IgG4-QPC: Quantitative Pathology-Consistent Virtual Staining from H&E to IgG4

Mengyue Luo[1], Dawen Wu[2], Yuxuan Ji[1], Mingxuan Tian[1], and Haixian Zhang[1] (✉)

[1] Machine Intelligence Lab, School of Computer Science and Technology, Sichuan University, Chengdu, Sichuan, China zhanghaixian@scu.edu.cn
[2] Department of Ophthalmology, West China Hospital, Sichuan University, Chengdu, Sichuan, China

Abstract. IgG4 immunohistochemical (IHC) staining plays a critical role in pathological diagnosis, where the quantitative assessment of IgG4-positive plasma cells—specifically their density and spatial distribution—serves as a pivotal pathological criterion for evaluating conditions like IgG4-related disease (IgG4-RD). H&E-to-IHC virtual staining offers a promising low-cost alternative to conventional immunostaining. However, most existing methods focus on global visual realism or continuous tissue textures (e.g., tumor regions), failing to capture the discrete and sparse diagnostic cell patterns required for IgG4 assessment. Moreover, real-world IHC images often contain non-specific staining (NSS), introducing noisy supervision that can be erroneously learned by generative models and compromise pathological consistency. Most existing virtual staining methods assume clean supervision or depend on strict data filtering, rather than explicitly accounting for NSS. To address these challenges, we propose IgG4-QPC, a virtual staining method that captures discrete IgG4-positive cell patterns from H&E images to achieve quantitative pathological consistency. Specifically, we introduce a Segmentation-Guided Instance-Level Alignment (SGIA) module, utilizing a segmentation prior to align cell-level features and enforce spatial density consistency, thereby enhancing sensitivity to diagnostically relevant plasma cells. Furthermore, we design a Noise-Aware Adversarial Learning strategy that identifies NSS artifacts and employs spatially-adaptive mask weighting to suppress background noise during optimization. Experimental results on paired H&E-IgG4 datasets demonstrate that our method significantly outperforms state-of-the-art approaches in image quality, cell density consistency, and spatial correlation.

Keywords: H&E-to-IHC Virtual Staining· IgG4 Quantitative Pathology · Cell Instance Consistency · Non-Specific Staining

1 Introduction

Immunohistochemistry (IHC) staining plays a critical role in pathological diagnosis by revealing functional and molecular information beyond morphology [14].

IHC staining for Immunoglobulin G4 (IgG4) [18] is essential to evaluate various inflammatory diseases. However, traditional IHC staining is time-consuming, expensive, and inevitably consumes tissue samples and laboratory resources [11,12]. H&E-to-IHC virtual staining based on generative models provides a promising and cost-effective alternative [1,16]. Nevertheless, in the context of IgG4-related disease (IgG4-RD), a systemic immune-mediated disease, diagnosis relies heavily on quantitative histopathological assessment of IgG4 staining. Specifically, the density and spatial distribution of IgG4-positive plasma cells serve as key pathological criteria for disease confirmation [8]. Consequently, preserving instance-level quantitative consistency of diagnostically relevant plasma cells is critical for reliable IgG4 virtual staining.

Virtual staining models are predominantly based on Generative Adversarial Networks (GANs). Early GAN-based approaches, such as CycleGAN and CUT [7,10,13,23], primarily focused on achieving visual similarity in terms of overall appearance(e.g., general tissue and cellular morphology) and staining style. However, unlike natural image translation, virtual staining places greater emphasis on preserving diagnostically relevant pathological structures [15]. Recent studies have attempted to integrate clinical knowledge, such as membrane staining intensity or protein expression levels, to guide generation [4,15]. Nevertheless, these methods are primarily designed to enforce consistency at the tissue or region level, where morphological patterns are relatively continuous. In contrast, IgG4 evaluation fundamentally depends on identifying and quantifying sparse and discrete plasma cell instances. This discrepancy in analytical granularity—from modeling general tissue or cellular morphology and continuous pathological regions to quantifying sparse, diagnostically relevant cellular instances—has received limited attention in existing virtual staining studies.

Achieving quantitative consistency is further complicated by the quality of supervision data. Real-world IHC slides frequently suffer from non-specific staining (NSS) artifacts [3], whereas most existing virtual staining models implicitly assume clean supervision or rely on strict data curation. As illustrated in Fig. 1, these artifacts (red boxes) manifest as diffuse background precipitates, which can sometimes visually resemble positive cells. For H&E-to-IgG4 virtual staining, such noisy supervision can mislead generative models to reproduce background artifacts as pathological signals, resulting in false-positive plasma cells. These effects directly undermine pathological consistency, limiting the clinical reliability of virtual staining.

