scanner [40]. Therefore, usually the images coming from a medical center show specific colours due to the prepara-

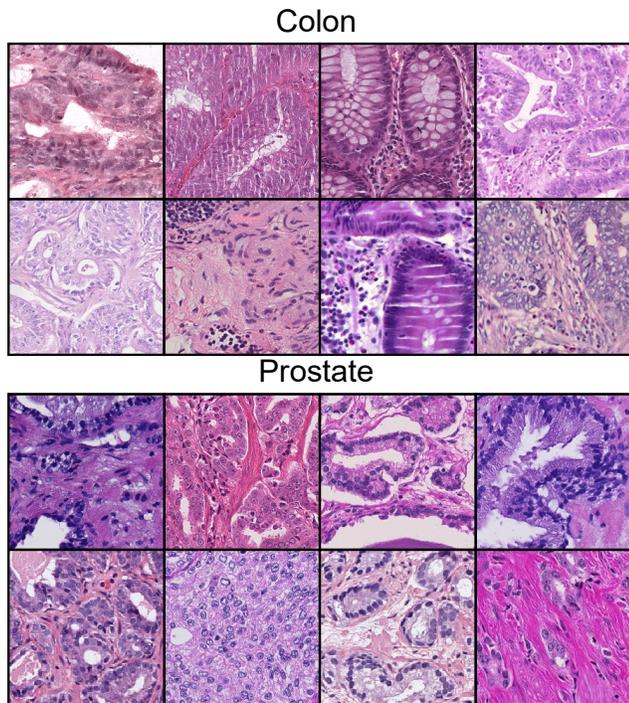


Figure 1. Example of colour variations in histopathology images, collected from publicly available repositories. The figure includes examples from colon tissue (above) and prostate tissue (below).

tion and the scanner adopted [38]. Figure 1 shows examples of stain variations.

Computational pathology is a domain that aims to develop algorithms to automatically analyze WSIs. Convolutional Neural Networks (CNN) are the state-of-the-art algorithm for most of the computational pathology tasks, reaching high performance in tasks such as classification and segmentation. Stain colour heterogeneity hinders the development of CNNs that generalize well on unseen data, especially when data originate from several medical centers [39, 41, 31, 19, 21]. CNN training aims to learn relevant features for solving a task despite the stain variations between the images [15]. However, CNNs usually underperform when tested on data with different colour variations than the ones used to train the network [39], leading to models that are not effective when using data originating from several medical centers [31]. This limitation increases the strain needed to develop models that produce consistent results and generalize well on heterogeneous data [19], with significant effects on the algorithms reliability [28].

Stain colour heterogeneity is still an unsolved problem in computational pathology, despite of the increasing number of methods developed to train CNNs that generalize better on stain heterogeneous data. Methods to tackle stain variations involve mainly two approaches: approaches using (1) the pixel-space and (2) the feature-space [41].

Pixel-space methods include techniques to modify the characteristics of the input images [10, 39, 28], such as colour normalization and colour augmentation. Colour normalization applies a transformation to the image in order to match the stain of another image template [10]. Traditional colour normalization approaches aim to estimate a colour matrix, that identifies the H&E components of the image, such as [22, 40, 30]. More recent approaches, such as [9, 33, 17, 4], involve Generative Adversarial Networks (GAN) to normalize the images, learning to match different stain distributions. Colour augmentation applies a random perturbation to the image in order to have several stain variations during training, working on the matrix that describes H&E components of the image or on the brightness and contrast of the image [39]. Among these two techniques, colour augmentation has shown to often reach higher performance [39]. However, the augmentation requires to tune the transformation parameters to avoid using artefacts to train the network, since the perturbation is often random. Feature-space methods include techniques, such as domain-adversarial networks [28, 31, 15, 21], that force the CNN to learn domain-invariant features during the training. Domain-adversarial networks work under the assumption that images coming from the same domain (such as the medical center, the patient or the tissue) share the same staining. The CNNs, in parallel to the main task, learn features that are invariant to the domain where the image originates from [13] and as a consequence invariant to the stain variations. Domain-adversarial networks have often shown higher performance when compared with colour augmentation, such as in [28, 21]. However, the definition of domain can be too strict, e.g. in the case where each patient is considered a domain [15], or too broad, e.g. in the case where each medical center is considered a domain [28, 21]. In the first case, images from different patients can have the same stain; in the latter case, images coming from the same medical center can include different stains.

This paper aims to describe a novel integrated method to train CNNs that better generalize on data including images with colour variations that are different from the ones used for training. The method is based on a multi-task CNN, referred as H&E-Adversarial Network (H&E-AN). H&E-AN includes two output branches: the first one classifies histopathology images, while the second one predicts H&E matrices linked to the images. H&E-AN training involves the optimization of a loss function including two terms, in order to learn discriminative and stain-invariant features through the adversarial optimization, similar to the one presented in domain-adversarial CNNs [13, 28, 21, 15]. The adversarial optimization goes in the opposite gradient direction that minimizes the patch-classifier loss (to learn discriminative features respect to the classes) and in the positive direction of the gradient for the H&E features, to max-

imize the H&E-regressor loss.

The H&E-AN CNN is evaluated on colon and prostate cancer WSI classification, comparing it with other algorithms developed to alleviate stain colour heterogeneity. Colon cancer is the fourth most commonly diagnosed cancer in the world [3]. Colon cancer diagnosis involves the detection of malignant glands and small agglomerations of cells [11], called polyps, on the colon border. Prostate cancer (PCa) is the second most frequent cancer in the male population worldwide [29]. Prostate cancer is diagnosed using the Gleason grading system [6] that involves the detection and the estimation of malignant glands in order to estimate the aggressiveness of cancer. The comparison is made testing the methods on data originating from several unseen medical centers, to assess the CNN capability to generalize on images including stain variations. H&E-AN CNN reaches higher performance in classifying both prostate and colon images, compared with other methods to handle histopathology stain variability, showing that the method can help better deal with stain colour heterogeneous images.

