



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Pathological Semantics-Preserving Learning for H&E-to-IHC Virtual Staining

Fuqiang Chen^{1,2}, Ranran Zhang¹, Boyun Zheng^{1,2}, Yiwen Sun^{1,3}, Jiahui He^{1,4}, and Wenjian Qin¹

¹ Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China

wj.qin@siat.ac.cn

² Shenzhen College of Advanced Technology, University of Chinese Academy of Sciences, Shenzhen, China

³ Department of Health Technology and Informatics, The Hong Kong Polytechnic University, Hong Kong SAR, China

⁴ School of Computer Science, University of Nottingham Ningbo China, Ningbo, China

Abstract. Conventional hematoxylin-eosin (H&E) staining is limited to revealing cell morphology and distribution, whereas immunohistochemical (IHC) staining provides precise and specific visualization of protein activation at the molecular level. Virtual staining technology has emerged as a solution for highly efficient IHC examination, which directly transforms H&E-stained images to IHC-stained images. However, virtual staining is challenged by the insufficient mining of pathological semantics and the spatial misalignment of pathological semantics. To address these issues, we propose the Pathological Semantics-Preserving Learning method for Virtual Staining (PSPStain), which directly incorporates the molecular-level semantic information and enhances semantics interaction despite any spatial inconsistency. Specifically, PSPStain comprises two novel learning strategies: 1) Protein-Aware Learning Strategy (PALS) with Focal Optical Density (FOD) map maintains the coherence of protein expression level, which represents molecular-level semantic information; 2) Prototype-Consistent Learning Strategy (PCLS), which enhances cross-image semantic interaction by prototypical consistency learning. We evaluate PSPStain on two public datasets using five metrics: three clinically relevant metrics and two for image quality. Extensive experiments indicate that PSPStain outperforms current state-of-the-art H&E-to-IHC virtual staining methods and demonstrates a high pathological correlation between the staging of real and virtual stains. Code is available at <https://github.com/ccitachi/PSPStain>.

Keywords: Semantics preserving · Protein awareness · Prototype consistency · Virtual stain.

1 Introduction

With the increasing demand for labor-intensive and time-consuming IHC examination [1], virtual staining has emerged as a viable solution, which directly

generates IHC-stained pathological images from conventional H&E-stained images. This innovative approach is expected to significantly improve the efficiency of cancer diagnosis. However, the inherent challenge of virtual staining is the lack of aligned ground truth (GT) pairs for training. Generally, the virtual staining GT pair is obtained from two depth-wise consecutive cuts of the same tissue and stained separately. This inevitably prevents pixel-perfect image correspondences due to the changes in cell morphology, and staining-induced degradation [7].




Fig. 1. Two key problems in H&E-to-IHC virtual staining: (a) Protein expression level varies greatly in each sample from the same-grade WSI, and only preserving grade information rather than protein expression level results in a less robust representation with semantic dispersion. (b) Spatial misalignment between the generated image and label causes the false response of patches with similar semantics.

Existing virtual staining methods can be regarded as a progressive process for extracting pathological semantic information in inconsistent GT pairs, where the most critical pathological semantics are those at the molecular level. Initially, GAN-based algorithms are directly applied in virtual staining without additional pathological constraint [6,10,19]. Recent algorithms utilize grade-level semantic information [15]. Subsequently, the patch-level semantic information is incorporated to enhance the molecular-level semantic preserving [8,16,7]. Moreover, the pixel-level semantic information is employed with the semantic mask [9].

