



Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology

David Tellez^{a,*}, Geert Litjens^a, Péter Bándi^a, Wouter Bulten^a, John-Melle Bokhorst^a,
Francesco Ciompi^a, Jeroen van der Laak^{a,b}

^aDiagnostic Image Analysis Group and the Department of Pathology, Radboud University Medical Center, Nijmegen, the Netherlands

^bCenter for Medical Image Science and Visualization, Linköping University, Linköping, Sweden

ARTICLE INFO

Article history:

Received 18 February 2019

Revised 25 June 2019

Accepted 19 August 2019

Available online 21 August 2019

Keywords:

Deep learning

Convolutional neural network

Computational pathology

ABSTRACT

Stain variation is a phenomenon observed when distinct pathology laboratories stain tissue slides that exhibit similar but not identical color appearance. Due to this color shift between laboratories, convolutional neural networks (CNNs) trained with images from one lab often underperform on unseen images from the other lab. Several techniques have been proposed to reduce the generalization error, mainly grouped into two categories: stain color augmentation and stain color normalization. The former simulates a wide variety of realistic stain variations during training, producing stain-invariant CNNs. The latter aims to match training and test color distributions in order to reduce stain variation. For the first time, we compared some of these techniques and quantified their effect on CNN classification performance using a heterogeneous dataset of hematoxylin and eosin histopathology images from 4 organs and 9 pathology laboratories. Additionally, we propose a novel unsupervised method to perform stain color normalization using a neural network. Based on our experimental results, we provide practical guidelines on how to use stain color augmentation and stain color normalization in future computational pathology applications.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Computational pathology aims at developing machine learning based tools to automate and streamline the analysis of whole-slide images (WSI), i.e. high-definition images of histological tissue sections. These sections consist of thin slices of tissue that are stained with different dyes so that tissue architecture becomes visible under the microscope. In this study, we focus on hematoxylin and eosin (H&E), the most widely used staining worldwide. It highlights cell nuclei in blue color (hematoxylin), and cytoplasm, connective tissue and muscle in various shades of pink (eosin). The eventual color distribution of the WSI depends on multiple steps of the staining process, resulting in slightly different color distributions depending on the laboratory where the sections were processed, see Fig. 1 for examples of H&E stain variation. This inter-center stain variation hampers the performance of machine learning algorithms used for automatic WSI analysis. Algorithms that were trained with images originated from a single pathology laboratory often underperform when applied to images from a dif-

ferent center, including state-of-the-art methods based on convolutional neural networks (CNNs) (Goodfellow et al., 2016; Komura and Ishikawa, 2018; Veta et al., 2019; Sirinukunwattana et al., 2017). Existing solutions to reduce the generalization error in this setting can be categorized into two groups: (1) *stain color augmentation*, and (2) *stain color normalization*.

1.1. Stain color augmentation

Stain color augmentation, and more generally data augmentation, has been proposed as a method to reduce CNN generalization error by simulating realistic variations of the training data. These artificial variations are hand-engineered to mimic the appearance of future test samples that deviate from the training manifold. Previous work on data augmentation for computational pathology has defined two main groups of augmentation techniques: (1) morphological and (2) color transformations (Liu et al., 2017; Tellez et al., 2018). Morphological augmentation spans from simple techniques such as 90° rotations, vertical and horizontal mirroring, or image scaling; to more advanced methods like elastic deformation (Simard et al., 2003), additive Gaussian noise, and Gaussian blurring. The common denominator among these transformations

* Corresponding author.

E-mail address: david.tellezmartin@radboudumc.nl (D. Tellez).




Fig. 1. Example images from training and test datasets. Applications are indicated by colors and keywords: tumor detection in lymph nodes (*lymph*), colorectal cancer tissue classification (*crc*), mitosis detection (*mitosis*) and prostate epithelium detection (*prostate*). Training set images are indicated by the keyword *rumc* and black outline. The rest belong to test sets from other centers. Stain variation can be observed between training and test images.

is the fact that only the morphology of the underlying image is modified and not the color appearance, e.g. Gaussian blurring simulates out of focus artifacts which is a common issue encountered with WSI scanners. Conversely, color augmentation leaves morphological features intact and focuses on simulating stain color variations instead. Common color augmentation techniques borrowed from Computer Vision include brightness, contrast and hue perturbations. Recently, researchers have proposed other approaches more tailored to mimic specific H&E stain variations, e.g. by perturbing the images directly in the H&E color space (Tellez et al., 2018), or perturbing the principal components of the pixel values (Bug et al., 2017).

1.2. Stain color normalization

Stain color normalization reduces stain variation by matching the color distribution of the training and test images. Traditional approaches try to normalize the color space by estimating a color deconvolution matrix that allows identifying the underlying stains (Reinhard et al., 2001; Macenko et al., 2009). More recent methods use machine learning algorithms to detect certain morphological structures, e.g. cell nuclei, that are associated with certain stains, improving the result of the normalization process (Khan et al., 2014; Bejnordi et al., 2016). Deep generative models, i.e. variational autoencoders and generative adversarial networks (Kingma and Welling, 2013; Goodfellow et al., 2014), have been used to generate new image samples that match the template data manifold (Cho et al., 2017; Zanjani et al., 2018). Moreover, color normalization has been formulated as a style transfer task where the style is defined as the color distribution produced by a particular lab (Bug et al., 2017). However, despite their success and widespread adoption as a preprocessing tool in a variety of computational pathology applications (Clarke and Treanor, 2017; Albarqouni et al., 2016; Janowczyk et al., 2017; Bándi et al., 2019), they are not always effective and can produce images with color distributions that deviate from the desired color template. In this study, we propose a novel unsupervised approach that leverages the power of deep learning to solve the problem of stain normalization. We reformulate the problem of stain normalization as an image-to-image translation task and train a neural network to solve it. We do so by feeding the network with heavily augmented H&E images and training the model to reconstruct the original image without augmentation. By learning to remove this color variation, the network effectively learns to perform *stain color normalization* in unseen images whose color distribution deviates from that of the training set.