2. Methods

The paper proposes H&E-AN CNN method that exploits the information related to H&E of the WSIs during the training of a CNN in WSI patch classification, to learn stain-invariant features. The information related to the H&E composition used to stain the image is a matrix M including the RGB components of each reagent.

$$M = \begin{pmatrix} H_R & H_G & H_B \\ E_R & E_G & E_B \end{pmatrix}$$

For each of the patches, the H&E components are evaluated using the method proposed by Macenko et al. [22], for each patch in the training set. Each element of the matrix includes values ranging from 0 to 1. The CNN is a multi-task network $f(\theta_{CNN}, X) \rightarrow (Y, M)$ with trainable parameters θ_{CNN} , that aims to predict the classes Y and the H&E components M from the input patches X . The CNN includes three main components: a convolutional layer block ($ConvL$), a classifier (Cl) and a regressor (r). The convolutional layers are a function $f(\theta_{ConvL}, x_i) \rightarrow feat$ with trainable parameters θ_{ConvL} that produces a feature vector $feat_i$ from an input patch x_i . The classifier is a function $f(\theta_{Cl}, feat_i) \rightarrow y'_i$ with trainable parameters θ_{Cl} that makes class predictions y'_i from $feat_i$. The regressor is a function $f(\theta_r, feat_i) \rightarrow m'_i$ with trainable parameters θ_r that predicts the H&E matrix component m'_i from $feat_i$.

The training of the network aims to learn features that 1) reach high performance in patches classification and 2) do not take into account the stain variations of the training samples. Therefore, the CNN must learn a feature representation $f(\theta_{ConvL})$ of the input data that is discriminative and stain-invariant [12, 28]. The training of the CNN involves

the optimization of θ_{ConvL} , θ_{Cl} and θ_r . For the training, the loss function $Loss$ involves two terms: $Loss_{Cl}$, that measures the error in classification task and $Loss_r$, that measures the error in the regression task. The losses are described as follows (where \mathcal{L}_{Cl} and \mathcal{L}_r are a measure of discrepancy):

$$Loss_{Cl} = \left(\frac{1}{N} \sum_{n=1}^N \mathcal{L}_{Cl}(f(\theta_{ConvL}, \theta_{Cl}, x_i), y_i) \right) \quad (1)$$

$$Loss_r = \left(\frac{1}{N} \sum_{n=1}^N \mathcal{L}_r(f(\theta_{ConvL}, \theta_r, x_i), m_i) \right) \quad (2)$$

$$Loss = Loss_{Cl} + \lambda * Loss_r \quad (3)$$

The equations show that both the classifier and the regressor share the same set of features θ_{ConvL} , that should be both discriminative with respect to the classes and stain-invariance. The mechanism to have discriminative features is to minimize the $Loss_{Cl}$ (using Cross-Entropy as $Loss_{Cl}$) term of the loss function. The mechanism to obtain stain-invariant features is to maximize the $Loss_r$ (using squared L2-Loss as $Loss_r$) of the loss function, as proposed in domain-adversarial networks [13, 21, 28, 31, 15]. While the minimization of regression loss function would lead to learn the features to discriminate among stain variations, its maximization leads to learn stain-invariant features, so that the stain variations are indistinguishable for the model. The squared L2-Loss function is chosen considering that the M space does not include outliers that can hinder the regression, since the values in M range from 0 to 1. Regression term maximization is achieved by modifying the CNN, including a layer that reverses the gradients (through the multiplication of the gradient with a negative scalar λ) during the backpropagation and leave them unchanged during the forward propagation [12]. The network iterative stochastic gradient mechanism for respectively θ_{ConvL} , θ_{Cl} , θ_r (μ is the learning rate) is described as follows:

$$\theta_{ConvL} \leftarrow \theta_{ConvL} - \mu \left(\frac{\partial \mathcal{L}_{Cl}^i}{\partial \theta_{ConvL}} - \lambda \frac{\partial \mathcal{L}_r^i}{\partial \theta_{ConvL}} \right) \quad (4)$$

$$\theta_{Cl} \leftarrow \theta_{Cl} - \mu \left(\frac{\partial \mathcal{L}_{Cl}^i}{\partial \theta_{Cl}} \right) \quad (5)$$

$$\theta_r \leftarrow \theta_r - \lambda \mu \left(\frac{\partial \mathcal{L}_r^i}{\partial \theta_r} \right) \quad (6)$$

The CNN architecture is shown in Figure 2.

3. Experiments

Datasets The method is trained and evaluated on highly-heterogeneous histopathological images, including colon and prostate tissue. In both partitions, the images originate from several medical sources, guaranteeing high heterogeneity in terms of stain and allowing to evaluate the

