



Data Alchemy: Mitigating Cross-Site Model Variability Through Test Time Data Calibration

Abhijeet Parida^{1,2}, Antonia Alomar³, Zhifan Jiang¹, Pooneh Roshanitabrizi¹, Austin Tapp¹, María J. Ledesma-Carbayo², Ziyue Xu⁴, Syed Muhammed Anwar^{1,5}, Marius George Linguraru^{1,5(✉)}, and Holger R. Roth⁴

¹ Children's National Hospital, Washington, DC, USA
mlingura@childrensnational.org

² Universidad Politécnica de Madrid, Madrid, Spain

³ Universitat Pompeu Fabra, Barcelona, Spain

⁴ Nvidia Corporation, Santa Clara, CA, USA

⁵ George Washington University, Washington, DC, USA

Abstract. Deploying deep learning-based imaging tools across various clinical sites poses significant challenges due to inherent domain shifts and regulatory hurdles associated with site-specific fine-tuning. For histopathology, stain normalization techniques can mitigate discrepancies, but they often fall short of eliminating inter-site variations. Therefore, we present *Data Alchemy*, an explainable stain normalization method combined with test time data calibration via a template learning framework to overcome barriers in cross-site analysis. *Data Alchemy* handles shifts inherent to multi-site data and minimizes them without needing to change the weights of the normalization or classifier networks. Our approach extends to unseen sites in various clinical settings where data domain discrepancies are unknown. Extensive experiments highlight the efficacy of our framework in tumor classification in hematoxylin and eosin-stained patches. Our explainable normalization method boosts classification tasks' area under the precision-recall curve (AUPR) by 0.165, 0.545 to 0.710. Additionally, *Data Alchemy* further reduces the multisite classification domain gap, by improving the 0.710 AUPR an additional 0.142, elevating classification performance further to 0.852, from 0.545. Our *Data Alchemy* framework can popularize precision medicine with minimal operational overhead by allowing for the seamless integration of pre-trained deep learning-based clinical tools across multiple sites.

Keywords: Calibration · Generalizibility · Explainability · Histopathology

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-73290-4_15.

1 Introduction

In recent years, deep learning-based methods have performed well for various medical imaging analysis tasks such as disease diagnosis, classification, and segmentation [1, 21]. However, according to the United States Food and Drugs Administration, there is no approval for artificial intelligence and machine learning-enabled medical devices in histopathology for the calendar year 2023 [4]. This suggests few of the developed methods are usable in a clinical setting – particularly in histopathology, due to known challenges of generalizability and robustness across sites. Data and protocol variability across sites further hamper the approval of regulatory compliance for such tools [3, 7]. The typical approach to improve a model’s performance and generalizability is to calibrate each model at every site before deployment [12, 18]. While effective in some circumstances, model weight calibration resulting in substantial parameter-related modifications necessitates regulatory re-approval. To overcome these challenges, we propose a different approach. Instead of performing weight calibration that would necessitate regulatory re-approval of the model, we perform data calibration/template learning using *Data Alchemy* to reduce the gap domain and hence, solve the generalizability problem at test time. To establish the efficacy of this approach, we address tumor classification in digital histopathology images.

In histopathology, cells and tissue samples must be stained to be visible under a microscope. Then, they are digitized using microscopic scanners. The resulting samples’ appearance varies depending on several factors such as the used reagents, staining procedure, and scanner specifications. Such variations directly affect analysis performed both by a pathologist or automated classification algorithms [22]. Stain normalization has been investigated as a pre-processing step to reduce color variations between histopathology samples. This involves transferring the color (stain) of a source histology patch to a target patch, while preserving the morphological tissue structure (content). Several studies have shown that data augmentation and stain normalization help increase the prediction accuracy [23, 27, 28]. However, striking the appropriate balance between structure preservation and color consistency is challenging, as the resulting samples either contain artifacts and hallucinations or suffer in color appearance.

Related Works: Conventional stain normalization methods are mostly based on histogram transformations or color deconvolution (stain separation) [6, 8]. Histogram transform-based methods usually impose the color characteristics of a reference patch to another source patch using linear transformations [19, 30]. Color deconvolution is a method for decoupling light-absorbance and stain concentration in each pixel using spectral characteristics of different stains [17]. In other works, such as [15], RGB images are transformed into optical flows for estimating the stain vectors using singular value decomposition (SVD). However, these methodologies tend to generate artifacts in the background and/or color discontinuities in the normalized images.