1.3. Multicenter evaluation

Despite the wide adoption of *stain color augmentation* and *stain color normalization* in the field of computational pathology, the effects on performance of these techniques have not been systematically evaluated. Existing literature focuses on particular applications, and does not quantify the relationship between these techniques and CNN performance (Komura and Ishikawa, 2018; Wang et al., 2015; Zhu et al., 2014; Veta et al., 2019). In this study, we aim to overcome this limitation by comparing these techniques across four representative applications including multicenter data. We selected four patch-based classification tasks where a CNN was trained with data from a single center only, and evaluated in unseen data from multiple external pathology laboratories. We chose four relevant applications from the literature: (1) detecting the presence of mitotic figures in breast tissue (Tellez et al., 2018); (2) detecting the presence of tumor metastases in breast lymph node tissue (Bándi et al., 2019); (3) detecting the presence of epithelial cells in prostate tissue (Bulten et al., 2019); and (4) distinguishing among 9 tissue classes in colorectal cancer (CRC) tissue (Ciompi et al., 2017). All test datasets presented a substantial and challenging stain color deviation from the training set, as can be seen in Fig. 1. We trained a series of CNN classifiers following an identical training protocol while varying the *stain color normalization* and *stain color augmentation* techniques used during training. This thorough evaluation allowed us to establish a ranking among the methods and measure relative performance improvements among them.

1.4. Contributions

Our contributions can be summarized as follows:

- We systematically evaluated several well-known *stain color augmentation* and *stain color normalization* algorithms in order to quantify their effects on CNN classification performance.
- We conducted the previous evaluation using data from a total of 9 different centers spanning 4 relevant classification tasks: mitosis detection, tumor metastasis detection in lymph nodes, prostate epithelium detection, and multiclass colorectal cancer tissue classification.
- We formulated the problem of *stain color normalization* as an unsupervised image-to-image translation task and trained a neural network to solve it.

The paper is organized as follows. Sections 2 and 3 describe the materials and methods thoroughly. Experimental results are

explained in Section 4, followed by Sections 5 and 6 where the discussion and final conclusion are stated.

2. Materials

We collected data from a variety of pathology laboratories for four different applications. In all cases, we used images from the Radboud University Medical Centre (Radboudumc or *rumc*) exclusively to train the models for each of the four classification tasks. Images from the remaining centers were used for testing purposes only. We considered RGB patches of 128x128 pixels extracted from annotated regions. Examples of these patches are shown in Fig. 1. The following sections describe each of the four classification tasks.

2.1. Mitotic figure detection

In this binary classification task, the goal was to accurately classify as positive samples those patches containing a mitotic figure in their center, i.e. a cell undergoing division. In order to train the classifier, we used 14 H&E WSIs from triple negative breast cancer patients, scanned at 0.25µm/pixel resolution, with annotations of mitotic figures obtained as described in Tellez et al. (2018). We split the slides into training (6), validation (4) and test (4), and extracted a total of 1M patches. We refer to this set as *mitosis-rumc*.

For the external dataset, we used publicly available data from the TUPAC Challenge (Veta et al., 2019), i.e. 50 cases of invasive breast cancer with manual annotations of mitotic figures scanned at 0.25µm/pixel resolution. We extracted a total of 300 K patches, and refer to this dataset as *mitosis-tupac*.

2.2. Tumor metastasis detection

The aim of this binary classification task was to identify patches containing metastatic tumor cells. We used publicly available WSIs from the Camelyon17 Challenge (Bándi et al., 2019). This cohort consisted of 50 exhaustively annotated H&E slides of breast lymph node resections from breast cancer patients from 5 different centers (10 slides per center), including Radboudumc. They were scanned at 0.25µm/pixel resolution and the tumor metastases were manually delineated by experts.

We used the 10 WSIs from the Radboudumc to train the classifier, split into training (4), validation (3) and test (3), and extracted a total of 300K patches. We refer to this dataset as *lymph-rumc*. We used the remaining 40 WSIs as external test data, extracting a total of 1.2M patches, and assembling 4 different test sets (one for each center). We named them according to their center's name acronym: *lymph-umcu*, *lymph-cwh*, *lymph-rh* and *lymph-lpe*.

2.3. Prostate epithelium detection

The goal of this binary classification task was to identify patches containing epithelial cells in prostate tissue. We trained the classifier with 25 H&E WSIs of prostate resections from the Radboudumc scanned at 0.5 µm/pixel resolution, with annotations of epithelial tissue as described in Bulten et al. (2019). We split this cohort into training (13), validation (6) and test (6), and extracted a total of 250 K patches. We refer to it as *prostate-rumc*.

We used two test datasets for this task. First, we selected 10 H&E slides of prostate resections from the Radboudumc with different staining and scanning conditions, resulting in substantially different stain appearance (see *prostate-rumc2* in Fig. 1). This test set was manually annotated as described in Bulten et al. (2019) and named *prostate-rumc2*. We extracted 75 K patches from these WSIs. Second, we used publicly available images from 20 H&E slides of prostatectomy specimens with manual annotations of epithelial tissue obtained as described

in Bulten et al. (2019) and Gertych et al. (2015). We extracted 65K patches from them and named the test set *prostate-cedar*.

2.4. Colorectal cancer tissue type classification

In this multiclass classification task, the goal was to distinguish among 9 different colorectal cancer (CRC) tissue classes, namely: (1) tumor, (2) stroma, (3) muscle, (4) lymphocytes, (5) healthy glands, (6) fat, (7) blood cells, (8) necrosis and debris, and (9) mucus. We used 54 H&E WSIs of colorectal carcinoma tissue from the Radboudumc scanned at 0.5 µm/pixel resolution to train the classifier, with manual annotations of the 9 different tissue classes. We split this cohort into training (24), validation (15) and test (15), extracted a total of 450 K patches, and named it *crc-rumc*.

We used two external datasets for this task. First, a set of 74 H&E WSIs from rectal carcinoma patients with annotations of the same 9 classes, as described in Ciompi et al. (2017). We extracted 35K patches and refer to this dataset as *crc-labpon*. Second, we used a publicly available set of H&E image patches from colorectal carcinoma patients (Kather et al., 2016). Annotations for 6 tissue types were available: (1) tumor, (2) stroma, (3) lymph, (4) healthy glands, (5) fat, and (6) blood cells, debris and mucus. We extracted 4K patches in total, and refer to this dataset as *crc-heidelberg*.