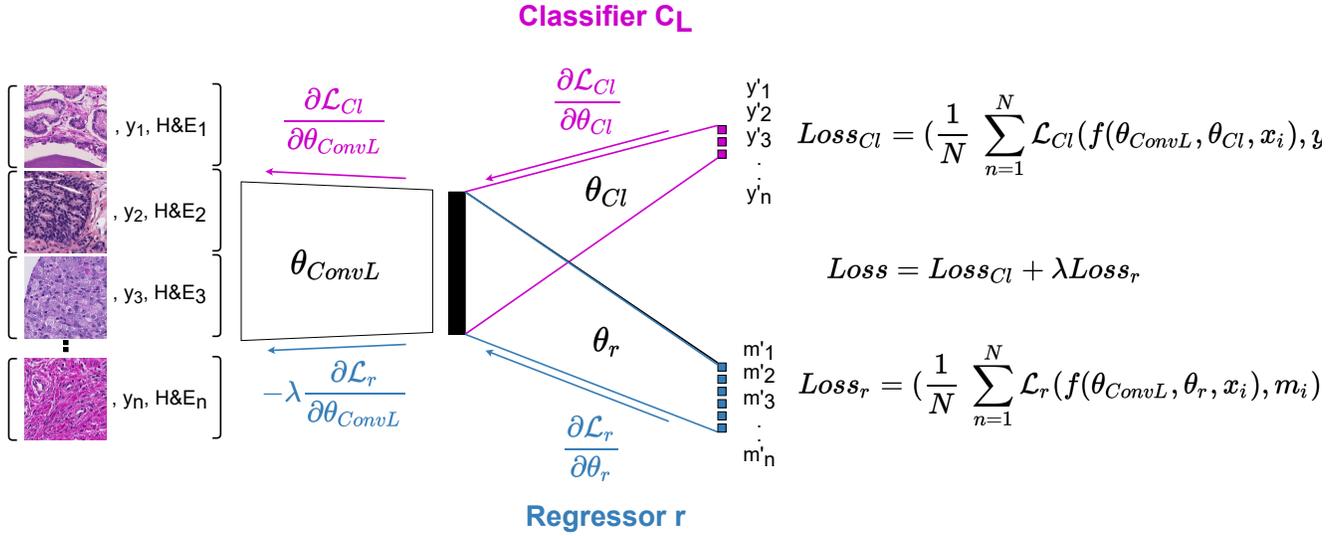


Figure 2. Overview of the H&E-adversarial CNN and of the training schema. The network includes three main components: convolutional layers $ConvL$, that produce discriminative and stain-invariant features, the classifier Cl (purple) that predicts the class for the patches and the regressor r (blue) that predicts the H&E matrix component of the patches. The training involves the optimization of two loss functions: the minimization of $Loss_{Cl}$ for the classifier and the maximization of $Loss_r$ for the regressor. The latter term of the loss function is maximized through a gradient reverse layer, that applies a negative scalar to the gradients during the backpropagation. This schema aims to obtain discriminative (thanks to the classifier) and stain-invariant (thanks to the regressor) features.

Table 1. Composition of the colon dataset. The colon dataset originates from seven medical sources and it is annotated with three classes: Cancer, Dysplasia and Normal. Data originating from hospitals (AOEC and Radboudumc) are used to train, validate and as internal test set for the CNN. Data originating from publicly available repositories are used as external test set to evaluate the capability of the CNN to generalize on stain heterogeneous data.

Medical Source	Cancer	Dysplasia	Normal	Total
AOEC	6158	23312	3853	33323
Training	4059	13170	3402	20631
Validation	844	4005	78	4927
Testing	1255	6137	373	7765
Radboudumc	4430	3542	1998	9970
Training	2995	2498	1304	6797
Validation	643	707	365	1715
Testing	792	337	329	1458
Internal Testing data	2047	6474	702	9970
AIDA	7881	3296	31859	43036
GlaS	450	0	210	660
CRC	1507	0	1144	2651
UNITOPATHO	0	18064	3487	21551
CAMEL	0	15757	12030	27787
External Testing data	9838	37117	48730	95685

capability of the CNN to generalize on images with unseen stain variations. Image heterogeneity is shown in Figure 1 and Figure 3. As a basic overview, Figure 3 shows the distribution of the patches according to their H&E components. The components, highlighted with a different colour to identify the dataset where they originate from, are pro-

Table 2. Composition of the prostate dataset. The prostate dataset originates from four medical sources and it is annotated with four classes: Benign, Gleason Pattern 3 (GP3), Gleason pattern 4 (GP4) and Gleason pattern 5 (GP5). The TMAZ and Sicapv2 datasets are used to train, validate and as internal test set for the CNN. Gleason challenge and DiagSet datasets are used as external test set, to evaluate the capability of the CNN to generalize on stain heterogeneous data.

Medical Source	Benign	GP3	GP4	GP5	Total
TMAZ	3487	8946	7424	3610	23467
Training	2010	5992	4472	2766	15240
Validation	1350	1352	831	457	4927
Testing	127	1602	2121	387	4237
Sicapv2	11069	10784	2979	2767	27599
Training	9432	6499	2250	2011	20192
Validation	604	819	302	210	1935
Testing	1033	3466	427	546	5472
Internal Testing Data	1160	5068	2548	933	9709
Gleason Challenge	1080	2431	3649	100	7260
Diagset	8783	1243	4334	696	15056
External Testing Data	9863	3674	7983	796	22316

jected on a bidimensional space using the Principal Component Analysis (PCA) [16].

The colon partition includes images collected from seven medical sources. The images are WSIs (private dataset from Azienda Ospedaliera per le Emergenze e Urgenze Cannizzaro, AOEC, private dataset from Radboudumc, Analytic Imaging Diagnostics Arena [37], AIDA) and small cropped sections of WSIs (Gland Segmentation in Colon Histol-

