



Semi-supervised Adversarial Learning for Stain Normalisation in Histopathology Images

Cong Cong¹, Sidong Liu^{2,3}, Antonio Di Ieva³, Maurice Pagnucco¹, Shlomo Berkovsky², and Yang Song¹(✉)

¹ School of Computer Science and Engineering, University of New South Wales, Sydney, Australia

yang.song1@unsw.edu.au

² Centre for Health Informatics, Macquarie University, Sydney, Australia

³ Computational NeuroSurgery Lab, Macquarie University, Sydney, Australia

Abstract. Hematoxylin and Eosin (H&E) stained histopathology images provide important clues for diagnostic and prognostic assessment of diseases. However, similar tissues can be stained with inconsistent colours which significantly hinder the diagnostic process and training of deep learning models. Various Generative Adversarial Network (GAN) based stain normalisation methods have thus been proposed as a preprocessing step for the downstream classification or detection tasks. However, most of these methods are based on either unsupervised learning which suffers from large discrepancy between domains or supervised learning which requires a target domain and only utilises the target domain images. In this work, we propose to leverage Semi-supervised Learning with GAN to incorporate the source domain images in the learning of stain normalisation without requiring their corresponding ground truth data. Our approach achieves highly effective performance on two classification tasks for brain and breast cancers.

Keywords: Semi-supervised learning · Stain normalisation · Conditional generative adversarial networks

1 Introduction

Tissue staining is used to facilitate effective interpretation of histopathology images. However, the appearance of stained tissue slides can be highly heterogeneous due to the different staining protocols and the subsequent digitisation of the images. This undesired colour variance in histopathology images motivates the study of stain normalisation, which normalises the images to reduce the impact of colour heterogeneity.

Existing studies of stain normalisation normally fall into two main categories, traditional and deep learning-based methods. Specifically, traditional stain normalisation methods use mathematical frameworks to match the colour distribution of an input image with the selected reference images. Examples include the

colour normalisation in the LAB colour space [18], minimisation of the Wasserstein Barycenter between colour distributions [15], and normalisation of stain vectors [12, 29]. However, most of these conventional methods rely on a reference image to calculate the stain statistics, therefore the stain normalisation results can be heavily biased by a less representative reference image.

Deep learning-based stain normalisation is used more often these days. Most of these methods leverage the Generative Adversarial Network (GAN) [8] to formulate stain normalisation as an image-to-image translation task. Unsupervised stain normalisation has been widely studied, as it eases the requirement of a reference image. For example, StainGAN [24] utilises the CycleGAN model [31] for unsupervised stain normalisation. SAASN [25] incorporates the self-attention mechanism in CycleGAN to achieve finer details. Another method [13] combines CycleGAN with a pretrained segmentation network to better preserve the cellular structures in the images. While these methods produce satisfactory results in most cases, unsupervised methods still suffer from the large discrepancy between domains in heterogeneous datasets. Thus, supervised stain normalisation, which uses paired images from the target domain, has been developed. A HRNet-based [27] neural style transfer GAN is developed [16], which trains with selected reference images and preserves finer details by including a perceptual loss [9]. Others treat stain normalisation as an image colourisation task, in which they learn a model to repaint the input images with the target domain colour style [23, 30]. Compared to unsupervised approaches, these supervised methods produce substantially better normalised images. However, the colourisation task is only performed on the target domain images, without considering the source domain. Such a colourisation formulation does not truly represent the objective of stain normalisation, which is to align the colour distributions between source and target domains.

In this work, we propose to use Semi-Supervised Learning (SSL) with GAN to solve stain normalisation as an image colourisation task. Our model learns to repaint the input images with the target domain colours and we use the image from the source domain to enhance this mapping. Specifically, given the hematoxylin component [21] or grayscale image as input, our GAN-based stain normalisation network learns to repaint it with the target colour styles. In order to incorporate source domain images into the training and explicitly optimise the colourisation for source domain images, a two-decoder generator with a shared encoder is proposed. This design allows us to apply a novel consistency regularisation which enforces the model to produce consistent outputs for the source domain and also minimises their difference from the target domain colour distribution.

Contributions: 1) We design a novel semi-supervised colourisation model for stain normalisation, so that source domain images can be incorporated into the model learning without requiring paired ground truth data for the source domain. 2) Our model adopts a novel two-decoder design with consistency loss to enforce the generator to normalise the unlabelled source images with the desired target domain colour. 3) We conduct extensive experiments on the public

TCGA [6] and BreakHis [26] datasets, and demonstrate that our method leads to large improvement for the downstream classification task on both types of histopathology images.