However, rethinking the existing methods and challenges of H&E-to-IHC virtual staining, there are two significant problems (Fig. 1). **(a) Insufficient Mining of Pathological Semantics:** The level of protein expression, which is the representation of the molecular-level semantic information, has not been directly extracted for pathological constraints. For instance, the protein expression level of human epidermal growth factor receptor-2 (HER2) varies greatly in the samples within the 3+-grade whole slide images (WSI). By only preserving grade-level pathological information rather than protein expression level for learning, the molecular-level pathological semantics become dispersed, resulting in a less robust representation. **(b) Spatial Misalignment of Pathological Semantics:** The input (H&E) and label (IHC) inevitably exhibit spatial inconsistency, implying that corresponding patches in the paired images contain different normal and tumor cells. Once directly assuming that similar semantics exist in corresponding patches for training, representations with similar seman-

tics may falsely separate, and those with opposite semantics may falsely group. This interference hinders the learning of pathological representations.

To address the above problems, we propose a novel Pathological Semantics-Preserving Learning method for H&E-to-IHC Virtual Staining (PSPStain). PSPStain comprises two novel learning strategies to solve the corresponding problems. For problem (a), we directly quantify the protein expression of each IHC-stained image by determining the optical density in the DAB channel [12], preserving the molecular-level pathological semantics. Additionally, we introduce a novel Focal Optical Density (FOD) map to re-weight the contribution of tumor and non-tumor areas during quantification. For problem (b), we make the assumption that the pathological content of generated image needs to be consistent with label. To enhance pathological semantic interaction despite any spatial inconsistency, we facilitate the convergence of pathological features towards cross-corresponding prototypes in the both generated image and label. During the verification process, we empirically observed that metrics such as PSNR and SSIM do not always strictly correlate with high-quality virtual staining results [4]. This led us to incorporate three pathological relevant metrics.

The main contributions are as follows:

1) We propose a Protein-Aware Learning Strategy (PALS) with a Focal Optical Density (FOD) map to extract molecular-level pathological semantics, constraining the protein expression level between the generated image and label.

2) We propose a Prototype-Consistent Learning Strategy (PCLS), which establishes a prototypical correlation between the generated tumor and the label.

3) Extensive experiments on BCI and MIST-her2 datasets have demonstrated that PSPStain effectively preserves pathological semantics and improves staining performance without additional annotation.

2 Methods

Our PSPStain (Fig. 2) contains two learning strategies: PALS, aimed at preserving the consistency of protein expression level, and PCLS, focus on enhancing tumor prototype alignment. In this framework, given an input pair consisting of an H&E image $I \in \mathbb{R}^{H \times W \times C}$ and its corresponding label (real IHC) $K^R \in \mathbb{R}^{H \times W \times C}$. The backbone produces a generated image (fake IHC) $K^F \in \mathbb{R}^{H \times W \times C}$.

2.1 Protein-Aware Learning Strategy

In most instances, protein as an antigen is conjugated to an enzyme, such as peroxidase [3]. Among the stains employed for IHC images, the DAB stain yields an intense brown staining when it reacts with the enzyme. Therefore, the fundamental concept behind PALS is to be accurately aware of molecular-level information, specifically protein expression level, by focusing on the DAB channel.

Optical Density Determination: Before presenting PALS, we introduce the optical density (OD). Each stain is characterized by an absorption factor to RGB.