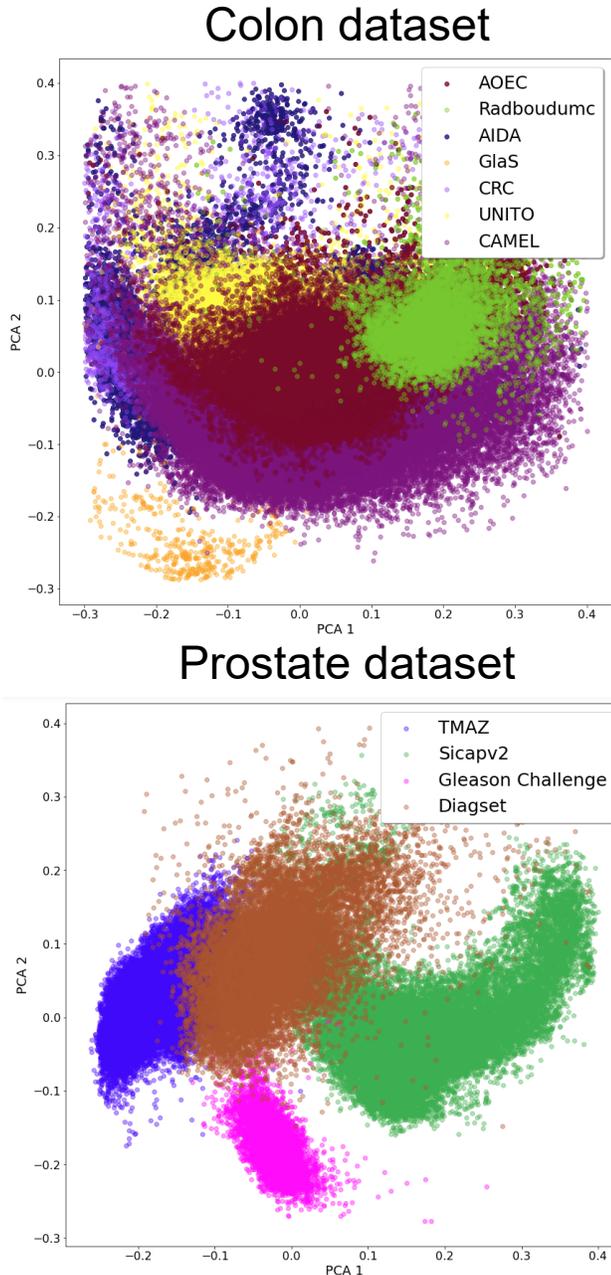


Figure 3. Stain heterogeneity across datasets. The Figure shows the H&E components of the patches for colon (upper part) and the prostate (lower part) datasets. The components are projected in a bidimensional space through PCA analysis. Each dot is a patch, highlighted with a different colour according to the source dataset.

ogy Images Challenge [35], GlaS, ColoRectal Cancer Tissue Phenotyping [2], CRC, UNITOPATHO [5], CAMEL dataset [43]). The WSIs are pixel-wise annotated by pathologists, while the cropped sections are annotated at section level. Among all the annotations available, the classes chosen to test the method are cancer, dysplasia and normal

tissue. AOEC and Radboudumc private datasets originate from two hospitals and are used to train, validate and test the method, as internal testing data. The other five datasets originate from publicly available repositories and are used as external test data. Table 1 summarizes the colon partition composition.

The prostate partition includes images collected from four publicly available datasets. The images are Tissue Micro Arrays (TMAs) (The Tissue Micro Array Zurich [1], TMAZ, and Gleason Challenge [27, 18]) and WSIs (Sicapv2 [34] and DiagSet [20]). The images are pixel-wise annotated by pathologists, with benign, Gleason Pattern 3 (GP3), Gleason Pattern 4 (GP4) and Gleason Pattern 5 as classes. TMAZ and Sicapv2 and are used to train, validate and test the method, as internal testing data. Gleason Challenge and DiagSet are used as external test data. Table 2 summarizes the prostate partition composition.

Pre-processing The image-preprocessing involves WSI splitting into patches and the extraction of the patches coming from tissue regions, for both colon and prostate. The patches are 224x224 pixels in size, in order to facilitate the input for the CNN. The colon dataset includes WSIs and cropped sections of WSIs. The WSIs come with a tissue mask, including pixel-wise annotations. WSIs are split into a grid of patches at magnification 10x. The grid is built based on the highest magnification level available M_m , with patches of p_s in size. Afterwards, grid patches are resized to 224x224, using the following equation, as shown in [24]:

$$224 : 10 = p_s : M_m \quad (7)$$

The cropped sections come without a tissue mask, but being small images, the section’s label is assigned to the tissue. In order to avoid the extraction of patches from background regions, a tissue mask is generated using the method proposed by [32]. Patches are randomly extracted, with a variable number depending on the size of the images, to avoid strong overlap of the patches: 20 from CRC, 2 from GlaS, 5 from UNITO and CAMEL. The numbers are chosen considering the amount of pixels included in each of the datasets.

The prostate dataset includes pixel-wise annotated TMAs and WSIs. The patches are extracted with a size of 750x750 pixels from magnification 40x and resized to 224x224, as shown in [1, 23]. In TMAs (TMAZ and Gleason Challenge), 30 patches are randomly extracted from each core. In WSIs (Sicapv2 and Diagset), the images are split into a grid of patches from the highest magnification level available, as shown for colon WSIs, using the following equation:

$$750 : 40 = p_s : M_m \quad (8)$$

Comparison with other methods The method proposed in this paper is compared with other methods developed

to train CNNs that generalize on data with high stain heterogeneity. All the methods are evaluated using quadratic weighted Cohen’s Kappa score (κ -score) [25] as metric and the Wilcoxon Rank-Sum test [42] is applied to check if the improvement obtained by the H&E-AN is statistically significant. The comparison involves six methods: a CNN trained without any augmentation or normalization of the images, a CNN trained using a stain normalization method, a CNN trained using stain normalization through a StainGAN, a CNN trained using colour augmentation methods, a CNN trained using stain augmentation, a CNN trained using a domain-adversarial network. For each method, the performance is evaluated on the internal test set, on the external test and on their combination. The stain normalization method is the one proposed by Macenko et al. [22]. For each type of tissue, a target image is randomly selected so that the stains of all the images in training and testing sets matches the target image’s stain. The StainGAN network is the one proposed by Shaban et al. [33], to normalize images from different domains. In this case, the StainGAN is pre-trained to normalize images through the training datasets, for colon and prostate. Afterwards, the network is used to normalize images during CNN training. The stain augmentation method involves the perturbation of the colours within the images [39]. The parameters involved in the random perturbation are the brightness, hue, saturation. The stain augmentation method involves the perturbation of the H&E matrix [39]. Two parameters are involved: a random constant σ_1 (multiplied to the H&E matrix) that scales the colour intensity and a random constant σ_2 (added to the H&E matrix * σ_1 product) that increase or reduce the colour intensity. The domain-adversarial CNN is the one proposed by Otolara et al. [28]. The network is a multi-task CNN, that predicts the patches and predicts the domain where they originate from. As described for the StainGAN implementation, the domain-adversarial CNN is trained using data from two domains.

The κ -score is a metric that measures the reliability between annotations, commonly used in histopathology to assess the pathologists’ performance in the evaluation of images. In this case, the annotations are the predictions made by the CNN and the ground truth made by pathologists. The optimal value of the κ -score is 1, which means a complete agreement between the annotations, while a κ -score equal -1 means a complete disagreement between the annotations. κ -score equal 0 means random agreement, as the measure is normalized with respect to agreement by chance.