Recent efforts have focused on deep learning-based methods, especially those using generative adversarial networks (GANs). GAN-based methods target stain

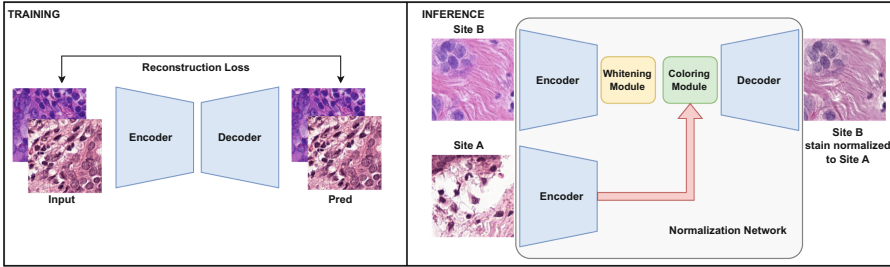


Fig. 1. Stain normalization during training (left) and inference (right).

normalization as a style-transfer problem [5]. Some proposed methods have used cycle-consistent generative adversarial networks (cycleGAN) to match the target distribution [10, 20, 25]. In another approach, content was disentangled from style, opening the possibility of multiple stain representations and, in classification tasks, outperforming conventional color augmentation techniques [29]. However, GANs are computationally expensive, are prone to mode collapse, and can lead to undesired changes in the underlying morphological structures [8].

Our Contribution: 1) A **stain normalization** method that combines the advantages of Singular Value Decomposition (SVD) transformations in the latent space with the non-linearity of convolutional networks to ensure structure preservation in a simple, interpretable, and computationally efficient manner. 2) We propose a **test time data calibration** method via template learning called *Data Alchemy* that improves model generalizability without altering parameters during testing, thus maintaining regulatory compliance. 3) We demonstrate the effectiveness of our strategies by evaluating them on histopathological tumor classification data.

2 Methods and Experimental Settings

2.1 Explainable Stain Normalization

We approached histopathology stain normalization as an image reconstruction task using feature transformation during inference, as shown in Fig. 1. Specifically, an image reconstruction network was trained using image I , such that $I = dec(enc(I))$, where $enc(\cdot)$ and $dec(\cdot)$ are the encoder and decoder, respectively. Feature transformations were done using whitening and coloring transforms, proposed for arbitrary style transfer between natural images [14].

The whitening transform was defined as $f_c = E_c D_c^{-\frac{1}{2}} E_c^T enc(I_c)$, where D_c is a diagonal matrix of eigenvalues and E_c is the orthogonal matrix of eigenvectors of the covariance matrix $enc(I_c) \cdot enc(I_c)^T$, and I_c represent the patches that need to be re-stained. The covariance matrix is positive semi-definite, ensuring all eigenvalues ≥ 0 . This whitening transform removed stain-specific information




Fig. 2. Data Alchemy uses the normalization network and a classifier to learn a template at test time to improve the classifier performance when deployed at a site.

while preserving structure-related information from the patch that needs to be re-stained.

The “whitened” f_c was then “colorized” using the coloring transforms, defined as $f_{cs} = E_s D_s^{-\frac{1}{2}} E_s^T f_c$ [9, 14], where D_s is a diagonal matrix of eigenvalues and E_s is the orthogonal matrix of eigenvectors of the covariance matrix $enc(I_s) \cdot enc(I_s)^T$, and I_s represents the patch whose staining parameters are used to stain the patch I_c . The coloring transform added stain-specific information from I_s to the “whitened” f_c .

The features of I_c can be blended using a parameter α with re-stained features, f_{cs} , to control the stylization effect [14], as $f_{cs} = \alpha f_{cs} + (1 - \alpha) enc(I_c)$. For patch staining, we set $\alpha = 1$ as we aim to produce a stained patch and not control stylization.

Implementation Details: We used all layers of VGG-19 [26] upto ‘*conv_3_3*’ as the encoder and the exact inverted architecture of the encoder as the decoder (Fig. 1). We minimized L1 as a reconstruction loss using the AdamW optimizer for 10 epochs with a batch size of 96 and a learning rate of $1e^{-4}$. The best-performing model on the validation set was saved for stain normalization.