To address these challenges, we propose IgG4-QPC to capture discrete IgG4-positive cell patterns and achieve quantitative pathological consistency under noisy supervision. Specifically, we introduce a Segmentation-Guided Instance-Level Alignment module(SGIA) module. Incorporating a pre-trained positive cell segmentation network as a semantic prior, SGIA explicitly aligns generator features with cell-level semantics and enforces density consistency, thereby enhancing sensitivity to plasma cells. Furthermore, to mitigate NSS interference, we design a Noise-Aware Adversarial Learning strategy. By identifying NSS to construct refined supervision signals and introducing a spatially adaptive

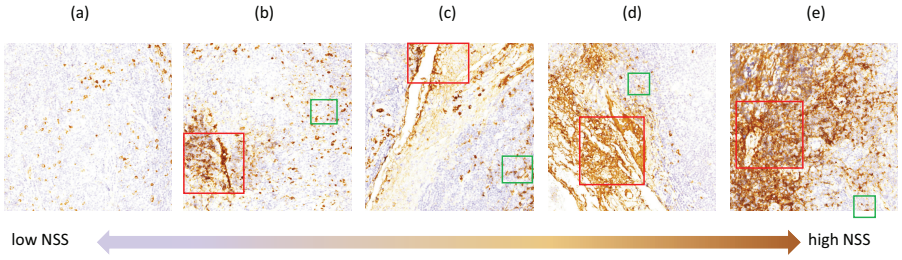

Fig. 1. Non-specific staining (NSS) artifacts in IgG4 staining with increasing severity from (a) to (e). Green boxes highlight true positive plasma cells, while red boxes indicate NSS artifacts.

re-weighting mechanism, the proposed approach suppresses the contribution of background noise during adversarial training, effectively inhibiting false-positive generation.

Overall, our primary contributions can be summarized as follows:

1. We frame the pathological consistency of IgG4 virtual staining as a quantitative generation problem of discrete diagnostic cell instances, and propose a tailored solution under noisy supervision.
2. We propose a segmentation-guided instance-level alignment module (SGIA) that explicitly injects cell-level semantic priors into the virtual staining generation process.
3. We design a noise-aware adversarial learning scheme that effectively disentangles true pathological signals from staining artifacts, improving robustness and reliability under noisy labels.

2 Methods

2.1 Problem Definition and Overview

Given an H&E image $x \in X \subset \mathbb{R}^{H\times W\times C}$, and its paired IgG4 IHC image $y \in Y \subset \mathbb{R}^{H\times W\times C}$, our goal is to learn a virtual staining mapping $G : X \to Y$ that generates a virtual IgG4 IHC image $\hat{y} = G(x)$, achieving quantitative pathological consistency at the level of diagnostically relevant plasma cell instances while preserving tissue morphology.

As illustrated in Fig. 2, our proposed framework incorporates two core strategies: (a) Segmentation-Guided Instance-Level Alignment(SGIA), which enforces semantic consistency by mapping generator-encoded features into the semantic space defined by a pre-trained segmentation network S_G and imposing density constraints; and (b) Noise-Aware Adversarial Learning, which utilizes a dynamically generated spatial weight map to suppress false-positive responses caused by non-specific staining.

2.2 Segmentation-Guided Instance-Level Alignment(SGIA)

To explicitly model discrete cell instances, we employ a pre-trained IgG4-positive plasma cell segmentation network S_G operating on real IgG4 IHC images [21]. The segmentation model is trained independently and kept frozen during virtual staining to provide stable instance-level semantic priors.

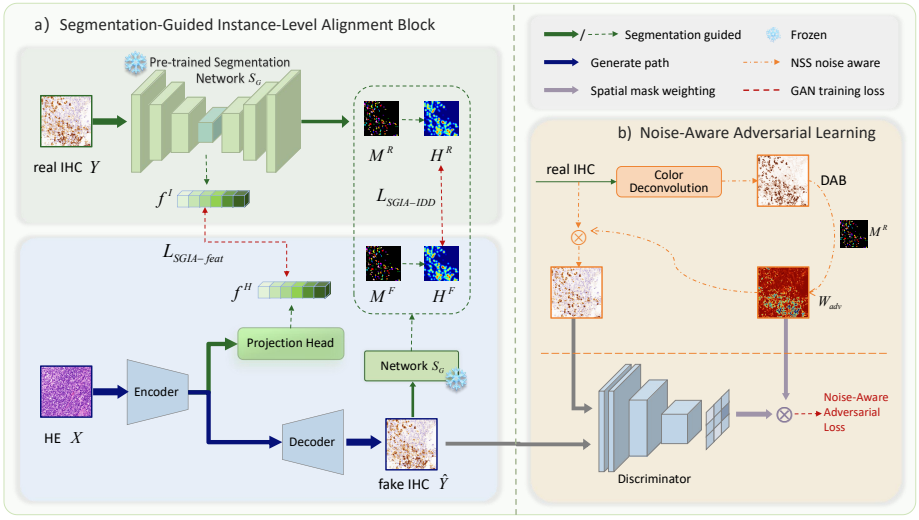

Fig. 2. Overview of the proposed method. The method consists of a Segmentation-Guided Instance-Level Alignment (SGIA) module (left) and a Noise-Aware Adversarial Learning strategy (right).