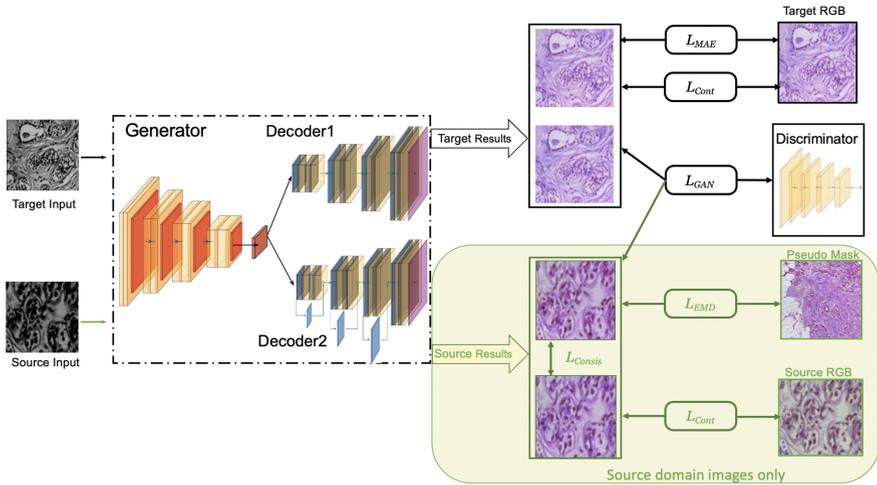


Fig. 1. Our proposed semi-supervised GAN-based stain normalisation model, by incorporating source domain images into the training with pseudo masks in a semi-supervised learning manner.

2 Methods

We design a semi-supervised GAN-based colourisation model for stain normalisation. The objective of this model is to colourise the inputs with the target domain stain colours. For target domain images, we use a supervised learning strategy with paired training data (original H&E stained colour image as ground truth and corresponding grayscale or hematoxylin component as input). Such paired training data are only available in the target domain but not in the source domain. Therefore, to explicitly enhance the colourisation performance for source domain images, we design a semi-supervised learning approach to incorporate the source domain images with pseudo masks into the model training.

Figure 1 shows the overall design of our model. The generator network colourises the input images with target stain style. The PatchGAN [23] discriminator judges whether the generator output is fake or real. The model is trained alternately with batches of labelled target domain images and unlabelled source domain images. For the labelled target domain, our model is trained in a supervised fashion with adversarial loss and content loss [9], whereas on the unlabelled

source domain, an additional supervised loss is introduced for consistency regularisation based on the generated RGB images. The average of the two decoder branches' outputs in the generator as the final normalisation results.

2.1 Normalisation via Colourisation

We first describe the stain normalisation process by considering the target domain only. Specifically, the target domain hematoxylin component or grayscale images X and their corresponding RGB images Y are used as labelled paired data to train a Conditional GAN [14] in which the generator learns to colourise the inputs with the target domain colour statistics. The generator, which has an encoder-decoder architecture like U-Net [19] encodes the input images to extract high-level features and then decodes the extracted features to get the colourised images. Note that in Fig. 1, this generator design consists of Encoder and Decoder1, while Decoder2 is introduced for SSL as explained in the Sect. 2.2. Then the discriminator judges the image pairs $(X; Y)$ or $(X; G(X))$, as real or fake by maximising the following the adversarial loss function:

$$\mathcal{L}_{GAN}(X, Y) = \log(D(X, Y)) + \log(1 - D(X, G(X))) \quad (1)$$

where, the term $\log(D(X, Y))$ will not be calculated for the unlabelled source domain image pairs.

To further regulate the model, such that it not only learns the correct colour mapping but also preserves high-level structural image content, we minimise the content loss \mathcal{L}_{Cont} between perceptual features of generated RGB images $\{\hat{Y}\}$ and the original RGB images $\{Y\}$. In our case, the feature presentations of the last four convolutional layers ($n = 4$) in a pretrained VGG16 are used to form the high-level features to derive the content loss:

$$\mathcal{L}_{Cont}(\hat{Y}, Y) = \sum_j^n \omega_j \frac{1}{C_j H_j W_j} \|\phi_j(\hat{Y}) - \phi_j(Y)\| \quad (2)$$

where ϕ_j is the feature map produced from j_{th} convolution layer before the max pooling layer, $C_j H_j W_j$ is the shape of ϕ_j and $\omega_j = 1/n$.