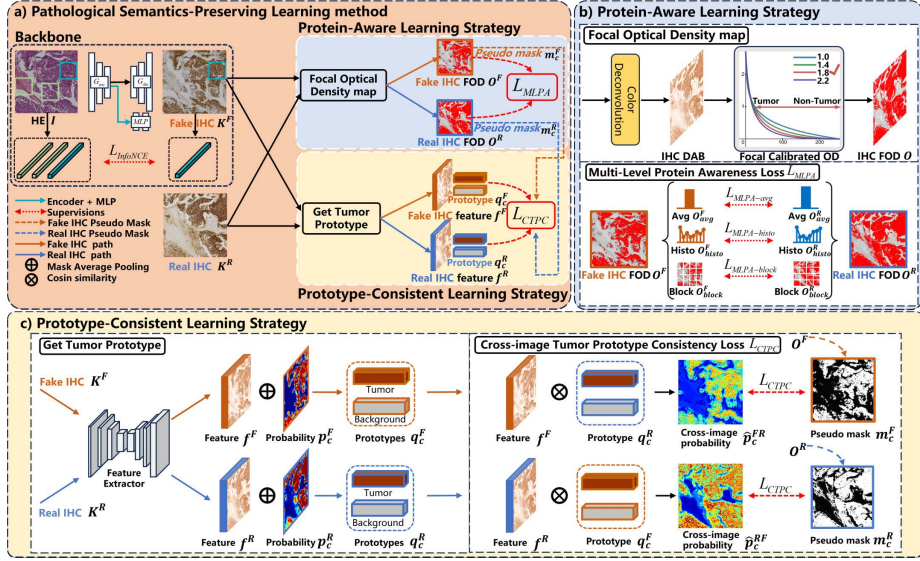


Fig. 2. Our proposed framework consists of two newly learning strategies for preserving protein expression consistency and enhancing semantic alignment.

Since the OD is proportional to the concentration of the stain, the amount of stain is the factor determining the OD at a wavelength as per the LambertBeer law [12].

$$OD_C = -\log_{10}(I_C/I_{0,C}) = A * c_C \quad (1)$$

where $I_{0,C}$ and I_C denote the light intensity entering and passing the specimen. Subscript C denotes the channel. A is stain amount with absorption factor c .

Focal Optical Density map: In this module, we first use the traditional color deconvolution [11] for stain separation. Then we specifically select DAB stain’s OD values to generate the RGB image (IHC DAB).

In most samples, the tumor area is significantly smaller than the non-tumor area, resulting in a notable imbalance. Hence, training with just the normal OD map Eq.1 becomes inefficient as the majority of locations represent easy non-tumor areas, contributing no useful learning signal. Instead, we propose to reshape the map function to down-weight non-tumor area and meanwhile focus training on the tumor area. For FOD, we simulate Eq.1 by converting the IHC DAB to grayscale and using the focal calibrated map to assign gray values to positive signal (FOD) as followed:

$$O_C = (-\log_{10}((I_C)/I_{0,C}))^\alpha \quad (2)$$

where O is the FOD with tunable focusing parameter $\alpha > 1$. The FOD map is visualized for several values of α in Fig.2(b). With an increase in the focusing parameter α , the de-emphasis of non-tumor areas and the emphasis on tumor areas are enhanced. Among the tested values, $\alpha = 1.8$ yields the best results.

Multi-Level Protein Awareness: For pathology consistency, a Multi-Level Protein Awareness (MLPA) loss is adopted to match fake IHC FOD O^F with real IHC FOD O^R .

Initially, we align the average protein expression level and set a tolerance value β of 0.2 to account for inherent differences. If the difference in protein expression level is less than β multiplied by the average expression level in reference image O_{avg}^R , there is no contribution to model learning.

$$L_{MLPA-avg} = \begin{cases} \|O_{avg}^F - O_{avg}^R\|_2, & \text{if } |O_{avg}^F - O_{avg}^R| \geq \beta \cdot O_{avg}^R \\ 0, & \text{if } |O_{avg}^F - O_{avg}^R| < \beta \cdot O_{avg}^R \end{cases} \quad (3)$$

The O_{avg} denotes the average amount of IHC FOD O , $O_{avg} = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W O_{h,w}$.

Secondly, we divide the FOD range from 0 to e into intervals of $N_h = 20$ to ensure consistency in the distribution of protein expression between generated images and labels.

$$L_{MLPA-histo} = \frac{1}{N_h} \sum_{i=1}^{N_h} \|O_{histo_i}^F - O_{histo_i}^R\|_2 \quad (4)$$

where O_{histo_i} denotes the i_{th} histogram accumulation of the IHC FOD O . $O_{histo_i} = \sum_{j=1}^{M_h} O_j$, if $\frac{(i-1) \cdot e}{N_h} < O_j \leq \frac{i \cdot e}{N_h}$, where O_j denotes the j_{th} pixel of IHC FOD O and M_h is the number of pixels in the i_{th} histogram.