2.5. Multi-organ dataset

For the purpose of training a network to solve the problem of *stain color normalization*, we created an auxiliary dataset by aggregating patches from *mitosis-rumc*, *lymph-rumc*, *prostate-rumc* and *crc-rumc* in a randomized and balanced manner. We discarded all labels since they were not needed for this purpose. We preserved a total of 500 K patches for this set and called it the *multi-organ* dataset.

3. Methods

In this study, we evaluated the effect in classification performance of several methods for *stain color augmentation* and *stain color normalization*. This section describes these methods.

3.1. Stain color augmentation

We assume a homogeneous stain color distribution ϕ_{train} for the training images and a more varied color distribution ϕ_{test} for the test images. Note that it is challenging for a classification model trained solely with ϕ_{train} to generalize well to ϕ_{test} due to potential stain differences among sets. To solve this problem, stain color augmentation defines a preprocessing function f that transforms images of the training set to present an alternative and more diverse color distribution ϕ_{augment} :

$$\phi_{\text{train}} \xrightarrow{f} \phi_{\text{augment}} \quad (1)$$

on the condition that:

$$(\phi_{\text{augment}} \supseteq \phi_{\text{train}}) \wedge (\phi_{\text{augment}} \supseteq \phi_{\text{test}}) \quad (2)$$

In practice, heavy data augmentation is used to satisfy Eq. (2). In order to simplify our experimental setup, we grouped several data augmentation techniques into the following categories attending to the nature of the image transformations. Examples of the resulting augmented images are shown in Fig. 2.

Basic. This group included 90 ° rotations, and vertical and horizontal mirroring.

Morphology. We extended *basic* with several transformations that simulate morphological perturbations, i.e. alterations in shape, texture or size of the imaged tissue structures, including scanning




Fig. 2. Summary of the data augmentation techniques and datasets used in this study, organized in columns and rows respectively. Patches on the leftmost column depict the original input images and the rest of patches are augmented versions of them. Augmentations performed in the *grayscale* color space are depicted on the right for one sample dataset only. *Basic* augmentation is included in all cases.

artifacts. We included *basic* augmentation, scaling, elastic deformation (Simard et al., 2003), additive Gaussian noise (perturbing the signal-to-noise ratio), and Gaussian blurring (simulating out-of-focus artifacts).

Brightness & contrast (BC). We extended *morphology* with random brightness and contrast image perturbations (Haeberli and Voorhies, 1994).

Hue-Saturation-Value (HSV). We extended the *BC* augmentation by randomly shifting the hue and saturation channels in the HSV color space (Van der Walt et al., 2014). This transformation produced substantially different color distributions when applied to the training images. We tested two configurations depending on the observed color variation strength, called *HSV-light* and *HSV-strong*.

Hematoxylin-Eosin-DAB (HED). We extended the *BC* augmentation with a color variation routine specifically designed for H&E images (Tellez et al., 2018). This method followed three steps. First, it disentangled the hematoxylin and eosin color channels by means of color deconvolution using a fixed matrix. Second, it perturbed the hematoxylin and eosin stains independently. Third, it transformed the resulting stains into regular RGB color space. We tested two configurations depending on the observed color variation strength, called *HED-light* and *HED-strong*.

During training, we selected the value of the augmentation hyper-parameters randomly within certain ranges to achieve stain variation. We tuned all ranges manually via visual examination. In particular, we used a scaling factor between $[0.8, 1.2]$, elastic deformation parameters $\alpha \in [80, 120]$ and $\sigma \in [9.0, 11.0]$, additive Gaussian noise with $\sigma \in [0, 0.1]$, Gaussian blurring with $\sigma \in [0, 0.1]$, brightness intensity ratio between $[0.65, 1.35]$, and contrast intensity ratio between $[0.5, 1.5]$. For *HSV-light* and *HSV-strong*, we used hue and saturation intensity ratios between $[-0.1, 0.1]$ and $[-1, 1]$, respectively. For *HED-light* and *HED-strong*, we used intensity ratios between $[-0.05, 0.05]$ and $[-0.2, 0.2]$, respectively, for all HED channels.

3.2. Stain color normalization

Stain color normalization reduces color variation by transforming the color distribution of training and test images, i.e. ϕ_{train} and ϕ_{test} , to that of a template ϕ_{normal} . It performs such transformation using a normalization function g that maps any given color distribution to the template one:

$$(\phi_{\text{train}} \xrightarrow{g} \phi_{\text{normal}}) \wedge (\phi_{\text{test}} \xrightarrow{g} \phi_{\text{normal}}) \quad (3)$$

By matching ϕ_{train} and ϕ_{test} , the problem of stain variance vanishes and the model no longer requires to generalize to unseen stains in order to perform well. We evaluated several methods that implement g (see Fig. 3), and propose a novel technique based on neural networks.

Identity. We performed no transformation on the input patches, serving as a baseline method for the rest of techniques.

Grayscale. In this case, g transformed images from RGB to grayscale space, removing most of the color information present in the patches. We hypothesized that this color information is redundant since most of the signal in H&E images is present in morphological and structural patterns, e.g. the presence of a certain type of cell.

Deconv-based. We followed the color deconvolution approach proposed by Macenko et al. (2009). This method assumes that the hematoxylin and eosin stains are linearly separable in the optical density (OD) color space, as opposed to RGB space. This method finds the two largest singular value directions using singular value decomposition, and projects the OD pixel values onto this plane. This procedure allows to identify the underlying hematoxylin and eosin stain vectors, and use them to perform color deconvolution on a given image to decompose the RGB image into its normalized hematoxylin and eosin components.

LUT-based. We implemented an approach that uses tissue morphology to perform stain color normalization (Bejnordi et al., 2016). This popular method has been used by numerous re-




Fig. 3. Visual comparison of the stain color normalization techniques used in this study. Rows correspond to the different tested techniques and columns to datasets, with green for lymph, blue for mitosis, yellow for prostate and red for colorectal. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Network-based stain color normalization. From left to right: patches from the training set are transformed with heavy color augmentation and fed to a neural network. This network is trained to reconstruct the original appearance of the input images by removing color augmentation, effectively learning how to perform stain color normalization.

searchers in recent public challenges (Bándi et al., 2019; Bejnordi et al., 2017). It detects cell nuclei in order to precisely characterize the H&E chromatic distribution and density histogram for a given WSI. First, it does so for a given template WSI, e.g. an image from the training set, and a target WSI. Second, the color distributions of the template and target WSIs are matched, and the color correspondence is stored in a look-up table (LUT). Finally, this LUT is used to normalize the color of the target WSI.