2.2 Downstream Classification Task

To evaluate our stain normalization method, we reimplemented a downstream classification task from [13]. We used the ResNet-34 to identify tumor cells in small patches of whole slide images (WSIs). The classifier was trained using patches from one site, and tested on the unseen site. We compare the ResNet’s accuracy using stain normalization with a fixed patch template and our proposed test time data calibration to establish generalizability.

Implementation Details: We trained three models (ResNet-34): one on site, A, one on site B, and one on combined sites (A and B). The models were trained with augmentations from [13], which included color jitter, changes in brightness, hue and saturation, random flips, and rotation. We minimized the cross-entropy loss for tumor *vs.* healthy patches using the AdamW optimizer for 60 epochs

Table 1. Data summary of the splits, sites, and number of patches for the CAMELYON 16 [2]. The WSI for each of the sites is available in appendix A.

	Site A		Site B	
	healthy	tumor	healthy	tumor
Train	77,204	52,000	71,447	60,000
Validation	33,306	58,000	11,551	8,000
Test	20,183	34,000	8,449	12,000

with a batch size 256 and a learning rate of $1e^{-4}$. The model with the best validation metrics was chosen as our classifier.

2.3 Data Alchemy: Test Time Data Calibration

For classifiers trained in Sect. 2.2 to function optimally at different sites, a calibration step was necessary. We used the normalization method from Sect. 2.1 to adjust incoming patches to familiar stain parameters. Since we did not want to alter model weights, we propose adjusting the target template of the normalization network instead.

We randomly drew a real patch from the classifier’s training site to instantiate a template. We froze the normalization network and the classifier and set the template tensor as learnable. During calibration, labeled images from the test site were normalized to match the staining of the training site. These stain-normalized images were passed through the classifier to obtain class logits, which were used to calculate the losses with the labels of the patches. The only learnable parameter was the template, so gradients were calculated for it, and over multiple iterations, the optimizer learned a synthetic template. Thus, the classifier guided the normalization network in modifying the template to improve the site’s classification accuracy. This calibration step, was performed during deployment and is the test-time calibration called *Data Alchemy*. The schematic for test-time calibration is shown in Fig. 2.

Implementation Details: The validation set from the sites is used to learn a template for calibrating the classifier. Half of the dataset was used to learn the template and the rest is the validation set of the data calibration step. Optimization is performed for 10 epochs to minimize the cross-entropy loss using the AdamW optimizer with a batch size of 256 and a learning rate of $1e^{-4}$.

2.4 Dataset

We used CAMELYON 16 [2], a public dataset consisting of 400 WSIs of sentinel lymph nodes from two sites, site A - Radboud University Medical Center, Nijmegen, and site B - University Medical Center, Utrecht. Further, we used coordinates provided by Baidu Research [13] to determine the presence or absence of tumor cells in 256×256 patches. The site-wise sample distribution is presented in Table 1.

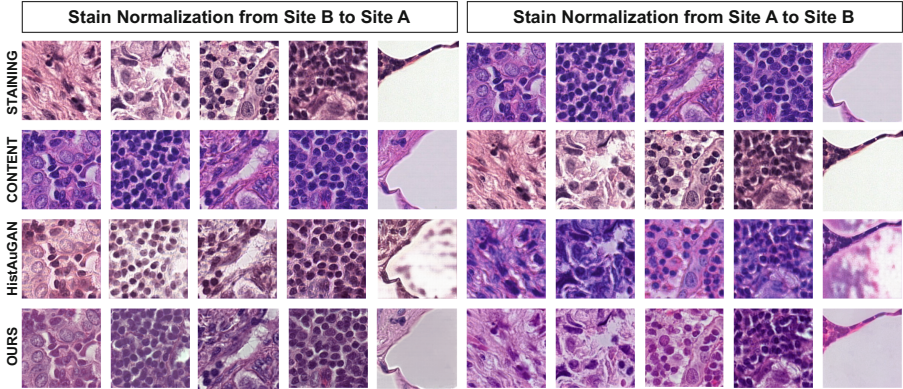


Fig. 3. Qualitative visualization of stain normalization. A detailed version is available in appendix C.