Finally, we divide the image into $N_b = 16$ blocks and independently calculate the average expression amount of each block. This preserves the intensity of expression in specific regions as regularization, referred to as *MLPA-block*.

$$L_{MLPA-block} = \frac{1}{N_b} \sum_{i=1}^{N_b} \|O_{block_i}^F - O_{block_i}^R\|_2 \quad (5)$$

where O_{block_i} denotes the i_{th} block average amount of the IHC FOD O . $O_{block_i} = \frac{1}{M_b} \sum_{j=1}^{M_b} O_j$, if $position(O_j) \in block_i$, where O_j denotes the j_{th} pixel of IHC FOD O and M_b is the number of pixels in the i_{th} block.

Thus, the overall objective of our MLPA is formulated as follows:

$$L_{MLPA} = L_{MLPA-avg} + L_{MLPA-histo} + L_{MLPA-block} \quad (6)$$

2.2 Prototype-Consistent Learning Strategy

Inspired by [17], PCLS effectively preserves pathological semantic consistency even in the presence of spatial misalignment. Specifically, when comparing the generated image K^F with the label K^R , it is highly likely that the tumor content within these images remains consistent. Therefore, we calculate the prototype consistency loss L_{CTPC} between them. We derive masks m_c^F and m_c^R by applying the same threshold to fake IHC FOD O^F and real IHC FOD O^R .

Get Tumor Prototype: By defining the tumor class and non-tumor class in the IHC image, the prototype of each class refers to the aggregated representation. We pretrain a segmentation UNet with the pseudo mask m_c^R on real IHC images and freeze it to extract tumor semantic. Initially, input IHC image is processed by UNet, yielding a seg probability map as output. Feature maps are then derived from the layer preceding the output layer. Let $f^F \in \mathbb{R}^{H \times W \times D}$ represents the feature map of output K^F , while $p_c^F(i)$ denotes the probability of pixel i belonging to class c . Similarly, $f^R \in \mathbb{R}^{H \times W \times D}$ represents the feature map of label K^R , and $p_c^R(i)$ denotes the probability of pixel i belonging to class c . Using these feature maps and probability maps, we aggregate class-wise prototypes representing pixel-wise features of tumor and background. The prototypes of output q_c^F and label q_c^R are as followed:

$$q_c^F = \frac{\sum_i p_c^F(i) \cdot f^F(i)}{\sum_i p_c^F(i)}, q_c^R = \frac{\sum_i p_c^R(i) \cdot f^R(i)}{\sum_i p_c^R(i)} \quad (7)$$

Cross-image Tumor Prototype Consistency: The feature similarity is calculated based on the output prototype q_c^F and the label prototype q_c^R . For semantic interaction, cosine similarity is computed leveraging image features with prototypes from another image. Specifically, \hat{s}_c^{FR} and \hat{s}_c^{RF} are two cross-image prototypical similarity maps obtained by calculating cosine similarity between the feature map f^F , f^R and the prototype vector q_c^R , q_c^F . These are defined as:

$$\hat{s}_c^{FR} = \frac{f^F \cdot q_c^R}{\|f^F\| \cdot \|q_c^R\|}, \hat{s}_c^{RF} = \frac{f^R \cdot q_c^F}{\|f^R\| \cdot \|q_c^F\|} \quad (8)$$

Then, we apply softmax operations to \hat{s}_c^{FR} and \hat{s}_c^{RF} to derive the corresponding cross-image probability prediction \hat{p}_c^{FR} and \hat{p}_c^{RF} respectively:

$$\hat{p}_c^{FR} = \frac{e^{\hat{s}_c^{FR}}}{\sum_c e^{\hat{s}_c^{FR}}}, \hat{p}_c^{RF} = \frac{e^{\hat{s}_c^{RF}}}{\sum_c e^{\hat{s}_c^{RF}}} \quad (9)$$