Style-based. Bug et al. (2017) proposed to use a neural network to perform stain color normalization based on the idea of style transfer. They transform the color distribution of RGB images by using feature-aware normalization, a mechanism that shifts and scales intermediate feature maps based on features extracted from the input image. This feature extractor is an ImageNet (Deng et al., 2009) pre-trained network, while the rest of the model is trained to reconstruct PCA-augmented histopathology images. We used the authors' implementation of the method and retrained the model using images from the *multi-organ* dataset.

Network-based. We developed a novel approach to perform stain color normalization based on unsupervised learning and neural networks (see Fig. 4). We parameterized the normalization function g with a neural network G and trained it end-to-end to remove the effects of data augmentation. Even though it is not possible to invert the many-to-many augmentation function f , we can learn a partial many-to-one function that maps any arbitrary color distribution ϕ_{augment} to a template distribution ϕ_{normal} :

$$\phi_{\text{augment}} \xrightarrow{G} \phi_{\text{normal}} \quad (4)$$

Since G can normalize ϕ_{augment} (Eq. (4)), and ϕ_{augment} is a superset of ϕ_{train} and ϕ_{test} (Eq. (2)), we conclude that G can effectively normalize ϕ_{train} and ϕ_{test} (Eq. (2)).

We trained G to perform image-to-image translation using the *multi-organ* dataset. During training, images were heavily augmented and fed to the network. The model was tasked with reconstructing the images with their original appearance, before augmentation. We used a special configuration of the *HSV* augmentation where we kept the color transformation only, i.e. did not include *basic*, *morphology* and *BC*. We used the maximum intensity for the transformation hyper-parameters, i.e. hue, saturation and value channel ratios between $[-1, 1]$. The strength of this augmentation resulted in images with drastically different color distributions, sometimes compressing all color information into grayscale. In order to invert this complex augmentation, we hypothesized that the network learned to associate certain tissue structures with their usual color appearance.

We used an architecture inspired by U-Net (Ronneberger et al., 2015), with a downward path of 5 layers of strided convolutions (Springenberg et al., 2014) with 32, 64, 128, 256 and 512 3×3 filters, stride of 2, batch normalization (BN) (Ioffe and Szegedy, 2015) and leaky-ReLU activation (LRA) (Maas et al., 2013). The upward path consisted of 5 upsampling layers, each one composed of a pair of nearest-neighbor upsampling and a convolutional operation (Odena et al., 2016), with 256, 128, 64, 32 and 3 3×3 filters, BN and LRA; except for the final convolutional layer that did not have BN and used the hyperbolic tangent (tanh) as activation function. We used long skip connections to ease the synthesis upward path (Ronneberger et al., 2015), and applied L2 regularization with a factor of 1×10^{-6} .

We minimized the mean squared error (MSE) loss using stochastic gradient descent with Adam optimization (Kingma and Ba, 2014) and 64-sample mini-batch, decreasing the learning rate by a factor of 10 starting from 1×10^{-2} every time the validation loss stopped improving for 4 consecutive epochs until 1×10^{-5} . Finally, we selected the weights corresponding to the model with the lowest validation loss during training.

Convergence to average solutions is a known effect with bottleneck architectures trained with MSE loss. Note, however, that our network-based normalization architecture includes long skip connections between the downward and the upward paths. These skip connections allow the model to copy spatial structures from the in-




Fig. 5. Constellations of internal and external datasets analyzed in this work. Each data point represents the mean and standard deviation pixel intensity of all image patches in a particular dataset in the HSV color space (hue and saturation in the x and y axis, respectively). Note how normalization methods tend to cluster the color distribution of the datasets, whereas color augmentation does the opposite. Color augmentation plot (bottom-right): patches from internal images are transformed with different color augmentation methods (grey points representing the original internal and external datasets are shown as reference).

put images to the output images with ease, and utilize the rest of the model to modify style-related color features. Since there is no bottleneck effect, i.e., the model has all the information necessary to reconstruct the input image, image reconstructions are highly accurate and do not show any blurriness in practice.

3.3. Color analysis

In order to understand how *stain color augmentation* and *stain color normalization* influenced the color differences between internal (*rumc*) and external datasets (*rest*), we analyzed the image patches in the HSV color space. We measured the mean and standard deviation pixel intensity along the hue and saturation dimensions, and plotted the results in a 2D plane, comparing images processed with the color normalization and augmentation techniques analyzed in this work (see Fig. 5). We confirmed the clustering effect of normalization algorithms, and the scattering effect of augmentation methods.

3.4. CNN classifiers

In order to measure the effect of *stain color augmentation* and *stain color normalization*, we trained a series of identical CNN classifiers to perform patch classification using different combinations of these techniques. For training and validation purposes, we used the *rumc* datasets described in Section 2.

The architecture of such CNN classifiers consisted of 9 layers of strided convolutions with 32, 64, 64, 128, 128, 256, 256, 512 and 512 3x3 filters, stride of 2 in the even layers, BN and LRA; followed by global average pooling; 50% dropout; a dense layer with 512 units, BN and LRA; and a linear dense layer with either 2 or 9 units depending on the classification task, followed by a softmax. We applied L2 regularization with a factor of 1×10^{-6} .

We minimized the cross-entropy loss using stochastic gradient descent with Adam optimization and 64-sample class-balanced mini-batch, decreasing the learning rate by a factor of 10 starting from 1×10^{-2} every time the validation loss stopped improving for 4 consecutive epochs until 1×10^{-5} . Finally, we selected the

weights corresponding to the model with the lowest validation loss during training.

4. Experimental results

We conducted a series of experiments in order to quantify the impact in performance of the different *stain color augmentation* and *stain color normalization* methods introduced in the previous section across four different classification tasks. We trained a CNN classifier for each combination of organ, color normalization and data augmentation method under consideration. In the case of *grayscale* normalization, we only tested *basic*, *morphology* and *BC* augmentation techniques. We conducted 152 different experiments, repeating each 5 times using different random initialization for the network parameters, accounting for a total of 760 trained CNN classifiers.