2.5 Evaluation Metrics

Stain Normalization: We used metrics of structural similarity index measure (SSIM) and peak-signal-to-noise ratio (PSNR) [24] for evaluation. We also used specialized metrics $cycleL1$ and $AP(i, p)$ [16] to quantify the preservation of structural information and the accuracy of stain normalization. $cycleL1 = \|I_c, sty(sty(I_c, I_s), I_c)\|$, is the norm between the original patch I_c and the reconstructed original patch after two stain normalizations using normalization network $sty(\cdot)$. $Sty(\cdot)$ stains I_c to the parameters of I_s , from another site. This stained patch is re-stained with the staining parameters of I_c to get the reconstructed original patch. For WSI, we adapt $AP(i, p)$ to measure changes in boundaries within patches highlighted using Sobel filters [11]. For ideal stain normalization, $cycleL1$ be 0 and $AP(i, p)$ should be 1.

Tumor Classification: We used the area under the precision-recall (AUPR) curve and the area under the receiver operating characteristics (AUROC) curve as metrics. Additionally, we reported a $F1$ score using the best threshold from the precision-recall curve.

3 Results

3.1 Comparison with Other Stain Normalization Techniques

For *Data Alchemy*, the staining method must be controllable and capable of handling unseen stains during testing. So, we compare the performance of our proposed stain normalization method with HistAugAN [29], both quantitatively and qualitatively. Figure 3 shows examples of stain normalization on different patches. Both approaches reduce the color appearance variations and create plausible stained samples while preserving the general structure visible in the

original patches. However, HistAuGAN does not preserve the exact structures present in the original patch. It hallucinates additional nuclei and generates artifacts in the white background (Appendix C). In contrast, our method preserves structural details better without any hallucinations or artifacts.

Table 2 shows that our proposed stain normalization has a lower *cycleL1* error compared to HistAuGAN. Moreover, our proposed method performs better in terms of *SSIM* and *PSNR*. These, together with higher values of $AP(i, p)$ and the qualitative examples, suggest that our proposed method is better at preserving the structural information present in the original patch, hence a better choice for stain normalization. Therefore, the subsequent classification tasks are performed using our proposed stain normalization module.

Exploring Explainability in Stain Normalization: In Fig. 4, we show an example of the normalization of site A to site B. We can see that post normalization the higher eigenvalues of site A become smaller with lower eigenvalues of site B. Also, site A normalized to site B looks much closer to site B than to site A. We hypothesize that eigenvalues and vector manipulation are sufficient in stain blending, as shown in Fig. 5. We can control the blending of two patches

Table 2. Quantitative performance comparison between various stain normalization methods. The reported numbers are average \pm standard deviation.

	cycleL1	SSIM	PSNR	AP(i, p)
	A to B			
HistAuGAN	0.060 \pm 0.014	0.691 \pm 0.087	14.893 \pm 3.319	0.713 \pm 1.123
Ours	0.046 \pm 0.018	0.918 \pm 0.060	17.570 \pm 3.061	0.744 \pm 0.269
	B to A			
HistAuGAN	0.062 \pm 0.013	0.720 \pm 0.068	14.465 \pm 2.714	0.553 \pm 1.556
Ours	0.043 \pm 0.030	0.896 \pm 0.139	18.827 \pm 3.783	0.845 \pm 0.123




Fig. 4. Eigenvalue and vectors of a content, stain, and stain normalized patch projected onto a 3D sphere. Each point of the surface represents an eigenvector and the color represents an eigenvalue. Red signifies a smaller and blue higher eigenvalues. (Color figure online)




Fig. 5. Eigenvalue blending of the content and stain patches produces different staining. The (%) in the arrows are the content and stain patches blending weights.

directly by purely using eigenvalues and vectors from two different sites. By controlling the effect of the eigenvalues and vectors, we can control the staining to one particular site or the other. The manipulation of the eigenvalues helps us understand why one site is stained in a particular way compared to the other.

3.2 Stain Normalization on Downstream Task

Table 3 shows that the classifier performs best when trained on data from both sites A and B. The upper bound model (UBM) represents the classifier trained and tested on the same data location, while the lower bound model (LBM) shows the performance drop (0.394 *AUPR*) when testing on a site different from the training site. Stain normalization to a single template from site A improves classifier performance beyond the LBM when the classifier is trained on site A

Table 3. Quantitative performance of tumor and non-tumor classification of patches in various settings. UBM and LBM are expected Upper Bound Model and Lower Bound Model performance when trained with one dataset only. Train, Norm, and Temp refer to the train data site, normalization, and normalization template site, respectively.