The cross-image tumor prototype consistency (CTPC) loss is defined as:

$$L_{CTPC} = \frac{1}{C \times H \times W} \sum_{i=1}^{H \times W} \sum_{c=1}^C (\|\hat{p}_c^{FR}(i) - m_c^F(i)\|_2 + \|\hat{p}_c^{RF}(i) - m_c^R(i)\|_2) \quad (10)$$

2.3 Loss Function of PSPStain.

The overall learning objective is as follows:

$$L_{total} = L_{adv} + L_{NCE} + \lambda_M L_{MLPA} + \lambda_C L_{CTPC} + \lambda_S L_{SSIM} + \lambda_G L_{GP} \quad (11)$$

where L_{MLPA} , L_{CTPC} , L_{GP} focus on pathological consistency, with L_{GP} originating from [8]. L_{adv} , L_{NCE} , L_{SSIM} contribute to image quality enhancement.

3 Experiments

Datasets: Experiments are conducted on two public datasets: the Breast Cancer Immunohistochemical (BCI) challenge dataset [8] and the Multi-IHC Stain Translation (MIST) dataset [7]. The BCI dataset consists of 3396 H&E-HER2 pairs of the training set images and 977 pairs of test set images from 51 WSIs. In the MIST dataset, we use the MIST_{HER2}, which contains 4642 paired samples for training and 1000 for testing from 64 WSIs.

Implementation Details: CUT [10] is selected as the baseline model. The generator is ResNet-6Blocks [5] and the discriminator is PatchGAN [6]. We trained our networks with random 512×512 crops and a batch size of four. Adam optimizer with the learning rate of 1×10^{-4} . The maximum number of training epochs was set to 80. The weight value of $\lambda_M, \lambda_C, \lambda_S, \lambda_G$ is 1.0, 2.5, 0.05, 10.0.

Table 1. Quantitative evaluations on two datasets using three pathology-related metrics and two image quality metrics. The mIOD and IOD are subtracted from the GT and the value closer to 0 indicates better results. The **best** values are highlighted.

Method	BCI-Her2					MIST-Her2				
	Pathological Relevance			Image Quality		Pathological Relevance			Image Quality	
	mIOD	IOD \times_{10^7}	Peason-R \uparrow	PSNR \uparrow	SSIM \uparrow	mIOD	IOD \times_{10^7}	Peason-R \uparrow	PSNR \uparrow	SSIM \uparrow
CUT [10]	-0.3528	-2.9126	0.0276	18.1246	0.4483	-0.0830	-2.9478	0.7164	13.8998	0.1680
Pix2Pix [6]	-0.3598	-2.8819	0.0431	15.8659	0.4263	-0.0414	-2.5782	0.2986	12.9840	0.1676
PyramidP2P [8]	-0.3105	-2.9056	0.1018	19.9488	0.4647	-0.0446	-4.5285	0.6894	14.9122	0.1995
ASP [7]	-0.2438	-2.9292	-0.0406	17.8651	0.4923	-0.1303	-5.7422	0.0659	14.1841	0.2004
PSPStain	-0.1800	-0.7808	0.7553	18.6220	0.4498	-0.0834	-2.5491	0.8303	14.1948	0.1876




Fig. 3. Qualitative comparison with samples virtually stained by various methods.

Evaluations: To evaluate image quality, we use the standard Peak Signal-to-Noise Ratio (PSNR) [2] and Structural Similarity Index Measure (SSIM) [14].




Fig. 4. The abscissa represents the index number of samples. The ordinate is the accumulated IOD value. **(a)** **(b)** refer to the curve on BCI and MIST datasets, respectively. **(c)** Strategy ablation. **(d)** FOD factor α ablation. **(e)** MLPA module ablation

Table 2. Ablation studies on Strategy, FOD factor, and MLPA module. All ablation study is performed on the BCI dataset. The **best** values are highlighted.

