

STRATEGIES FOR TRAINING STAIN INVARIANT CNNs

Thomas Lampert* Odysée Merveille* Jessica Schmitz[◇]
Germain Forestier[†] Friedrich Feuerhake^{◇,‡} Cédric Wemmert*

*ICube, Université de Strasbourg, France [◇]Institute of Pathology, Hannover Medical School, Germany
[†]IRIMAS, Université de Haute Alsace, Mulhouse, France [‡]University Clinic, Freiburg, Germany

ABSTRACT

An important part of Digital Pathology is the analysis of multiple digitised whole slide images from differently stained tissue sections. It is common practice to mount consecutive sections containing corresponding microscopic structures on glass slides, and to stain them differently to highlight specific tissue components. These multiple staining modalities result in very different images but include a significant amount of consistent image information. Deep learning approaches have recently been proposed to analyse these images in order to automatically identify objects of interest for pathologists. These supervised approaches require a vast amount of annotations, which are difficult and expensive to acquire—a problem that is multiplied with multiple stainings. This article presents several training strategies that make progress towards stain invariant networks. By training the network on one commonly used staining modality and applying it to images that include corresponding but differently stained tissue structures, the presented unsupervised strategies demonstrate significant improvements over standard training strategies.

Index Terms—multi-stain analysis, digital pathology, whole slide images, deep convolutional networks

1. INTRODUCTION

The field of digital pathology emerged with the introduction of whole slide imaging (WSI) scanners and led to the development of new tools for analysing histopathology slides [1]. Meanwhile, deep learning architectures were developed that outperformed most state-of-the-art algorithms in image analysis. Deep learning has thus been quickly applied to solve digital pathology problems [2, 3, 4]. Nonetheless supervised deep learning approaches require large amounts of annotated data. Unlike classic computer vision applications, annotating histopathology slides requires in-depth knowledge and therefore can only be done by trained experts such as pathologists. Consequently, the acquisition of large annotated histopathological datasets is a complicated task [5].

Image datasets in digital pathology applications often consist of consecutive slides stained differently, each staining providing specific information on the same region of interest (see Figure 1, first row). Even though differently stained slides appear very different, there is often a significant amount of consistent information between them. For example, they may both share the same counterstain (e.g. haematoxylin), or they may highlight different parts of the same structure.

The analysis and integration of information from different stainings is usually performed with reference to a specific organ, structure, or pattern observed in the tissue. For example, to diagnose pathologies such as breast cancer or kidney allograft rejection it is necessary to study the inflammatory micro-environment of the organ. In these cases, the relevant information is the distribution of immune cells (e.g. macrophages or lymphocytes) in relation to important structures of the organ, such as glomeruli for the kidney, or lobules for the breast.

To automatically perform such an analysis, the structure of interest (glomeruli, lobules, etc.) should be detected in each section irrespective of the individual staining modality. This is a typical application of deep learning, assuming that annotations for each staining are available. Thus, there is a high demand for annotations in digital pathology, currently limited by the availability of experts and the complexity of the required expert knowledge [5].

This article presents several training strategies that make progress towards stain invariant networks (i.e. a network that can generalise between stainings). By training the network using a commonly acquired staining in renal pathology Periodic acid-Schiff reaction (PAS, a histochemical reaction to detect certain types of sugar-containing tissue components) and applying the same network to images with different stainings, the presented strategies demonstrate significant improvements over standard training strategies. These can be used in settings where there is not enough data or annotations for the desired application staining. Alternatives in the literature are concerned with transfer learning or domain adaptation strategies where some take advantage of networks pretrained on huge natural image datasets [6]. Most of these works consider the problem of tissue variability between images that have the same staining [7, 8] or between images from different tissues [9]. These approaches are much

This work was supported by the SysMIFTA project (FKZ:031L-0085A). The authors are grateful to Nvidia Corporation for donating a Quadro P6000 GPU, which was used to conduct this research.

more complex to implement as they require specific network architectures.