	Norm	Testing on Site B					Testing on Site A				
		Train	Temp	AUPR	AUROC	F1	Train	Temp	AUPR	AUROC	F1
Best Case	No	A, B	–	0.939	0.904	0.850	A, B	–	0.939	0.912	0.878
UBM	No	B	–	0.916	0.888	0.840	A	–	0.930	0.881	0.848
LBM	No	A	–	0.545	0.525	0.771	B	–	0.841	0.774	0.787
Stain Normalization											
1 Template	Yes	A	A	0.710	0.704	0.792	B	B	0.648	0.578	0.771
10 Template (Ensemble)	Yes	A	A	0.672	0.689	0.799	B	B	0.686	0.630	0.773
1 Template (Test Data)	Yes	A	B	0.587	0.591	0.782	B	A	0.726	0.652	0.780
Data Alchemy: Test Time Data Calibration											
	Yes	A	Learnt	0.852	0.833	0.817	B	Learnt	0.928	0.890	0.856

and tested on site B. Using one or ten templates increases the *AUPR* scores by 0.165 and 0.127, respectively, over the LBM. When using a template from site B with a model trained on site A, we observe a negligible improvement of 0.042. This demonstrates that stain normalization improves classifier performance.

Table 3 also shows that in some scenarios, the LBM is close to the UBM, indicating that training on site B captures the necessary data diversity for good performance on site A. More visualizations of the phenomenon are in Appendix B. So When staining patches from site A to B, there is a drop of 0.193 and 0.155 *AUPR* using one or ten templates, respectively. A single template from site A only drops performance by 0.115. Overall, these findings suggest that static stain normalization may not always be beneficial for classifier performance.

3.3 Test Time Data Calibration

Since static stain normalization may not guarantee optimal performance and we cannot update the model parameters due to regulatory concerns, we apply *Data Alchemy* to the classifier. In Table 3, the classifier trained on site A and tested on site B, the learned template boosts performance by 0.307 *AUPR* over the LBM and is just 0.064 below the UBM. Additionally, the classifier trained on site B and tested on site A also improves performance by 0.087 *AUPR* over the LBM and is only 0.002 below the UBM. We also observe an improvement of 0.009 *AUPR* and 0.008 F1 score of the data-calibrated model over the UBM. This demonstrates that *Data Alchemy*'s learned template enhances classifier performance across different sites and has the potential to surpass the UBM.

4 Conclusion

We propose an effective and explainable stain normalization strategy that preserves the image structures and reduces stain variance between a template image and the original patch. Moreover, data calibration using *Data Alchemy* improves the classification accuracy without retraining of any kind. It serves as a step that enhances classifier generalizability, reducing the domain gap between multiple sites. Apart from easing regulatory approval hurdles, *Data Alchemy* may be used for onsite model weight calibration when it is difficult to access the model (e.g., API-based interaction) or update the model (e.g., black boxes that do not support retraining or continuous learning).

Acknowledgements. This work was supported by the National Cancer Institute award UG3CA236536.

References

1. Anaya-Isaza, A., Mera-Jiménez, L., Zequera-Diaz, M.: An overview of deep learning in medical imaging. *Inform. Med. Unlock.* **26**, 100723 (2021)