Method		BCI-Her2				
		Pathological Relevance			Image Quality	
		mIOD	IOD $\times 10^7$	Pearson-R \uparrow	PSNR \uparrow	SSIM \uparrow
Strategy ablation	Baseline	-0.3524	-2.9551	0.3611	21.1427	0.5043
	Baseline+PALS	-0.2894	-2.3197	0.6197	19.7659	0.4826
	Baseline+PALS+PCLS	-0.1800	-0.7808	0.7553	18.6220	0.4498
FOD factor α ablation	1.0	-0.2108	-2.2525	0.5884	19.5469	0.4419
	1.4	-0.2091	-1.4506	0.7223	19.3400	0.4366
	1.8	-0.1800	-0.7808	0.7553	18.6220	0.4498
	2.2	-0.2516	-2.2563	0.5630	19.5217	0.4656
MLPA module ablation	w/o MLPA-avg	-0.0687	-1.3915	0.0474	19.0137	0.4353
	w/o MLPA-histo	-0.2286	-1.9272	0.6972	19.1779	0.4338
	w/o MLPA-block	-0.1616	-1.9022	0.6728	17.0345	0.4080
	PSPStain	-0.1800	-0.7808	0.7553	18.6220	0.4498

To explore whether virtual staining meets clinical needs, we evaluate the positive signal prediction using integrated optical density (IOD) and mean integrated optical density (mIOD) [13,16], and calculate the Pearson correlation coefficient (Pearson-R) [9]. Specifically, IOD quantifies the amount of positive signal statistically, while mIOD represents the intensity of positive signal within the positive area. Pearson-R is used to assess the pathological correlation with protein expression level rather than just the size of the positive area as discussed in [9].

Comparison Study: We compare the performance with state-of-the-art methods. Empirical results demonstrate the superiority of our method. Table. 1 shows that PSPStain achieves the lowest deviation of mIOD and IOD and the highest Pearson-R score from the labels, indicating the highest pathological consistency. The higher PSNR and SSIM indicate higher image quality however not always in virtual staining. Specifically, the blurry IHC images with more averaged pixel value may cause them to be increased within spatial misalignment [18]. Fig. 3 shows the qualitative results, indicating that PSPStain effectively and robustly highlights tumor regions. Fig. 4(a)(b) compares the performance of different methods in two datasets by accumulating positive signals in each sample, revealing that PSPStain closely matches the ensemble of GT more than others.

Ablation Study: As demonstrated in Table. 2 and Fig. 4(c)(d)(e), we conduct the ablation study using the BCI dataset. (c) For strategy ablation, both learning strategies effectively improve the performance of virtual staining. (d) For FOD

ablation, we test various values of α . When α is set to 1.8, PSPStain performs the best. (e) For MLPA ablation, all three modules prove to be beneficial.

4 Conclusion

In this paper, we propose PSPStain which aims to extract molecular-level pathological information and align cross-image pathological semantics despite any spatial inconsistency. Specifically, protein expression level as the representation of molecular-level semantics is preserved by PALS. The FOD map, integrated into PALS, focuses training on the tumor rather than the non-tumor area. PCLS maintains semantic alignment via cross-image prototypical consistency learning. The superior virtual staining performance shows that PSPStain effectively and robustly preserves pathological semantics without additional expert annotation. Our method has been validated on the Her2-image dataset and can be extended to other DAB-stained IHC images.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (No. 62271475), Ministry of Science and Technology’s key research and development program (2023YFF0723400), the Youth Innovation Promotion Association CAS (2022365) and Shenzhen-Hong Kong Joint Lab on Intelligence Computational Analysis for Tumor Imaging (E3G111).

Disclosure of Interests. The authors declare no competing interests in the paper.