The Wilcoxon Rank-Sum test is used to assess if two probabilistic populations have the same distribution (null hypothesis). The null hypothesis is tested positive when the p -value > 0.05 , while it is tested negative (or rejected) when the p -value < 0.05 . In this case, the comparison is made between the H&E-AN and the method that reaches the highest

performance among the six other methods evaluated. If the null hypothesis is tested negative, the improvement is statistically significant.

Experimental parameters The proposed H&E-adversarial CNN and the methods used for comparison are implemented on the same backbone architecture and are trained multiple times, adopting the same strategy to set the hyper-parameters and to face the class imbalance. The backbone architecture is a DenseNet121 (pre-trained on ImageNet). The network produces a feature vector of size 1024 for each input patch. An intermediate fully connected layer with 128 nodes is inserted between the classifier and the output of the classifier. For each method, the CNN is trained ten times (average and standard deviation of the models are reported) to limit the non-deterministic effect of the stochastic gradient descent used to optimize the model. CNN hyper-parameters are chosen after the evaluation of a grid search algorithm [8]. Grid search aims to find an optimal configuration for the CNN hyperparameters. In this case, the optimal configuration is the one that allows the CNN to have the lowest loss function on the validation partition. The hyper-parameters involved in the grid search are the optimizer (Adam), the learning rate (10^{-3}), the decay rate (10^{-4}), the number of epochs (15 epochs with an early stop mechanism used to stop the training if the validation loss function has not decreased for more than five epochs), the number of nodes within the intermediate fully-connected layers (128), σ_1 and σ_2 for the stain augmentation algorithm (both 0.2), the λ parameter for the domain-adversarial CNN and for H&E-AN (respectively 0.5 and 1). Moreover, the grid search algorithm is adopted to identify the shifts for colour algorithm: in the colon data the hue shift is limited to be between -15 and 8, the saturation shift between -20 and 10 and the brightness shift between -8 and 8, while in prostate data the hue shift is limited to be between -9 and 9, the saturation shift is limited to be between -25 and 25 and the brightness shift is limited to be between -10 and 10. The effects of class imbalance are handled adopting a class-wise data augmentation method, involving three operations: rotations, flipping and colour augmentation. Data augmentation is implemented with the Albumentations library [7].

4. Results

The H&E-AN outperforms the CNNs trained using other methods, demonstrating the capability to better alleviate stain heterogeneity in colon and prostate histopathology image classification.

The evaluation on colon data involves the classification of cancer, dysplasia and normal glands on the internal test set (including AOEC and Radboudumc test partitions), the external test set (including AIDA, GlAS, CRC, UNITO and

Table 3. Overview of the CNN results on the colon test partitions, assessed with the κ -score for each of the methods tested. The results that are statistically significant (compared with the method reaching the highest performance among the ones used as comparison) are reported with an asterisk (*).

Method	Internal Test Set	External Test Set	Test Set
no augmentation	0.644 \pm 0.037	0.424 \pm 0.074	0.453 \pm 0.075
stain normalization [22]	0.650 \pm 0.036	0.492 \pm 0.037	0.510 \pm 0.034
StainGAN [33]	0.474 \pm 0.062	0.463 \pm 0.053	0.487 \pm 0.051
colour augmentation	0.636 \pm 0.032	0.488 \pm 0.049	0.518 \pm 0.046
stain augmentation	0.655 \pm 0.033	0.498 \pm 0.041	0.515 \pm 0.034
domain adversarial [28]	0.621 \pm 0.044	0.497 \pm 0.039	0.513 \pm 0.046
Our Method	0.661 \pm 0.030	0.532 \pm 0.038*	0.556 \pm 0.040*

Table 4. Overview of the CNN results on the prostate test partitions, assessed with the κ -score for each of the methods tested. The results that are statistically significant (compared with the method reaching highest performance among the ones used as comparison) are reported with an asterisk (*).

Method	Internal Test Set	External Test Set	Test Set
no augmentation	0.712 \pm 0.026	0.052 \pm 0.099	0.142 \pm 0.081
stain normalization [22]	0.679 \pm 0.066	0.271 \pm 0.101	0.336 \pm 0.078
StainGAN [33]	0.633 \pm 0.065	0.244 \pm 0.113	0.344 \pm 0.080
colour augmentation	0.714 \pm 0.044	0.418 \pm 0.043	0.476 \pm 0.040
stain augmentation	0.687 \pm 0.040	0.312 \pm 0.104	0.298 \pm 0.086
domain adversarial [28]	0.670 \pm 0.071	0.390 \pm 0.110	0.449 \pm 0.098
Our Method	0.725 \pm 0.035	0.474 \pm 0.066*	0.532 \pm 0.057*

CAMEL datasets) and the cumulative test set, including the combination of the internal and external test sets. The H&E-AN outperforms the other methods in each of the colon test sets. While in the internal test set the performance of the method is comparable with the ones of the other methods, on both the external and cumulative test sets the H&E-AN method obtains statistically significant improvements over the other methods. The CNN reaches respectively a κ -score = 0.532 \pm 0.038 (higher than the best among the other methods used for the comparison, the CNN trained with stain augmentation on the external test set) and κ -score = 0.556 \pm 0.040 (higher than the best among the other methods used for the comparison, the CNN trained with stain augmentation on the cumulative test set).

The evaluation on prostate involves the classification of benign, Gleason pattern 3 (GP3), Gleason pattern 4 (GP4) and Gleason pattern 5 (GP5) on the internal test set (including the TMAZ and Sicapv2 test partitions), the external test set (including Gleason Challenge and DiagSet datasets) and the cumulative test set, including the combination of the internal and external test sets. The H&E-AN outperforms the other methods in each of the prostate test sets. While on the internal test set the performance of the method is comparable with the ones of the other methods, in both the external and cumulative test sets H&E-AN the method obtains statistically significant improvements over the other methods. The CNN reaches respectively a κ -score = 0.474 \pm 0.066 (higher than the best among the other methods used for the comparison, colour augmentation, on the external test set)

and κ -score = 0.532 \pm 0.057 (higher than the best among the other methods used for the comparison, colour augmentation, on the cumulative test set).