4.1. Evaluation

We evaluated the area under the receiver-operating characteristic curve (AUC) of each CNN in each external test set. In the case of multiclass classification, we considered the unweighted average, i.e. we calculated the individual AUC per label (one-vs-all) and averaged the resulting values. We reported the mean and standard deviation of the resulting AUC for each experiment across five repetitions in Table 1.

In order to establish a global ranking among methods, shown in the rightmost column in Table 1, we performed the following calculation. We converted the AUC scores into ranking scores per test set column, and averaged these scores along the dataset dimension to obtain a global ranking score per method. Note that we performed an average across ranking scores, rather than AUC scores, following established procedures (Demšar, 2006). Data in Table 1 and the raw AUC scores are provided in machine-readable format as Supplementary Material to this article.

Table 1
Experimental results ranking *stain color augmentation* and *stain color normalization* methods. Values correspond to AUC scores, except for the last column, averaged across 5 repetitions with standard deviation shown between parenthesis. Each column represents a different external test dataset, with the last column *Ranking* indicating the position of each method within the global benchmark, computed as described in Section 4.1. Normalization methods: *Network* is our proposal; *Style* is from Bug et al. (2017); *LUT* is from Bejnordi et al. (2016); and *Deconvolution* is from Macenko et al. (2009).

Normalization	Augmentation	lymph-cwh	lymph-lpe	lymph-rh	lymph-umcu	mitosis-tupac	prostate-rumc2	prostate-cedar	crc-labpon	crc-heidelberg	Ranking
Identity	HED-light	0.952(0.004)	0.976(0.001)	0.946(0.009)	0.968(0.004)	0.996(0.001)	0.957(0.001)	0.879(0.011)	0.973(0.002)	0.895(0.002)	1.2(0.4)
Style	HED-light	0.961(0.002)	0.953(0.004)	0.952(0.001)	0.972(0.004)	0.991(0.003)	0.925(0.003)	0.879(0.006)	0.975(0.001)	0.917(0.003)	2.8(0.7)
Network	HSV-light	0.946(0.006)	0.962(0.001)	0.941(0.002)	0.965(0.004)	0.992(0.001)	0.957(0.000)	0.872(0.013)	0.980(0.001)	0.900(0.003)	3.9(1.9)
Network	HED-light	0.949(0.005)	0.968(0.001)	0.942(0.002)	0.963(0.004)	0.989(0.003)	0.958(0.001)	0.862(0.011)	0.980(0.001)	0.906(0.003)	4.1(1.6)
Identity	HSV-strong	0.955(0.003)	0.965(0.004)	0.929(0.002)	0.973(0.003)	0.988(0.003)	0.945(0.009)	0.886(0.005)	0.977(0.001)	0.902(0.003)	4.7(1.7)
Network	HSV-strong	0.953(0.002)	0.964(0.003)	0.946(0.002)	0.964(0.005)	0.991(0.003)	0.951(0.002)	0.852(0.006)	0.975(0.002)	0.894(0.005)	6.6(0.9)
Network	HED-strong	0.956(0.003)	0.959(0.002)	0.940(0.003)	0.965(0.004)	0.985(0.005)	0.943(0.003)	0.861(0.009)	0.974(0.002)	0.916(0.003)	7.9(1.9)
Identity	HED-strong	0.950(0.005)	0.959(0.005)	0.936(0.007)	0.957(0.007)	0.992(0.002)	0.945(0.003)	0.872(0.005)	0.967(0.003)	0.920(0.005)	8.1(2.8)
Style	HSV-strong	0.953(0.004)	0.956(0.004)	0.940(0.003)	0.959(0.007)	0.986(0.004)	0.932(0.005)	0.878(0.003)	0.976(0.001)	0.917(0.004)	9.0(2.6)
Style	HSV-light	0.940(0.011)	0.960(0.004)	0.944(0.007)	0.926(0.012)	0.992(0.001)	0.958(0.002)	0.852(0.008)	0.974(0.001)	0.921(0.003)	9.4(3.6)
Style	HED-strong	0.955(0.002)	0.949(0.004)	0.936(0.003)	0.954(0.005)	0.982(0.005)	0.942(0.002)	0.884(0.004)	0.975(0.000)	0.925(0.003)	9.9(1.2)
Grayscale	BC	0.956(0.003)	0.962(0.003)	0.935(0.005)	0.961(0.002)	0.989(0.002)	0.939(0.004)	0.851(0.002)	0.972(0.000)	0.884(0.003)	12.2(1.2)
Deconvolution	HSV-strong	0.955(0.003)	0.936(0.008)	0.941(0.004)	0.943(0.009)	0.991(0.001)	0.865(0.010)	0.867(0.004)	0.961(0.002)	0.928(0.001)	13.9(1.9)
LUT	HED-strong	0.934(0.006)	0.941(0.006)	0.925(0.006)	0.963(0.005)	0.989(0.002)	0.945(0.002)	0.871(0.005)	0.956(0.002)	0.945(0.001)	14.0(2.2)
Deconvolution	HED-strong	0.942(0.003)	0.962(0.003)	0.897(0.006)	0.967(0.003)	0.993(0.002)	0.827(0.018)	0.853(0.006)	0.969(0.001)	0.927(0.002)	14.4(1.2)
LUT	HSV-strong	0.923(0.009)	0.939(0.003)	0.928(0.005)	0.947(0.008)	0.987(0.002)	0.949(0.003)	0.862(0.007)	0.962(0.002)	0.940(0.002)	17.0(2.0)
Network	BC	0.944(0.003)	0.950(0.003)	0.903(0.003)	0.934(0.006)	0.983(0.005)	0.953(0.003)	0.869(0.009)	0.981(0.001)	0.881(0.005)	17.4(1.5)
Identity	HSV-light	0.888(0.013)	0.951(0.009)	0.942(0.004)	0.930(0.023)	0.962(0.015)	0.949(0.001)	0.905(0.005)	0.976(0.000)	0.894(0.003)	17.4(2.9)
LUT	HED-light	0.914(0.011)	0.926(0.011)	0.923(0.006)	0.932(0.019)	0.993(0.001)	0.948(0.003)	0.852(0.021)	0.966(0.003)	0.940(0.002)	17.9(2.2)