The adopted CNN architecture and proposed training strategies are presented in the next section, Section 3 presents inter-staining glomeruli segmentation results and analysis thereof, and Section 4 presents the study’s conclusions.

2. METHOD

2.1. CNN architecture

The U-Net architecture [10] is adopted as it has been proven successful in biomedical imaging [11] and, in particular, glomeruli detection [2]. Glomeruli segmentation is framed as a two classes problem: glomeruli (patches centred on a glomeruli), and tissue (that does not contain a glomeruli, randomly sampled). The slide background (non-tissue) is removed by thresholding each image by its mean value then removing small objects from the result and closing holes. Initial experiments revealed that the same training parameter values lead to convergence irrespective of the strategy used (batch size of 8, learning rate of 0.0001, 60 epochs, and the network that achieves the lowest validation loss is kept). The input patch size is 508×508 pixels.

To mitigate the lack of training samples, the use of elastic deformation [12] is proposed by the original authors. In addition to this (using the parameters $\sigma = 10$, $\alpha = 100$), the following augmentations are applied with an independent probability of 0.5 (batches are augmented ‘on the fly’):

affine: random rotation sampled from the interval $[0^\circ, 180^\circ]$, random shift sampled from $[-205, 205]$ pixels, random magnification sampled from $[0.8, 1.2]$, and horizontal/vertical flip;

noise: additive Gaussian noise¹ with $\sigma \in [0, 2.55]$;

blur: Gaussian filter¹ with $\sigma \in [0, 1]$;

brightness enhance² with a factor sampled from $[0.9, 1.1]$;

colour enhance² with a factor sampled from $[0.9, 1.1]$;

contrast enhance² with a factor sampled from $[0.9, 1.1]$;

stain variation by colour deconvolution [13], α sampled from $[-0.25, 0.25]$ and β from $[-0.05, 0.05]$.

These values were chosen as they result in realistic images. All samples are standardised to $[0, 1]$ and normalised by the mean and standard deviation of the training set.

2.2. Training strategies

The standard approach for segmenting histopathological slides is to present the three colour channels to the network to exploit both colour and texture/structure information (referred to herein as RGB). Henceforth ‘structure’ refers to tissue appearance (i.e. is not used in a biological sense).

It is hypothesised that since histopathological slides of the same staining exhibit very low colour variance when com-

pared to natural images, a deep learning approach will rapidly learn to rely upon structures being present in specific channels (or combinations thereof), which may limit the network’s ability to generalise to unseen stainings. The effects of this are already known when dealing with inter-laboratory variance of the same staining [13, 14].

Objects to be segmented, e.g. glomeruli, are generally easily identified between stainings as globally they exhibit the same structure and texture, see Figure 1. It should therefore be possible to bias the network to learn stain invariant features. This work investigates this possibility by modifying the data presented to the network in an unsupervised manner. Borrowing domain adaptation terminology, herein the staining used for training and validation is referred to as the *source staining* (irrespective of any transformations), and the stainings to which the network are applied as *target stainings*.

The simplest approach is to remove the colour component from the data (referred to herein as *Greyscale*). Each sample is converted to greyscale using the formula $I = 0.2125R + 0.7154G + 0.0721B$ ¹. The result has all colour information compressed into one channel, forcing the network to learn structural features.

A more biologically motivated approach is to extract the haematoxylin channel from each staining by deconvolving the image using Ruifrok and Johnston’s approach [15] (referred to herein as *Haematoxylin*). Haematoxylin is commonly used as a counterstain and should therefore highlight the same structures in each staining. This strategy should therefore remove stain specific information and force the network to learn features that generalise well.