5. Discussion

The results show that the H&E-AN obtains higher performance compared with the other methods proposed in scientific literature to train CNNs that generalize on stain heterogeneous data, learning a stain-invariant feature representation. The H&E-AN reaches the highest performance in both colon and prostate image classification, considering both the internal and the external test set. The result suggests that the model generalizes better on datasets including heterogeneous stain variations. This fact can be explained considering the training mechanism of the multi-task network. During the training, the loss functions measuring the error in classification is minimized, while the loss function measuring the error in the regression of the H&E matrix is maximized. This loss function encourages the CNN to learn features that are determinant for the classification (through the minimization of classifier loss) and that are not dependent on the stain variations (through maximization of the regressor). Domain-adversarial CNNs inspire the mechanism. However, domain-adversarial CNNs rely on the domain concept, which may be fuzzy considering that images from the same domain may have different stain variations or that images from different domains may share the same stain variations. On the other hand, the regression of the

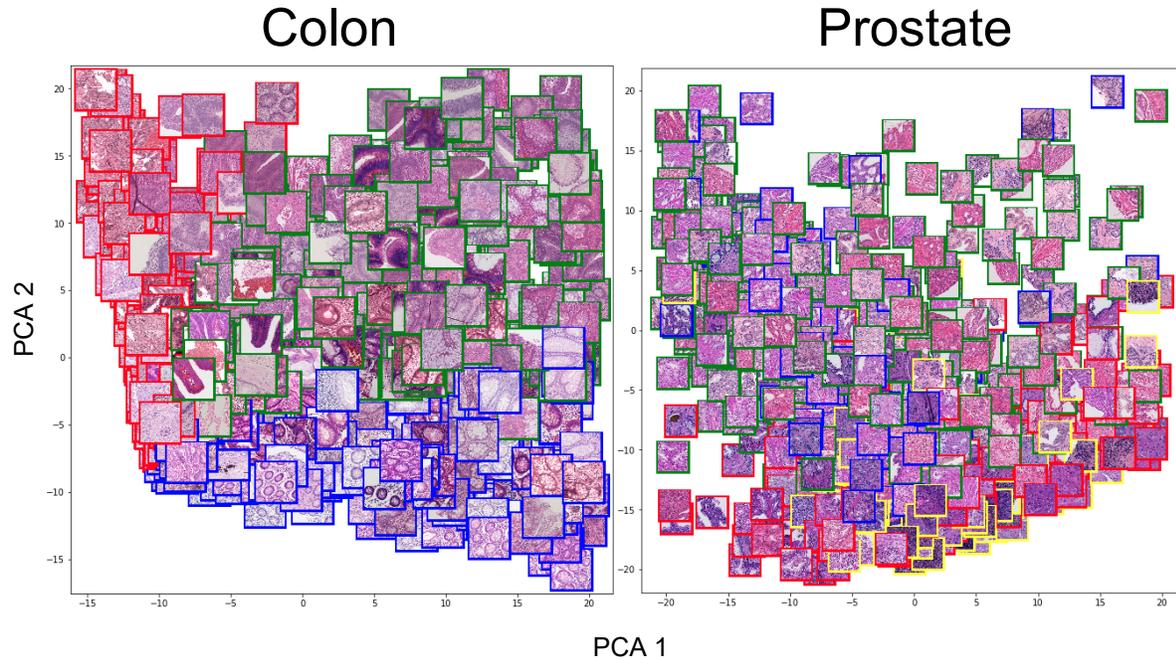


Figure 4. Stain-invariant feature representation of the images. The figure shows the distribution learnt by the CNN for the colon patches (left) and prostate (right), projecting the features in a bidimensional space through the PCA algorithm. The features are invariant to the stain, since it is possible to identify regions including a structure, such as glands, where several stain variations are present. The patches are highlighted with colour corresponding to the predicted class. In the colon subfigure: red is cancer, green is dysplasia, blue means normal glands. In prostate: red is benign, green is Gleason pattern 3, blue is Gleason pattern 4, yellow is Gleason pattern 5.

H&E matrix allows the model to directly learn features invariant to the stain instead of features invariant to the domain. The internal test set (including data coming from the same repository used to train the CNNs) and the external one (including data coming from unseen repositories) are heterogeneous. This fact can be seen looking at Figure 3, which shows the patch distribution highlighting the dataset source and considering the results reached by the CNN trained without any normalization or augmentation. Without keeping into account data heterogeneity, the CNN reaches good performance on the internal test set, where the stain variations are similar to the ones within the training set. However, the CNN underperforms (on prostate data the model can be considered as a random classifier) when tested on the external test set, including different stain variations.

The fact that the H&E-AN reaches the highest performance on the external data can be explained with the fact that the CNN learns stain-invariant features during the training, as shown in Figure 4. The Figure shows the feature representation of the H&E-AN, for both colon and prostate images. The features (from *ConvL* block) of each patch are projected on a bidimensional space, using the PCA algorithm. Within the distribution, it is possible to identify regions including similar structures, such as healthy glands in colon (bottom right) or malignant glands infiltrated with cells in prostate (top left). Those regions includes patches

with different stain variations, that are positioned in the same zone.

6. Conclusions

This paper introduces a novel H&E-AN to train CNNs that better generalize on stain heterogeneous data. The method exploits the H&E information of the images to learn stain invariant-features. The method involves a multi-task CNN, including a classifier (to predict patch classes) and a regressor (to predict the H&E matrix). The H&E-AN outperforms other methods considered in the comparison, including colour normalization, StainGAN, colour augmentation and domain-adversarial CNN, demonstrating its usefulness to handle stain variability in digital pathology images. We plan to test H&E-AN on other tasks (such as segmentation) and other datasets, evaluating the possibility to combine the method with other techniques and expanding the comparison to other methods. The code with the H&E-AN implementation is publicly available on Github (https://github.com/ilmaro8/HE_adversarial_CNN).