LUT	HSV-light	0.894(0.006)	0.936(0.006)	0.921(0.003)	0.942(0.007)	0.987(0.002)	0.951(0.002)	0.860(0.010)	0.971(0.002)	0.945(0.002)	19.2(1.2)
LUT	BC	0.925(0.025)	0.948(0.027)	0.853(0.016)	0.790(0.061)	0.985(0.004)	0.951(0.004)	0.848(0.018)	0.973(0.003)	0.924(0.005)	21.3(3.3)
Style	BC	0.949(0.005)	0.858(0.031)	0.938(0.001)	0.411(0.065)	0.987(0.004)	0.949(0.006)	0.764(0.047)	0.946(0.002)	0.903(0.005)	23.3(2.2)
Deconvolution	HSV-light	0.942(0.004)	0.930(0.009)	0.913(0.023)	0.961(0.002)	0.982(0.005)	0.850(0.019)	0.840(0.009)	0.958(0.006)	0.917(0.002)	23.5(1.0)
Network	Basic	0.944(0.003)	0.954(0.007)	0.887(0.010)	0.959(0.004)	0.969(0.005)	0.905(0.006)	0.815(0.019)	0.977(0.002)	0.855(0.006)	23.6(1.4)
Network	Morphology	0.939(0.010)	0.949(0.006)	0.890(0.012)	0.950(0.009)	0.980(0.006)	0.913(0.011)	0.823(0.022)	0.977(0.001)	0.868(0.002)	23.9(1.1)
Deconvolution	HED-light	0.930(0.005)	0.912(0.015)	0.916(0.005)	0.948(0.006)	0.982(0.002)	0.816(0.011)	0.834(0.004)	0.970(0.003)	0.927(0.005)	25.4(2.3)
Deconvolution	Morphology	0.951(0.003)	0.938(0.006)	0.849(0.021)	0.951(0.008)	0.993(0.002)	0.754(0.027)	0.749(0.037)	0.903(0.008)	0.865(0.015)	27.7(0.6)
Grayscale	Morphology	0.943(0.010)	0.820(0.021)	0.922(0.005)	0.941(0.011)	0.991(0.006)	0.910(0.009)	0.816(0.005)	0.929(0.006)	0.813(0.009)	27.7(1.2)
Style	Morphology	0.935(0.011)	0.725(0.082)	0.934(0.002)	0.361(0.113)	0.992(0.004)	0.918(0.006)	0.754(0.006)	0.873(0.010)	0.890(0.006)	28.6(2.4)
Grayscale	Basic	0.940(0.007)	0.692(0.064)	0.926(0.010)	0.938(0.019)	0.992(0.001)	0.882(0.008)	0.661(0.039)	0.934(0.002)	0.798(0.006)	30.0(0.6)
Deconvolution	BC	0.942(0.004)	0.896(0.008)	0.682(0.044)	0.949(0.005)	0.989(0.006)	0.794(0.021)	0.792(0.028)	0.930(0.007)	0.872(0.004)	30.4(1.2)
LUT	Morphology	0.898(0.007)	0.920(0.007)	0.801(0.021)	0.874(0.025)	0.969(0.008)	0.895(0.013)	0.803(0.007)	0.939(0.006)	0.906(0.006)	32.6(1.4)
Deconvolution	Basic	0.919(0.015)	0.896(0.038)	0.810(0.081)	0.902(0.026)	0.993(0.001)	0.753(0.006)	0.791(0.009)	0.903(0.003)	0.836(0.008)	32.8(0.7)
Style	Basic	0.918(0.002)	0.334(0.133)	0.926(0.004)	0.124(0.041)	0.991(0.003)	0.865(0.025)	0.723(0.024)	0.863(0.020)	0.857(0.010)	33.6(1.0)
LUT	Basic	0.908(0.010)	0.894(0.030)	0.809(0.022)	0.772(0.072)	0.951(0.009)	0.906(0.011)	0.741(0.018)	0.930(0.014)	0.890(0.013)	34.6(0.8)
Identity	BC	0.899(0.006)	0.634(0.100)	0.741(0.016)	0.177(0.047)	0.906(0.034)	0.936(0.006)	0.704(0.060)	0.684(0.009)	0.761(0.012)	36.2(0.4)
Identity	Morphology	0.811(0.026)	0.671(0.099)	0.673(0.027)	0.214(0.174)	0.986(0.006)	0.374(0.191)	0.602(0.023)	0.569(0.028)	0.720(0.009)	37.2(0.7)
Identity	Basic	0.811(0.009)	0.563(0.309)	0.790(0.047)	0.406(0.375)	0.965(0.009)	0.631(0.178)	0.624(0.053)	0.556(0.057)	0.701(0.028)	37.6(0.5)

4.2. Effects of stain color augmentation

Results in Table 1 show that *stain color augmentation* was crucial to obtain top classification performance, regardless of the *stain color normalization* technique used (see top-10 methods). Moreover, note that including color augmentation, either *HSV* or *HED*, was key to obtaining top performance since using *BC* augmentation alone produced mediocre results. We did not find, however, any substantial performance difference between using *HED* or *HSV* color augmentation. Similarly, we found that *strong* and *light* color augmentations achieved similar performance, with a slight advantage towards *light*. Heavy augmentation is known to reduce performance on images similar to those in the training set. However, we found less than 1% average performance reduction on the internal test set across organs. Regarding non-color augmentation, i.e. *basic*, *morphology* and *BC*, *BC* obtained the best results across almost all *stain color normalization* setups, followed by *morphology* and *basic* augmentation, as expected.

4.3. Effects of stain color normalization

According to results in Table 1, overall top performance was achieved without the use of color normalization. This piece of evidence suggests that color normalization is not a necessary condition to achieve high classification performance in histopathology images. However, we observed that color normalization generally produced classifiers that were more robust to different color augmentation techniques, e.g., *Identity* normalization performance diminished with *HSV-light* augmentation whereas *Network* normalization exhibited a high performance regardless of the color augmentation used.