The above-mentioned strategies either compress information contained in the colour channels (Greyscale) or focus on consistent inter-staining information (Haematoxylin) with the aim of simplifying the problem. They are therefore applied during training, validation, and testing. Another approach is to preserve the colour information, but modify it in such a way that the network does not learn to rely on colour specific structure. These are therefore only applied during training and validation. A brute-force approach to this is to randomly permute the colour channels. This strategy—*Channel Swap*—is simple to implement; however, the range of colours generated will be far greater than those encountered in application.

The last strategy, named *Colour Transfer*, restricts this to within meaningful bounds at the expense of being more complicated. The colour profile of each training patch is replaced with that of a randomly selected staining during training with probability $\frac{N-1}{N}$, where N is the number of stainings in the dataset. Therefore each ‘staining’ (including the source) is presented to the network with equal probability. The colour transfer is achieved by deconvolving the patch and applying its stain concentrations to stain vectors taken from another staining. Various approaches to this exist [16, 17, 18] (including computer vision approaches [19]), in this work we use that proposed by Macenko et al. due to its simplicity (automatic

¹Using the appropriate [scikit-image](#) functions.

²Using the appropriate [PIL](#) functions.

determination of stain vectors) and realistic output (see Figure 1, bottom row). It is necessary, however, to have samples of the target stainings to determine their stain vectors (this requirement can be removed by using predetermined stain vectors and Ruifrok and Johnston’s approach [15]). During training, a random patch from a random target staining is selected and the target stain vectors are determined.

N.B. the intention is to present plausible colour profiles to the network. This does not necessarily correspond to biologically relevant image information but is explored for any potential for the tested staining modalities. Potential limitations are addressed, e.g. in Figure 1 the bright red colour (chromogenic reaction used to detect CD34, a marker of blood vessel inner lining) could be successfully transferred but it is not localised as in the original image.

3. EXPERIMENTS AND RESULTS

3.1. Data

Tissue samples were collected from a cohort of 10 patients who underwent allograft nephrectomy after complete loss of function of a transplanted kidney due to chronic graft failure. These kidneys displayed several signs of rejection. The paraffin-embedded samples were cut into 3 μm thick sections and stained with either Jones H&E, PAS or Sirius Red, in addition to two immunohistochemistry markers (CD34 and CD68), $N = 5$, using an automated staining instrument (Ventana Benchmark Ultra). Whole slide images were acquired using an Aperio AT2 scanner at 40 \times magnification (a resolution of 0.253 μm /pixel). All the glomeruli in each WSI were annotated and validated by pathology experts by outlining them using Cytomine [20]. The dataset was divided into 4 training, 2 validation, and 4 test patients. The number of glomeruli in each staining dataset was: PAS – 662 (train.), 588 (valid.), 1092 (test); Jones H&E – 1042 (test); Sirius Red – 1037 (test); CD34 – 1019 (test); CD68 – 1046 (test). The training set comprised all glomeruli from the source staining training patients (662) and 4634 tissue patches (to account for the variance observed in non-glomeruli tissue).

3.2. Experiments

For each strategy described in Section 2.2, a CNN was trained on the source staining (PAS) and applied to the target stainings (PAS, Jones H&E, Sirius Red, CD34, and CD68). Test patient whole slides were segmented and precision, recall and F_1 score of the glomeruli were taken. To evaluate robustness, five repetitions of each experiment were performed. Table 1 shows the means and standard deviations of the repetitions for each measure.

The segmentation results trained on and applied to the source staining (PAS) form the controls. State-of-the-art performance [2, 21, 22] is achieved or surpassed and this establishes a baseline for the best achievable inter-staining perfor-

mance. The segmentation results trained on the source staining and applied to the target stainings (RGB strategy) form the negative controls and, as the network essentially fails to detect anything, demonstrate the need for a transfer strategy.

The first strategy relies on compressing structural information into one band: Greyscale. This vastly increases target performance compared to RGB (with no affect on source performance) and leads to decent segmentation on Jones H&E and CD34, which are structurally the most similar to PAS.