2.2 Semi-supervised Source Domain Normalisation

We then introduce SSL into the model so that the unlabelled source domain images can also be used in learning the correct colourisation mapping. Recent deep learning-based SSL methods normally implement a consistency regularisation that encourages the model to perform consistently with disturbance in the input or model [10, 22, 28] However, it is likely that such methods could produce consistent but incorrect labels for unlabelled data. Therefore, we propose an additional colour distance-based regularisation which enforces the model output distribution to be closer to the ground truth distribution.

To do this, inspired by [7, 10], we force the consistency between the two decoders by minimising the Mean Absolute Distance (L_{MAE}) between their outputs. We further maximise the mutual information [5] by distribution alignment, which is firstly used in [4] by scaling the output to match the label distribution. In our work, we propose to minimise the Earth Mover’s Distance [20] (L_{EMD}) between the decoders’ outputs and a pseudo mask selected from the target domain images as a form of colour distribution alignment.

We explicitly design **Decoder2** as an improved version of **Decoder1** such that the introduction of **Decoder2** will not degrade the model performance. A simple and effective way of such a modification can be achieved by adding residual blocks which short-circuit the concatenated inputs to the outputs of the convolutional layers. In terms of L_{EMD} calculation, a pseudo mask $\{Y'\}$ with target domain colour statistics is required for the normalised RGB images of unlabelled source domain images to match with. In order to obtain the mask that can best represent the colour statistic of the target domain, we choose the target domain image whose mean and standard deviation of pixel colours are closest to the overall target domain mean and standard deviation as this pseudo mask $\{Y'\}$. This pseudo mask is likely to have a different tissue pattern from the source domain image, but aims to provide a guidance on the colour distribution. Then, for the outputs of the two decoders on unlabelled source domain images $\hat{Y}^s = (\hat{Y}_1^s, \hat{Y}_2^s)$, we apply the following consistency loss:

$$\mathcal{L}_{Consist}(\hat{Y}^s, Y') = L_{MAE}(\hat{Y}_1^s, \hat{Y}_2^s) + L_{EMD}(\hat{Y}_1^s, Y') + L_{EMD}(\hat{Y}_2^s, Y') \quad (3)$$

2.3 Training Pipeline

We evaluate our model with two different inputs, the hematoxylin component of the H&E stained images, which can be extracted from the RGB images using the Beer-Lambert’s Law [17], and the grayscale image. They tend to perform differently on different datasets. The generator then colourises the inputs with the target colour style. We use different loss functions based on the input domain. If the inputs are from the target domain, we use adversarial loss, content loss and $L1$ loss to regularise the generator:

$$\mathcal{L}_{target}(X, Y^T) = \mathcal{L}_{GAN}(X, Y^T) + \mathcal{L}_{Cont}(G(X), Y^T) + \mathcal{L}_{MAE}(G(X), Y^T) \quad (4)$$

If the inputs are from the unlabelled source domain, we then incorporate the additional consistency loss in place of the L_{MAE} loss to encourage the network to produce high-quality content-preserved normalised RGB images:

$$\mathcal{L}_{source}(X, Y^S, Y') = \mathcal{L}_{GAN}(X, Y^S) + \mathcal{L}_{Cont}(G(X), Y^S) + \mathcal{L}_{Consist}(G(X), Y') \quad (5)$$

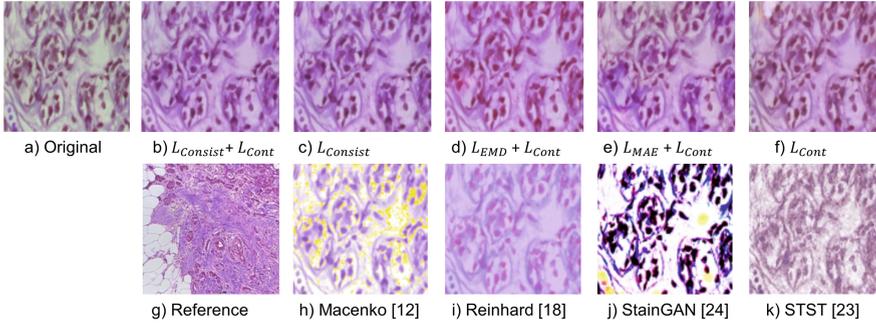


Fig. 2. Comparison of stain normalised results.