We did not find any substantial performance difference between neural network based normalization algorithms, *Network* and *Style*. Nevertheless, we observed that none of the classical approaches, *LUT* or *Deconvolution*, surpassed the performance of *Grayscale*. We hypothesize that these classical normalization methods can hide certain useful features from the images, resulting in added input noise that can affect classification performance.

Additionally, we measured the extra time required to normalize a regular whole-slide image composed of $50,000 \times 50,000$ RGB pixels. We found *LUT-based* to be the fastest taking 21.8 min, followed closely by *network-based* with 26.0 min, and the slower *deconv-based* and *style-based* taking 111.2min and 217.8min, respectively, excluding I/O delays.

5. Discussion

Our experimental results indicate that *stain color augmentation* improved classification performance drastically by increasing the CNN's ability to generalize to unseen stain variations. This was true for most of the experiments regardless of the type of *stain color normalization* technique used. Moreover, we found *HSV* and *HED* color transformations to be the key ingredients to improve performance since removing them, i.e. using *BC* augmentation, yielded a lower AUC under all circumstances; suggesting that inter-lab stain differences were mainly caused by color variations rather than morphological features. Remarkably, we observed hardly any performance difference between *HSV* or *HED*, and *strong* or *light* variation intensity.

Based on these observations, we concluded that CNNs are mostly insensitive to the type and intensity of the color augmentation used in this setup, as long as one of the methods is used. However, CNNs trained with simpler *stain color normalization* techniques exhibited more sensitivity to the intensity of color augmentation, i.e. they required a stronger augmentation in order to perform well. Finally, the fact that experiments with *grayscale* images

achieved mediocre performance was an indication that color provided useful information to the model. The worst performance was achieved with *morphology* and *identity* configurations, which was an indication that color information can act as noise when no augmentation is used, increasing overfitting and generalization error due to stain variation.

Regarding *stain color normalization*, we found that the best performing method did not use any normalization. This result challenged the common assumption that color normalization is a necessary step to achieve top classification performance in the histopathology setting; especially considering that color normalization added a computational overhead that can substantially reduce the overall classification speed. Neural network based methods, both *Network* and *Style*, achieved similar high performance on the benchmark, supporting the idea of reformulating the problem of *stain color normalization* as an image-to-image translation task.

Furthermore, we observed that all *stain color normalization* techniques obtained a poor performance when no color augmentation was used (below that of *Grayscale* with *BC*). We hypothesize that even in the case of excellent stain normalization, color information can serve as a source of overfitting, worsening with suboptimal normalization. We concluded that using the *stain color normalization* methods evaluated in this paper without proper *stain color augmentation* is insufficient to reduce the generalization error caused by stain variation and results in poor model performance.

Due to computational constraints, we limited the type and number of experiments performed in this study to patch-based classification tasks, ignoring other modalities such as segmentation, instance detection or WSI classification. However, we believe this limitation to have little impact in the conclusions of this study since the problem of generalization error has identical causes and effects in other modalities. In order to reduce the number of experiments, we avoided quantifying the impact of individual augmentation techniques, e.g. scaling augmentation alone, but grouped them into categories instead. Similarly, we limited the hyper-parameters' ranges to certain set of values, e.g. *light* or *strong* stain augmentation intensity. Nevertheless, according to the experimental results, we believe that testing a wider range of hyper-parameter values would not alter the main conclusions of this study.

6. Conclusion

For the first time, we quantified the effect of *stain color augmentation* and *stain color normalization* in classification performance across four relevant computational pathology applications using data from 9 different centers. Based on our empirical evaluation, we found that any type of *stain color augmentation*, i.e. *HSV* or *HED* transformation, should always be used. In addition, color augmentation can be combined with neural network based *stain color normalization* to achieve a more robust classification performance. In setups with reduced computational resources, color normalization could be omitted, resulting in a negligible performance reduction and a substantial improvement in processing speed. Finally, we recommend tuning the intensity of the color augmentation to *light* or *strong* in case color normalization is *enabled* or *disabled*, respectively.

Declaration of Competing Interest

All authors have participated in (a) conception and design, or analysis and interpretation of the data; (b) drafting the article or revising it critically for important intellectual content; and (c) approval of the final version.

This manuscript has not been submitted to, nor is under review at, another journal or other publishing venue.

The following authors have affiliations with organizations with direct or indirect financial interest in the subject matter discussed in the manuscript:

Acknowledgments

This study was supported by a Junior Researcher grant from the Radboud Institute of Health Sciences (RIHS), Nijmegen, The Netherlands; a grant from the Dutch Cancer Society (KUN 2015–7970); and another grant from the Dutch Cancer Society and the Alpe d'HuZes fund (KUN 2014–7032); this project has also received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825292. The authors would like to thank Dr. Babak Ehteshami Bejnordi for providing the code for the *LUT-based* stain color normalization algorithm; and Canisius-Wilhelmina Ziekenhuis Nijmegen, Laboratorium Pathologie Oost Nederland, University Medical Center Utrecht, and Rijnstate Hospital Arnhem for kindly providing tissue sections for this study.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.media.2019.101544](https://doi.org/10.1016/j.media.2019.101544).