Next is a strategy to extract biological information that is consistent across different stainings—the haematoxylin counterstain. Despite this biological plausibility, it does not result in good results. Indeed, the haematoxylin transformed images vary greatly between the different stainings (see Figure 1, second row). Several factors may explain this: 1) haematoxylin’s concentration relative to the primary stain may vary from one staining to another, resulting in different shades of blue and fixation amount; 2) as a counterstain, haematoxylin may become mixed with another stain in structures that are targeted by both. These result in a color mixing that is in practice not perfectly unmixed by a color deconvolution algorithm.

Although training a network using compressed structural information vastly improves inter-staining generalisation, acceptable segmentation performance is not achieved. The third strategy attempts to preserve color information but force the network not to specify the structure to particular colours. The Channel Swap strategy provides results equal to or better than the previous strategies in all but one staining (CD34), however this is achieved at the expense of stability due to the randomised colour channels.

Finally, colour variance is restricted to a more realistic range using colour transfer, giving the best results in all cases. It should be noted that the quality of the result is highly dependent on the degree of structural similarity between the target and source staining. Visually, there is a large difference between histochemical staining methods (e.g. H&E, PAS, Jones) and immunohistochemistry (antibody-mediated detection of certain structures like immune cells of blood vessel inner lining) with weak blue counterstain. Therefore, some limitations of the approach are expected, and the degree of variation may explain why detection of glomeruli in the CD68 staining did not work well with PAS as source staining.

4. CONCLUSIONS

It has been shown that the efficacy of simple transfer strategies depend on the degree of structural similarity between the source and target staining. The presented work dramatically improves the inter-staining segmentation performance when compared to standard training approaches, and this gain is correlated with the complexity of the strategy used. Nevertheless, it appears that a limit has been reached and strategies other than modifying training data should be the focus of future research into developing stain invariant networks.