3 Experiments and Results

3.1 Dataset and Implementation

We evaluate our method using two datasets, the TCGA¹ glioma cohort for Isocitrate Dehydrogenase (IDH) prediction [11] and the BreakHis database for breast cancer histopathological image classification [26]. We are grateful to the authors of [11] for providing us the IDH dataset, which is a subset of the public TCGA dataset [6]. It consists of 22,229 images each of 1024×1024 pixels extracted from 921 patients' whole slide images (WSIs) at $10\times$ magnification level. All WSIs have been labelled as either IDH wildtype or mutant confirmed by immunohistochemistry and/or genetic sequencing. The BreakHis dataset contains 7,909 images of 460×700 pixels collected from 82 patients. These images are stored with four magnification levels and are annotated as benign or malignant. In order to train our stain normalisation network, we split the IDH dataset based on the tissue source site (TSS) and choose the largest TSS as the target domain (3,414 images) and the rest as the source domain (18,805 images). The BreakHis dataset does not contain the tissue source site label, thus, we use k -means ($k = 5$) clustering to cluster the images based on the mean and standard deviation of image pixel colours. The largest cluster is used as the target domain (350 images) and the remaining 7,559 images as considered the source domain.

We fine-tune an ImageNet pretrained ResNet50 for the two binary classification tasks using the stain normalised images. For IDH classification, images from 738 randomly selected patients (80% of the cohort) are used for training, those from 91 patients for validation and the remaining 92 patients for testing. For the BreakHis dataset, we follow the 5-fold cross validation setting in [26] with 70% of data for training and 30% for testing in each split, and images of different magnification levels are mixed together during the training and testing. We measure the impact of stain normalisation on the classification performance using F1-score, accuracy and Area Under the Curve (AUC). We also measure the colour consistency of stain normalised images by computing the standard

¹ <https://portal.gdc.cancer.gov/>.

deviation (NMI_{SD}) and coefficient of variation (NMI_{CV}) of Normalised Median Intensity (NMI) [1].

Both the stain normalisation model and ResNet50 classifier are developed using TensorFlow Keras on NVIDIA Tesla P100 GPUs. We reshape the image to 256×256 pixels for both datasets and then train the models with the Adam optimiser with an initial learning rate of 0.0002. We train our stain normalisation model for 200 epochs and fine-tune the ResNet50 classifier for 40 epochs. The weights of loss terms are set to $\lambda_{L1} = 0.25$ and $\lambda_{content} = 0.75$.

3.2 Results

We compare our model with state-of-the-art methods that have experimental results reported on the two datasets. Specifically, for the IDH classification task, we compare with the recent study [11] that collected the IDH dataset. Different from our approach, [11] applies GAN for data augmentation to improve the classification performance without stain normalisation. For BreakHis, many approaches have been reported over the recent years. We choose to compare with the deep learning approaches [2, 3] that use the same experimental settings as ours. In particular, [3] has incorporated a traditional stain normalisation method [18] before performing classification. As shown in Table 1, our method (using either the hematoxylin component or grayscale image as input) achieves about 5%–7% performance improvement over [11] and [3]. We find that using the hematoxylin component as inputs is more suitable for images with low magnification levels. Though, using hematoxylin component can enhance the contrast at cell boundaries, it may not always appear in an image patch extracted from high magnification levels.

Table 1. Performance comparison with state-of-the-art on two datasets.

	IDH [11]					BreakHis [26]				
	F1	Acc	Auc	NMI_{SD}	NMI_{CV}	F1	Acc	Auc	NMI_{SD}	NMI_{CV}
w/o normalisation	0.815	0.821	0.90	0.087	0.114	0.823	0.823	0.909	0.072	0.098
IDH study [11]	–	0.870	0.938	–	–	–	–	–	–	–
BreakHis study [3]	–	–	–	–	–	–	0.890	–	–	–
BreakHis study [2]	–	–	–	–	–	–	0.834	–	–	–
Macenko [12]	0.878	0.860	0.914	0.064	0.077	0.899	0.938	0.885	0.025	0.028
Reinhard [18]	0.815	0.833	0.912	0.054	0.052	0.911	0.910	0.918	0.022	0.023
StainGAN [24]	0.878	0.870	0.917	0.044	0.052	0.902	0.898	0.944	0.024	0.028
STST [23]	0.891	0.880	0.918	0.041	0.054	0.937	0.935	0.972	0.018	0.022
Ours-Hematoxylin	0.937	0.934	0.984	0.035	0.040	0.948	0.962	0.950	0.020	0.021
Ours-Grayscale	0.878	0.882	0.920	0.060	0.071	0.980	0.980	0.996	0.017	0.019

We have also trained the same ResNet50 classifier with the stain normalised images using other stain normalisation methods [12, 18, 23, 24]. As shown in Table 1, our method provides better colour consistency (*i.e.* lower NMI_{SD} and NMI_{CV}) and improves the classification performance (*i.e.* 2%–7%). It can also

be seen that without (w/o) stain normalisation, the results are consistently lower than using the stain normalised images.