References

1. Anglade, F., Milner Jr, D.A., Brock, J.E.: Can pathology diagnostic services for cancer be stratified and serve global health? *Cancer* **126**(S10), 2431–2438 (2020)
2. Avciabas, I.s., Sankur, B.I., Sayood, K.: Statistical evaluation of image quality measures. *Journal of Electronic Imaging* **11**(2), 206–223 (2002)
3. Di Cataldo, S., Ficarra, E., Macii, E.: Computer-aided techniques for chromogenic immunohistochemistry: status and directions. *Computers in biology and medicine* **42**(10), 1012–1025 (2012)
4. Dubey, S., Kataria, T., Knudsen, B., Elhabian, S.Y.: Structural cycle gan for virtual immunohistochemistry staining of gland markers in the colon. In: *Machine Learning in Medical Imaging: 14th International Workshop, MLMI 2023, Held in Conjunction with MICCAI 2023, Vancouver, BC, Canada, October 8, 2023, Proceedings, Part II*. pp. 447–456. Springer (2023)
5. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
6. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1125–1134 (2017)
7. Li, F., Hu, Z., Chen, W., Kak, A.: Adaptive supervised patchnce loss for learning h&e-to-ihc stain translation with inconsistent groundtruth image pairs. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023: 26th International Conference, Vancouver, BC, Canada, October 8–12, 2023, Proceedings, Part VI*. pp. 632–641. Springer (2023)

8. Liu, S., Zhu, C., Xu, F., Jia, X., Shi, Z., Jin, M.: Bci: Breast cancer immunohistochemical image generation through pyramid pix2pix. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1815–1824 (2022)
9. Liu, S., Zhang, B., Liu, Y., Han, A., Shi, H., Guan, T., He, Y.: Unpaired stain transfer using pathology-consistent constrained generative adversarial networks. *IEEE Transactions on Medical Imaging* **40**(8), 1977–1989 (2021)
10. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16. pp. 319–345. Springer (2020)
11. Ruifrok, A.C., Johnston, D.A., et al.: Quantification of histochemical staining by color deconvolution. *Analytical and quantitative cytology and histology* **23**(4), 291–299 (2001)
12. Varghese, F., Bukhari, A.B., Malhotra, R., De, A.: Ihc profiler: an open source plugin for the quantitative evaluation and automated scoring of immunohistochemistry images of human tissue samples. *PloS one* **9**(5), e96801 (2014)
13. Wang, C.J., Zhou, Z.G., Holmqvist, A., Zhang, H., Li, Y., Adell, G., Sun, X.F.: Survivin expression quantified by image pro-plus compared with visual assessment. *Applied Immunohistochemistry & Molecular Morphology* **17**(6), 530–535 (2009)
14. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
15. Zeng, B., Lin, Y., Wang, Y., Chen, Y., Dong, J., Li, X., Zhang, Y.: Semi-supervised pr virtual staining for breast histopathological images. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II. pp. 232–241. Springer (2022)
16. Zhang, R., Cao, Y., Li, Y., Liu, Z., Wang, J., He, J., Zhang, C., Sui, X., Zhang, P., Cui, L., et al.: Mvfstain: multiple virtual functional stain histopathology images generation based on specific domain mapping. *Medical Image Analysis* **80**, 102520 (2022)
17. Zhang, Z., Ran, R., Tian, C., Zhou, H., Li, X., Yang, F., Jiao, Z.: Self-aware and cross-sample prototypical learning for semi-supervised medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023: 26th International Conference, Vancouver, BC, Canada, October 8–12, 2023, Proceedings, Part II. pp. 192–201. Springer (2023)
18. Zhu, C., Liu, S., Yu, Z., Xu, F., Aggarwal, A., Corredor, G., Madabhushi, A., Qu, Q., Fan, H., Li, F., et al.: Breast cancer immunohistochemical image generation: a benchmark dataset and challenge review. *arXiv preprint arXiv:2305.03546* (2023)
19. 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/CVF International Conference on Computer Vision. pp. 2223–2232 (2017)