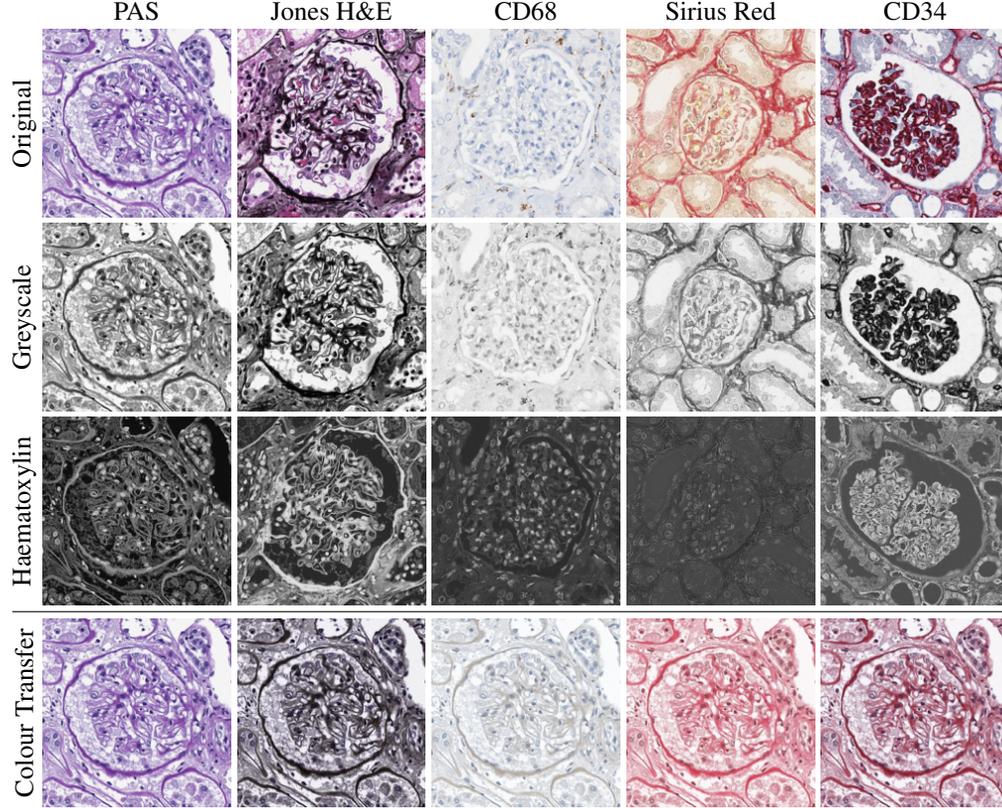


Fig. 1: Glomeruli patch examples with different transformations. Each column corresponds to a different staining, and each row a transformation. The colour transfer (4th row) results are applied to a PAS patch, using each staining as a target.

Training Strategy	Score	Test Staining				
		PAS	Jones H&E	CD68	Sirius Red	CD34
RGB	F ₁	0.901 (0.006)	0.031 (0.015)	0.000 (0.000)	0.015 (0.012)	0.011 (0.009)
	precision	0.873 (0.014)	0.023 (0.014)	0.372 (0.292)	0.009 (0.007)	0.052 (0.017)
	recall	0.932 (0.021)	0.075 (0.031)	0.000 (0.000)	0.063 (0.067)	0.006 (0.006)
Greyscale	F ₁	0.899 (0.004)	0.599 (0.080)	0.030 (0.033)	0.301 (0.059)	0.596 (0.050)
	precision	0.867 (0.013)	0.444 (0.088)	0.154 (0.098)	0.464 (0.072)	0.703 (0.058)
	recall	0.934 (0.007)	0.944 (0.007)	0.017 (0.020)	0.238 (0.078)	0.523 (0.071)
Haematoxylin	F ₁	0.872 (0.010)	0.676 (0.036)	0.002 (0.001)	0.001 (0.000)	0.042 (0.019)
	precision	0.807 (0.019)	0.609 (0.096)	0.029 (0.024)	0.003 (0.003)	0.079 (0.041)
	recall	0.948 (0.003)	0.785 (0.075)	0.001 (0.001)	0.000 (0.000)	0.034 (0.020)
Channel Swap	F ₁	0.889 (0.005)	0.671 (0.050)	0.106 (0.111)	0.508 (0.133)	0.379 (0.139)
	precision	0.846 (0.012)	0.659 (0.063)	0.167 (0.124)	0.658 (0.116)	0.605 (0.139)
	recall	0.936 (0.013)	0.700 (0.111)	0.090 (0.108)	0.457 (0.156)	0.299 (0.146)
Colour Transfer	F ₁	0.882 (0.015)	0.813 (0.042)	0.153 (0.083)	0.739 (0.070)	0.709 (0.028)
	precision	0.834 (0.039)	0.748 (0.078)	0.417 (0.133)	0.742 (0.106)	0.726 (0.062)
	recall	0.938 (0.015)	0.899 (0.017)	0.099 (0.059)	0.746 (0.073)	0.706 (0.089)

Table 1: Quantitative results for each strategy trained on PAS (source staining) and tested on different stainings (target stainings). The numbers in parentheses are the standard deviations of the corresponding scores.