References

- Albarqouni, S., Baur, C., Achilles, F., Belagiannis, V., Demirci, S., Navab, N., 2016. Agnet: deep learning from crowds for mitosis detection in breast cancer histology images. *IEEE Trans. Med. Imaging* 35 (5), 1313–1321.
- Bándi, P., Geessink, O., Manson, Q., van Dijk, M., Balkenhol, M., Hermesen, M., Bejnordi, B.E., Lee, B., Paeng, K., Zhong, A., et al., 2019. From detection of individual metastases to classification of lymph node status at the patient level: the CAMELYON17 challenge. *IEEE Trans. Med. Imaging* 38 (2), 550–560.
- Bejnordi, B.E., Litjens, G., Timofeeva, N., Otte-Höller, I., Homeyer, A., Karssemeijer, N., van der Laak, J.A., 2016. Stain specific standardization of whole-slide histopathological images. *IEEE Trans. Med. Imaging* 35 (2), 404–415.
- Bejnordi, B.E., Veta, M., Van Diest, P.J., Van Ginneken, B., Karssemeijer, N., Litjens, G., Van Der Laak, J.A., Hermesen, M., Manson, Q.F., Balkenhol, M., et al., 2017. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* 318 (22), 2199–2210.
- Bug, D., Schneider, S., Grote, A., Oswald, E., Feuerhake, F., Schüller, J., Merhof, D., 2017. Context-based normalization of histological stains using deep convolutional features. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, pp. 135–142.
- Bulten, W., Bándi, P., Hoven, J., van de Loo, R., Lotz, J., Weiss, N., van der Laak, J., van Ginneken, B., Hulsbergen-van de Kaa, C., Litjens, G., 2019. Epithelium segmentation using deep learning in h&e-stained prostate specimens with immunohistochemistry as reference standard. *Sci. Rep.* 9 (1), 864.
- Cho, H., Lim, S., Choi, G., Min, H., 2017. Neural stain-style transfer learning using GAN for histopathological images. In: *Proceedings of the Asian Conference on Machine Learning*.
- Ciampi, F., Geessink, O., Bejnordi, B.E., de Souza, G.S., Baidoshvili, A., Litjens, G., van Ginneken, B., Nagtegaal, I., van der Laak, J., 2017. The importance of stain normalization in colorectal tissue classification with convolutional networks. In: *Proceedings of the IEEE 14th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 160–163.
- Clarke, E.L., Treanor, D., 2017. Colour in digital pathology: a review. *Histopathology* 70 (2), 153–163.
- Demšar, J., 2006. Statistical comparisons of classifiers over multiple data sets. *J. Mach. Learn. Res.* 7 (Jan), 1–30.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition CVPR*. IEEE, pp. 248–255.
- Gertych, A., Ing, N., Ma, Z., Fuchs, T.J., Salman, S., Mohanty, S., Bhele, S., Velásquez-Vacca, A., Amin, M.B., Knudsen, B.S., 2015. Machine learning approaches to analyze histological images of tissues from radical prostatectomies. *Comput. Med. Imaging Graph.* 46, 197–208.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. In: *Proceedings of the Advances in Neural Information Processing Systems*, pp. 2672–2680.
- Goodfellow, I., et al., 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Haeberli, P., Voorhies, D., 1994. Image processing by linear interpolation and extrapolation. *IRIS Uni. Mag.* 28, 8–9.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proceedings of the International Conference on Machine Learning*, pp. 448–456.
- Janowczyk, A., Basavanthally, A., Madabhushi, A., 2017. Stain normalization using sparse autoencoders (stanoSA): application to digital pathology. *Comput. Med. Imaging Graph.* 57, 50–61.
- Kather, J.N., Weis, C.-A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G., 2016. Multi-class texture analysis in colorectal cancer histology. *Sci. Rep.* 6, 27988.
- Khan, A.M., et al., 2014. A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution. *IEEE Trans. Biomed. Eng.* 61 (6), 1729–1738.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. In: *Proceedings of the International Conference on Learning Representations*.
- Kingma, D.P., Welling, M., 2013. Auto-encoding variational Bayes. In: *Proceedings of the International Conference on Learning Representations*.
- Komura, D., Ishikawa, S., 2018. Machine learning methods for histopathological image analysis. *Comput. Struct. Biotechnol. J.* 16, 34–42.
- Liu, Y., Gadepalli, K., Norouzi, M., Dahl, G. E., Kohlberger, T., Boyko, A., Venugopalan, S., Timofeev, A., Nelson, P. Q., Corrado, G. S., et al., 2017. Detecting cancer metastases on Gigapixel pathology images, arXiv:1703.02442.
- Maas, A.L., Hannun, A.Y., Ng, A.Y., 2013. Rectifier nonlinearities improve neural network acoustic models. In: *Proceedings of the International Conference on Machine Learning*.
- Macenko, M., et al., 2009. A method for normalizing histology slides for quantitative analysis. In: *Proceedings of the IEEE International Symposium on Biomedical Imaging: From Nano to Macro/ISBI'09*. IEEE, pp. 1107–1110.
- Odena, A., Dumoulin, V., Olah, C., 2016. Deconvolution and checkerboard artifacts. *Distill* 1 (10).
- Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P., 2001. Color transfer between images. *IEEE Comput. Graph. Appl.* 21 (5), 34–41.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 234–241.
- Simard, P.Y., Steinkraus, D., Platt, J.C., et al., 2003. Best practices for convolutional neural networks applied to visual document analysis. In: *Proceedings of the International Conference on Document Analysis and Recognition*, 3. IEEE, pp. 958–962.
- Sirinukunwattana, K., et al., 2017. Gland segmentation in colon histology images: the Glas challenge contest. *Med. Image Anal.* 35, 489–502.
- Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M., 2014. Striving for simplicity: The all convolutional net. In: *Proceedings of the International Conference on Learning Representations*.
- Tellez, D., Balkenhol, M., Otte-Höller, I., van de Loo, R., Vogels, R., Bult, P., Wauters, C., Vreuls, W., Mol, S., Karssemeijer, N., et al., 2018. Whole-Slide mitosis detection in h&e breast histology using PHH3 as a reference to train distilled stain-Invariant convolutional networks. *IEEE Trans. Med. Imaging* 37 (9), 2126–2136.
- Veta, M., Heng, Y.J., Stathonikos, N., Bejnordi, B.E., Beca, F., Wollmann, T., Rohr, K., Shah, M.A., Wang, D., Rousson, M., et al., 2019. Predicting breast tumor proliferation from whole-slide images: the TUPAC16 challenge. *Med. Image Anal.* 54, 111–121.
- Van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T., 2014. Scikit-image: image processing in python. *PeerJ* 2.
- Wang, D., Foran, D.J., Ren, J., Zhong, H., Kim, I.Y., Qi, X., 2015. Exploring automatic prostate histopathology image Gleason grading via local structure modeling. In: *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, pp. 2649–2652.
- Zanjani, F.G., Zinger, S., Bejnordi, B.E., van der Laak, J.A., de With, P.H.N., 2018. Stain normalization of histopathology images using generative adversarial networks. In: *Proceedings of the International Symposium on Biomedical Imaging*. IEEE, pp. 573–577.
- Zhu, Y., Zhang, S., Liu, W., Metaxas, D.N., 2014. Scalable histopathological image analysis via active learning. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 369–376.