2. Bejnordi, B.E., et al.: Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* **318**(22), 2199–2210 (2017)
3. Cheng, J.Y., Abel, J.T., Balis, U.G., McClintock, D.S., Pantanowitz, L.: Challenges in the development, deployment, and regulation of artificial intelligence in anatomic pathology. *Am. J. Pathol.* **191**(10), 1684–1692 (2021)
4. Center for devices and radiological health, US FDA, artificial intelligence and machine learning (AI/ml)-enabled medical device, December 2023. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>
5. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2414–2423 (2016). <https://doi.org/10.1109/CVPR.2016.265>
6. Haub, P., Meckel, T.: A model based survey of colour deconvolution in diagnostic brightfield microscopy: error estimation and spectral consideration. *Sci. Rep.* **5**, 12096 (2015). <https://doi.org/10.1038/srep12096>
7. He, M., Li, Z., Liu, C., Shi, D., Tan, Z.: Deployment of artificial intelligence in real-world practice: opportunity and challenge. *Asia-Pacific J. Ophthalmol.* **9**(4), 299–307 (2020)
8. Hoque, M.Z., Keskinarkaus, A., Nyberg, P., Seppänen, T.: Stain normalization methods for histopathology image analysis: a comprehensive review and experimental comparison. *Inf. Fusion* **102**, 101997 (2024). <https://doi.org/10.1016/j.inffus.2023.101997>
9. Hossain, M.: Whitening and coloring transformations for multivariate Gaussian data. In: A Selecture Partly Based on the ECE662 Spring (2014)
10. Kang, H., et al.: Stainnet: a fast and robust stain normalization network. *Front. Med.* **8** (2021). <https://doi.org/10.3389/fmed.2021.746307>
11. Kanopoulos, N., Vasanthavada, N., Baker, R.: Design of an image edge detection filter using the sobel operator. *IEEE J. Solid-State Circuits* **23**(2), 358–367 (1988). <https://doi.org/10.1109/4.996>
12. Laves, M.H., Ihler, S., Fast, J.F., Kahrs, L.A., Ortmaier, T.: Well-calibrated regression uncertainty in medical imaging with deep learning. In: *Medical Imaging with Deep Learning*, pp. 393–412. PMLR (2020)
13. Li, Y., Ping, W.: Cancer metastasis detection with neural conditional random field. In: *Medical Imaging with Deep Learning* (2018)
14. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. *Adv. Neural Inf. Process. Syst.* **30** (2017)
15. Macenko, M., et al.: A method for normalizing histology slides for quantitative analysis. In: 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1107–1110 (2009). <https://doi.org/10.1109/ISBI.2009.5193250>
16. Parida, A., Jiang, Z., Packer, R.J., Avery, R.A., Anwar, S.M., Linguraru, M.G.: Quantitative metrics for benchmarking medical image harmonization. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI), pp. 1–5 (2024). <https://doi.org/10.1109/ISBI56570.2024.10635289>
17. Rabinovich, A., Agarwal, S., Laris, C., Price, J., Belongie, S.: Unsupervised color decomposition of histologically stained tissue samples. In: Thrun, S., Saul, L., Schölkopf, B. (eds.) *Advances in Neural Information Processing Systems*, vol. 16. MIT Press (2003)
18. Rajaraman, S., Ganesan, P., Antani, S.: Deep learning model calibration for improving performance in class-imbalanced medical image classification tasks. *PLoS ONE* **17**(1), e0262838 (2022)

19. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Comput. Graphics Appl.* **21**(5), 34–41 (2001). <https://doi.org/10.1109/38.946629>
20. Runz, M., Rusche, D., Schmidt, S., Weihrauch, M., Hesser, J., Weis, C.A.: Normalization of he-stained histological images using cycle consistent generative adversarial networks. *Diagnost. Pathol.* **16**, 71 (2021). <https://doi.org/10.1186/s13000-021-01126-y>
21. Sahiner, B., et al.: Deep learning in medical imaging and radiation therapy. *Med. Phys.* **46**(1), e1–e36 (2019)
22. Salvi, M., Acharya, U.R., Molinari, F., Meiburger, K.M.: The impact of pre- and post-image processing techniques on deep learning frameworks: a comprehensive review for digital pathology image analysis. *Comput. Biol. Med.* **128**, 104129 (2021)
23. Salvi, M., et al.: Impact of stain normalization on pathologist assessment of prostate cancer: a comparative study. *Cancers* **15**(5), 1503 (2023). <https://doi.org/10.3390/cancers15051503>
24. Setiadi, D.R.I.M.: PSNR vs SSIM: imperceptibility quality assessment for image steganography. *Multimedia Tools Appl.* **80**(6), 8423–8444 (2021)
25. Shaban, M.T., Baur, C., Navab, N., Albarqouni, S.: Staingan: stain style transfer for digital histological images. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), pp. 953–956. IEEE (2019)
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
27. Tellez, D., et al.: Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Med. Image Anal.* **58**, 101544 (2019). <https://api.semanticscholar.org/CorpusID:62841444>
28. Voon, W., et al.: Evaluating the effectiveness of stain normalization techniques in automated grading of invasive ductal carcinoma histopathological images. *Sci. Reports* **13**, 20518 (2023). <https://doi.org/10.1038/s41598-023-46619-6>
29. Wagner, S.J., et al.: Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12908, pp. 257–266. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87237-3_25
30. Zheng, Y., et al.: Histopathological whole slide image analysis using context-based CBIR. *IEEE Trans. Med. Imaging* **37**(7), 1641–1652 (2018). <https://doi.org/10.1109/TMI.2018.2796130>