5. REFERENCES

- [1] M. Gurcan, L. Boucheron, A. Can, et al., “Histopathological image analysis: A review,” *Reviews in Biomedical Engineering*, vol. 2, pp. 147–171, 2009.
- [2] T. de Bel, M. Hermsen, B. Smeets, et al., “Automatic segmentation of histopathological slides of renal tissue using deep learning,” in *Medical Imaging 2018: Digital Pathology*, 2018, vol. 10581, p. 1058112.
- [3] A. Janowczyk and A. Madabhushi, “Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases,” *Journal of Pathology Informatics*, vol. 7, no. 1, pp. 29, 2016.
- [4] K. Das, S. Conjeti, A. Roy, et al., “Multiple instance learning of deep convolutional neural networks for breast histopathology whole slide classification,” in *IEEE International Symposium on Biomedical Imaging*, 2018, pp. 578–581.
- [5] N. Schaadt, A. Grote, G. Forestier, et al., “Role of task complexity and training in crowdsourced image annotation,” in *Workshop on Computational Pathology, MIC-CAI*, 2018, pp. 44–51.
- [6] N. Bayramoglu and J. Heikkilä, “Transfer learning for cell nuclei classification in histopathology images,” in *European Conference on Computer Vision*, 2016, pp. 532–539.
- [7] M. Gadermayr, M. Strauch, B. Klinkhammer, et al., “Domain adaptive classification for compensating variability in histopathological whole slide images,” in *International Conference on Image Analysis and Recognition*, 2016, pp. 616–622.
- [8] M. Lafarge, J. Pluim, K. Eppenhof, et al., “Domain-adversarial neural networks to address the appearance variability of histopathology images,” in *International Workshop on Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 83–91. 2017.
- [9] M. Kandemir, “Asymmetric transfer learning with deep gaussian processes,” in *International Conference on Machine Learning*, 2015, pp. 730–738.
- [10] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [11] G. Litjens, T. Kooi, B. Bejnordi, et al., “A survey on deep learning in medical image analysis,” *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [12] P. Simard, D. Steinkraus, and J. Platt, “Best practices for convolutional neural networks applied to visual document analysis,” in *International Conference on Document Analysis and Recognition*, 2003, pp. 958–963.
- [13] D. Tellez, M. Balkenhol, I. Otte-Höllerand, et al., “Whole-slide mitosis detection in H&E breast histology using PHH3 as a reference to train distilled stain-invariant convolutional networks,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 9, pp. 2126–2136, 2018.
- [14] J. Gallego, A. Pedraza, S. Lopez, et al., “Glomerulus classification and detection based on convolutional neural networks,” *Journal of Imaging*, vol. 4, no. 1, pp. 20, 2018.
- [15] A. Ruifrok and D. Johnston, “Quantification of histochemical staining by color deconvolution,” *Analytical and Quantitative Cytology and Histology*, vol. 23, no. 4, pp. 291–299, 2001.
- [16] M. Macenko, M. Niethammer, J.S. Marron, et al., “A method for normalizing histology slides for quantitative analysis,” in *International Symposium on Biomedical Imaging*, 2009, pp. 1107–1110.
- [17] A. Vahadane et al., “Whole-slide mitosis detection in H&E breast histology using PHH3 as a reference to train distilled stain-invariant convolutional networks,” *IEEE Transactions on Medical Imaging*, vol. 35, no. 8, pp. 1962–1971, 2016.
- [18] C. Wemmert, J.M. Kruger, G. Forestier, et al., “Stain unmixing in brightfield multiplexed immunohistochemistry,” in *Image Processing (ICIP), 2013 20th IEEE International Conference on*. IEEE, 2013, pp. 1125–1129.
- [19] E. Reinhard, M. Ashikhmin, B. Gooch, et al., “Color transfer between images,” *IEEE Computer Graphics and Applications*, vol. 21, no. 5, pp. 34–41, 2001.
- [20] R. Marée, L. Rollus, B. Stévens, et al., “Collaborative analysis of multi-gigapixel imaging data using cytomine,” *Bioinformatics*, vol. 32, no. 9, pp. 1395–1401, 2016.
- [21] J. Bukowy, A. Dayton, D. Cloutier, et al., “Region-based convolutional neural nets for localization of glomeruli in trichrome-stained whole kidney sections,” *Journal of the American Society of Nephrology*, vol. 29, no. 8, pp. 2081–2088, 2018.
- [22] D. Govind, B. Ginley, B. Lutnick, et al., “Glomerular detection and segmentation from multimodal microscopy images using a butterworth band-pass filter,” in *SPIE Medical Imaging, 2018*, 2018, p. 1058114.