Table 2. Performance comparison with different pseudo masks.

	IDH [11]			BreakHis [26]		
	P_1	P_2	P_3	P_1	P_2	P_3
Macenko [12]	0.860	0.854	0.818	0.938	0.922	0.912
Reinhard [18]	0.833	0.819	0.780	0.910	0.894	0.882
Ours	0.934	0.928	0.931	0.980	0.978	0.974

We further test the choice of pseudo masks. Besides our pseudo mask design (P_1) described in Sect. 2.2, we evaluated two alternatives: (P_2) a target domain image that is structurally most similar (measured by SSIM) to the source domain, and (P_3) a randomly chosen image from the target domain. The results in Table 2 show that our model has less than 1% accuracy drop when using different masks, whereas other competing methods show 2%–5% drop. This validates our design of the pseudo mask and shows robustness compared to other methods.

Table 3. Performance comparison with different loss functions.

	$L_{Consist}$	L_{EMD}	L_{Cont}	F1	Accuracy	AUC
IDH [11]			✓	0.878	0.880	0.945
	✓		✓	0.923	0.931	0.972
		✓	✓	0.928	0.929	0.969
	✓	✓		0.860	0.865	0.938
	✓	✓	✓	0.937	0.934	0.984
BreakHis [26]			✓	0.944	0.950	0.956
	✓		✓	0.958	0.968	0.991
		✓	✓	0.969	0.970	0.992
	✓	✓		0.914	0.908	0.964
	✓	✓	✓	0.980	0.980	0.997

Table 3 and Fig. 2 show the results with different loss functions. It can be observed that each term contributes to the performance as removing a single component degrades the model. Visually, using the proposed consistency regularisation alone can already reach a satisfactory result, but adding the content loss brings extra benefits. However, using the content loss alone is not enough to produce satisfactory normalised masks for the unlabelled source domain images.

This can be remitted by adding $L_{Consist}$ or L_{EMD} regularisation which indicates the usefulness of the proposed consistency regularisation. Figure 3 shows the impact of introducing different numbers of source images into training. The number a on the x -axis indicates the ratio of the total number of introduced source domain images relative to the number of target domain images. Generally, the model performs better when more source images are introduced into the training.

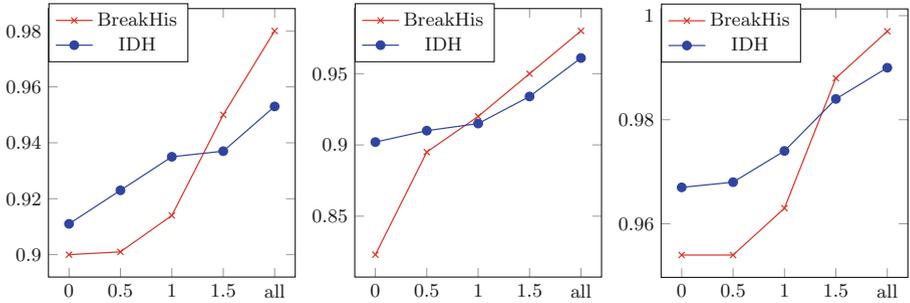


Fig. 3. Performance comparison with different numbers of source image used. Left to right: F1-score, accuracy and AUC.

4 Conclusions

In this paper, we propose a semi-supervised stain normalisation framework. The proposed model learns to colourise input images with the colour style of the target domain using both labelled target domain images and unlabelled source domain images. Our evaluation results show that our model produces higher quality images with high colour consistency with the target domain. We also show performance improvement over the prior art for two different histopathology image classification tasks.