References

- [1] Eirini Arvaniti, Kim S Fricker, Michael Moret, Niels Rupp, Thomas Hermanns, Christian Fankhauser, Norbert Wey, Peter J Wild, Jan H Rueschoff, and Manfred Claassen. Auto-

- mated gleason grading of prostate cancer tissue microarrays via deep learning. *Scientific reports*, 8(1):1–11, 2018. 5
- [2] Ruqayya Awan, Korsuk Sirinukunwattana, David Epstein, Samuel Jefferyes, Uvais Qidwai, Zia Aftab, Imaad Mujeeb, David Snead, and Nasir Rajpoot. Glandular morphometrics for objective grading of colorectal adenocarcinoma histology images. *Scientific reports*, 7(1):1–12, 2017. 5
- [3] Al B Benson, Alan P Venook, Mahmoud M Al-Hawary, Lynette Cederquist, Yi-Jen Chen, Kristen K Ciombor, Stacey Cohen, Harry S Cooper, Dustin Deming, Paul F Engstrom, et al. Nccn guidelines insights: colon cancer, version 2.2018. *Journal of the National Comprehensive Cancer Network*, 16(4):359–369, 2018. 3
- [4] Aïcha BenTaieb and Ghassan Hamarneh. Adversarial stain transfer for histopathology image analysis. *IEEE transactions on medical imaging*, 37(3):792–802, 2017. 2
- [5] Luca Bertero, Carlo Alberto Barbano, Daniele Perlo, Enzo Tartaglione, Paola Cassoni, Marco Grangetto, Attilio Fian-drotti, Alessandro Gambella, and Luca Cavallo. Unitopatho, 2021. 5
- [6] David G Bostwick. Grading prostate cancer. *American journal of clinical pathology*, 102(4 Suppl 1):S38–56, 1994. 3
- [7] A. Buslaev, A. Parinov, E. Khvedchenya, V. I. Iglovikov, and A. A. Kalinin. Alumentations: fast and flexible image augmentations. *ArXiv e-prints*, 2018. 6
- [8] Davide Chicco. Ten quick tips for machine learning in computational biology. *BioData mining*, 10(1):35, 2017. 6
- [9] Hyungjoo Cho, Sungbin Lim, Gunho Choi, and Hyunseok Min. Neural stain-style transfer learning using gan for histopathological images. *arXiv preprint arXiv:1710.08543*, 2017. 2
- [10] Francesco Ciompi, Oscar Geessink, Babak Ehteshami Bejnordi, Gabriel Silva De Souza, Alexi Baidoshvili, Geert Litjens, Bram Van Ginneken, Iris Nagtegaal, and Jeroen Van Der Laak. The importance of stain normalization in colorectal tissue classification with convolutional networks. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 160–163. IEEE, 2017. 2
- [11] Monika Ferlitsch, Alan Moss, Cesare Hassan, Pradeep Bhandari, Jean-Marc Dumonceau, Gregorios Paspatis, Rodrigo Jover, Cord Langner, Maxime Bronzwaer, Kumanan Nalankilli, et al. Colorectal polypectomy and endoscopic mucosal resection (emr): European society of gastrointestinal endoscopy (esge) clinical guideline. *Endoscopy*, 49(3):270–297, 2017. 3
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 3
- [13] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 2, 3
- [14] Metin N Gurcan, Laura E Boucheron, Ali Can, Anant Madabhushi, Nasir M Rajpoot, and Bulent Yener. Histopathological image analysis: A review. *IEEE reviews in biomedical engineering*, 2:147–171, 2009. 1
- [15] Noriaki Hashimoto, Daisuke Fukushima, Ryoichi Koga, Yusuke Takagi, Kaho Ko, Kei Kohno, Masato Nakaguro, Shigeo Nakamura, Hidekata Hontani, and Ichiro Takeuchi. Multi-scale domain-adversarial multiple-instance cnn for cancer subtype classification with unannotated histopathological images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3852–3861, 2020. 2, 3
- [16] Ian Jolliffe. Principal component analysis. *Encyclopedia of statistics in behavioral science*, 2005. 4
- [17] Hongtao Kang, Die Luo, Weihua Feng, Junbo Hu, Shaoqun Zeng, Tingwei Quan, and Xiuli Liu. Stainnet: a fast and robust stain normalization network. *arXiv preprint arXiv:2012.12535*, 2020. 2
- [18] Davood Karimi, Guy Nir, Ladan Fazli, Peter C Black, Larry Goldenberg, and Septimiu E Salcudean. Deep learning-based gleason grading of prostate cancer from histopathology images—role of multiscale decision aggregation and data augmentation. *IEEE journal of biomedical and health informatics*, 24(5):1413–1426, 2019. 5
- [19] Amjad Khan, Manfredo Atzori, Sebastian Otálora, Vincent Andrearczyk, and Henning Müller. Generalizing convolution neural networks on stain color heterogeneous data for computational pathology. In *Medical Imaging 2020: Digital Pathology*, volume 11320, page 113200R. International Society for Optics and Photonics, 2020. 1, 2
- [20] Michał Koziarski, Bogusław Cyganek, Bogusław Olborski, Zbigniew Antosz, Marcin Zydak, Bogdan Kwolek, Paweł Wkasowicz, Andrzej Bukala, Jakub Swadzba, and Piotr Sitkowski. Diagset: a dataset for prostate cancer histopathological image classification. *arXiv preprint arXiv:2105.04014*, 2021. 5
- [21] Maxime W Lafarge, Josien PW Pluim, Koen AJ Eppenhof, and Mitko Veta. Learning domain-invariant representations of histological images. *Frontiers in medicine*, 6:162, 2019. 1, 2, 3
- [22] Marc Macenko, Marc Niethammer, James S Marron, David Borland, John T Woosley, Xiaojun Guan, Charles Schmitt, and Nancy E Thomas. A method for normalizing histology slides for quantitative analysis. In *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1107–1110. IEEE, 2009. 2, 3, 6, 7
- [23] Niccolò Marini, Sebastian Otálora, Henning Müller, and Manfredo Atzori. Semi-supervised training of deep convolutional neural networks with heterogeneous data and few local annotations: An experiment on prostate histopathology image classification. *Medical image analysis*, page 102165, 2021. 5
- [24] Niccolò Marini, Sebastian Otálora, Damian Podareanu, Mart van Rijthoven, Jeroen van der Laak, Francesco Ciompi, Henning Müller, and Manfredo Atzori. Multi_scale_tools: a python library to exploit multi-scale whole slide images. *Frontiers in Computer Science*, 3:68, 2021. 5
- [25] Mary L McHugh. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282, 2012. 6
- [26] Yuri Murakami, Hikaru Gunji, Fumikazu Kimura, Masahiro Yamaguchi, Yoshiko Yamashita, Akira Saito, Tokiya Abe,