References

1. Basavanthally, A., Madabhushi, A.: EM-based segmentation-driven color standardization of digitized histopathology. In: Medical Imaging 2013: Digital Pathology, vol. 8676, p. 86760G (2013)
2. Bayramoglu, N., Kannala, J., Heikkilä, J.: Deep learning for magnification independent breast cancer histopathology image classification. In: International Conference on Pattern Recognition (ICPR), pp. 2440–2445 (2016)
3. Benhammou, Y., Achchab, B., Herrera, F., Tabik, S.: BreakHis based breast cancer automatic diagnosis using deep learning: taxonomy, survey and insights. *Neurocomputing* **375**, 9–24 (2020)

4. Berthelot, D., et al.: RemixMatch: semi-supervised learning with distribution alignment and augmentation anchoring. In: International Conference on Learning Representations (ICLR) (2019)
5. Bridle, J.S., Heading, A.J., MacKay, D.J.: Unsupervised classifiers, mutual information and ‘phantom targets’. In: Advances in Neural Information Processing Systems (NIPS) (1992)
6. Clark, K., et al.: The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digit. Imaging (JDI)* **26**(6), 1045–1057 (2013)
7. Fang, K., Li, W.-J.: DMNet: difference minimization network for semi-supervised segmentation in medical images. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 532–541. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_52
8. Goodfellow, I.J., et al.: Generative adversarial networks. In: Conference on Neural Information Processing Systems (NIPS) (2014)
9. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
10. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: International Conference on Learning Representations (ICLR) (2016)
11. Liu, S., et al.: Isocitrate dehydrogenase (IDH) status prediction in histopathology images of gliomas using deep learning. *Sci. Rep.* **10**(1), 1–11 (2020)
12. Macenko, M., et al.: A method for normalizing histology slides for quantitative analysis. In: IEEE International Symposium on Biomedical Imaging (ISBI), pp. 1107–1110 (2009)
13. Mahapatra, D., Bozorgtabar, B., Thiran, J.-P., Shao, L.: Structure preserving stain normalization of histopathology images using self supervised semantic guidance. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12265, pp. 309–319. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59722-1_30
14. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784) (2014)
15. Nadeem, S., Hollmann, T., Tannenbaum, A.: Multimarginal Wasserstein barycenter for stain normalization and augmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12265, pp. 362–371. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59722-1_35
16. Nishar, H., Chavanke, N., Singhal, N.: Histopathological stain transfer using style transfer network with adversarial loss. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12265, pp. 330–340. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59722-1_32
17. Parson, W.W.: *Modern Optical Spectroscopy*, vol. 2. Springer, Heidelberg (2007)
18. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Comput. Graph. Appl. (CG&A)* **21**(5), 34–41 (2001)
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
20. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover’s distance as a metric for image retrieval. *Int. J. Comput. Vis. (IJCV)* **40**(2), 99–121 (2000)
21. Ruifrok, A.C., Johnston, D.A.: Quantification of histochemical staining by color deconvolution. *Anal. Quantit. Cytol. Histol. (AQCH)* **23**(4), 291–299 (2001)

22. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In: Conference on Neural Information Processing Systems (NIPS) (2016)
23. Salehi, P., Chalechale, A.: Pix2pix-based stain-to-stain translation: a solution for robust stain normalization in histopathology images analysis. In: International Conference on Machine Vision and Image Processing (MVIP), pp. 1–7 (2020)
24. Shaban, M.T., Baur, C., Navab, N., Albarqouni, S.: StainGAN: stain style transfer for digital histological images. In: International Symposium on Biomedical Imaging (ISBI), pp. 953–956 (2019)
25. Shrivastava, A., et al.: Self-attentive adversarial stain normalization. In: Del Bimbo, A., et al. (eds.) ICPR 2021. LNCS, vol. 12661, pp. 120–140. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-68763-2_10
26. Spanhol, F.A., Oliveira, L.S., Petitjean, C., Heutte, L.: A dataset for breast cancer histopathological image classification. *IEEE Trans. Biomed. Eng. (TBME)* **63**(7), 1455–1462 (2015)
27. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5693–5703 (2019)
28. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Conferences on Neural Information Processing Systems (NIPS) (2017)
29. Vahadane, A., et al.: Structure-preserving color normalization and sparse stain separation for histological images. *IEEE Trans. Med. Imaging* **35**(8), 1962–1971 (2016)
30. Yuan, E., Suh, J.: Neural stain normalization and unsupervised classification of cell nuclei in histopathological breast cancer images. arXiv preprint [arXiv:1811.03815](https://arxiv.org/abs/1811.03815) (2018)
31. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 2223–2232 (2017)