- Michiie Sakamoto, Pinky A Bautista, and Yukako Yagi. Color correction in whole slide digital pathology. In *Color and Imaging Conference*, volume 2012, pages 253–258. Society for Imaging Science and Technology, 2012. [1](#)
- [27] Guy Nir, Soheil Hor, Davood Karimi, Ladan Fazli, Brian F Skinnider, Peyman Tavassoli, Dmitry Turbin, Carlos F Villamil, Gang Wang, R Storey Wilson, et al. Automatic grading of prostate cancer in digitized histopathology images: Learning from multiple experts. *Medical image analysis*, 50:167–180, 2018. [5](#)
- [28] Sebastian Otálora, Manfredo Atzori, Vincent Andrearczyk, Amjad Khan, and Henning Müller. Staining invariant features for improving generalization of deep convolutional neural networks in computational pathology. *Frontiers in bioengineering and biotechnology*, 7:198, 2019. [1](#), [2](#), [3](#), [6](#), [7](#)
- [29] Prashanth Rawla. Epidemiology of prostate cancer. *World journal of oncology*, 10(2):63, 2019. [3](#)
- [30] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer graphics and applications*, 21(5):34–41, 2001. [2](#)
- [31] Jian Ren, Ilker Hacihaliloglu, Eric A Singer, David J Foran, and Xin Qi. Unsupervised domain adaptation for classification of histopathology whole-slide images. *Frontiers in bioengineering and biotechnology*, 7:102, 2019. [1](#), [2](#), [3](#)
- [32] Mikael Rousson, Martin Hedlund, Mats Andersson, Ludvig Jacobsson, Gunnar Lathen, Bjorn Norell, Oscar Jimenez-del Toro, Henning Mueller, and Manfredo Atzori. Tumor proliferation assessment of whole slide images. In *Medical imaging 2018: Digital pathology*, volume 10581, page 105810Y. International Society for Optics and Photonics, 2018. [5](#)
- [33] M Tarek Shaban, Christoph Baur, Nassir Navab, and Shadi Albarqouni. Staingan: Stain style transfer for digital histological images. In *2019 Ieee 16th international symposium on biomedical imaging (Isbi 2019)*, pages 953–956. IEEE, 2019. [2](#), [6](#), [7](#)
- [34] Julio Silva-Rodríguez, Adrián Colomer, María A Sales, Rafael Molina, and Valery Naranjo. Going deeper through the gleason scoring scale: An automatic end-to-end system for histology prostate grading and cribriform pattern detection. *Computer Methods and Programs in Biomedicine*, 195:105637, 2020. [5](#)
- [35] Korsuk Sirinukunwattana, David RJ Snead, and Nasir M Rajpoot. A stochastic polygons model for glandular structures in colon histology images. *IEEE transactions on medical imaging*, 34(11):2366–2378, 2015. [5](#)
- [36] Karin Stacke, Gabriel Eilertsen, Jonas Unger, and Claes Lundström. Measuring domain shift for deep learning in histopathology. *IEEE journal of biomedical and health informatics*, 25(2):325–336, 2020. [1](#)
- [37] Caroline Bivik Stadler, Martin Lindvall, Claes Lundström, Anna Bodén, Karin Lindman, Jeronimo Rose, Darren Treanor, Johan Blomma, Karin Stacke, Nicolas Pinchaud, et al. Proactive construction of an annotated imaging database for artificial intelligence training. *Journal of digital imaging*, 34(1):105–115, 2021. [4](#)
- [38] Zaneta Swiderska-Chadaj, Thomas de Bel, Lionel Blanchet, Alexi Baidoshvili, Dirk Vossen, Jeroen van der Laak, and Geert Litjens. Impact of rescanning and normalization on convolutional neural network performance in multi-center, whole-slide classification of prostate cancer. *Scientific Reports*, 10(1):1–14, 2020. [2](#)
- [39] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen van der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical image analysis*, 58:101544, 2019. [1](#), [2](#), [6](#)
- [40] Abhishek Vahadane, Tingying Peng, Amit Sethi, Shadi Albarqouni, Lichao Wang, Maximilian Baust, Katja Steiger, Anna Melissa Schlitter, Irene Esposito, and Nassir Navab. Structure-preserving color normalization and sparse stain separation for histological images. *IEEE transactions on medical imaging*, 35(8):1962–1971, 2016. [1](#), [2](#)
- [41] Jelica Vasiljević, Friedrich Feuerhake, Cédric Wemmert, and Thomas Lampert. Towards histopathological stain invariance by unsupervised domain augmentation using generative adversarial networks. *arXiv preprint arXiv:2012.12413*, 2020. [2](#)
- [42] Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pages 196–202. Springer, 1992. [6](#)
- [43] Gang Xu, Zhigang Song, Zhuo Sun, Calvin Ku, Zhe Yang, Cancheng Liu, Shuhao Wang, Jianpeng Ma, and Wei Xu. Camel: A weakly supervised learning framework for histopathology image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10682–10691, 2019. [5](#)