Instance Feature Alignment. The standard CUT framework preserves local structural correspondence between H&E and IHC images through the PatchNCE loss [13]. However, such patch-level constraints are often insufficient to capture the fine-grained morphology and functional semantics of plasma cells. To inject instance-level priors into the generation process and enhance the awareness of discrete instances, we introduce a lightweight cell semantic alignment branch C attached to the generator encoder E_G .

Given an input H&E image x , the generator encoder produces feature representations $E_G(x)$, which are projected by the alignment branch to obtain instance-aware features $f^H = C(E_G(x))$. Meanwhile, the paired real IHC image y (specifically, the cleaned version y_{clean}, detailed in Sec. 2.3) is fed into the frozen segmentation network S_G, from which intermediate instance-level semantic features $f^I = E_S(y_{clean})$ are extracted.

By enforcing consistency between f^H and f^I , the generator's encoded morphological features are explicitly aligned with the discrete cell semantics of real

IHC images, enhancing the model's sensitivity to cell-level morphology and functional heterogeneity.

$$\mathcal{L}_{SGIA\text{-}feat} = E_{x,y}\left[\|f^H - f^I\|_1\right] \tag{1}$$

Instance Density Distribution Consistency. Beyond feature alignment, we further constrain the generation process using instance-aware density distribution consistency objectives. The generated IHC image $\hat{y}$ is fed into the frozen S_G to obtain the predicted cell probability map M^F from which a density distribution map H^F is derived. The real IHC image y is processed in the same manner to produce the corresponding density map H^R. Consistency between H^F and H^R is enforced to align the spatial distribution of IgG4-positive cell instances.

The overall loss of the SGIA module is defined as:

$$\mathcal{L}_{SGIA}(G) = \lambda_{SGIA\text{-}feat}\mathcal{L}_{SGIA\text{-}feat} + \lambda_{SGIA\text{-}IDD}\|H^R - H^F\|_2^2 \tag{2}$$

2.3 Noise-Aware Adversarial Learning

IgG4-stained images commonly exhibit non-specific staining (NSS), manifested as diffuse background responses that intensify in regions with high positive cell density. If treated as genuine supervision signals, these background artifacts may be mistakenly learned by the generator, leading to false-positive cell generation and degraded pathological consistency.

Identification of Non-Specific Staining. We design an automated cleaning algorithm based on color deconvolution and semantic priors to construct reliable supervision signals. For a real IHC image y, we first extract the DAB channel y_{dab} via color deconvolution [17] and convert it to the HSV space. A soft weighting mask for high-DAB regions is then computed as $W_{\mathrm{DAB}} = \mathcal{H}(y_{dab} - \tau)$,where τ=0.2 is empirically chosen based on visual inspection of DAB saturation. The artifact mask is defined as: $M_{\mathrm{ns}} = W_{\mathrm{DAB}} \odot (\mathbf{1} - M^R)$, which captures regions with strong DAB signal but low cell likelihood. A noise-reduced target image y_{clean} is then generated through spatially adaptive interpolation:

$$y_{\mathrm{clean}} = y \odot (\mathbf{1} - \alpha M_{\mathrm{ns}}) + \alpha M_{\mathrm{ns}}, \tag{3}$$

with fading strength α=0.85, balances effective NSS suppression with preservation of genuine tissue structures.

Mask-Guided Noise-Aware Adversarial Optimization. To further suppress noise during generation, the artifact mask M_{ns} is transformed into a spatial weight map W_{adv} and incorporated into the adversarial loss. Cell regions are assigned higher weights to preserve realism, while artifact regions receive lower weights to reduce their influence. The weighted adversarial loss is formulated as:

$$\mathcal{L}_{GAN}^{weighted}(G, D) = \mathbb{E}_{y\sim\mathcal{Y}}\left[W_{adv} \cdot \log D(y_{clean})\right] + \mathbb{E}_{x\sim\mathcal{X}}\left[W_{adv} \cdot \log(1 - D(G(x)))\right] \tag{4}$$

This strategy encourages the discriminator to focus on stable, pathology-relevant cell structures while ignoring background staining noise, thereby effectively suppressing false-positive patterns during generation.

2.4 Total Loss

The comprehensive loss function in our model is formulated as:

$$\mathcal{L}_{total} = \lambda_{GAN}\mathcal{L}_{GAN}^{weighted} + \lambda_{NCE}\mathcal{L}_{NCE} + \lambda_{GP}\mathcal{L}_{GP} + \mathcal{L}_{SGIA} \tag{5}$$

where $\mathcal{L}_{GAN}^{weighted}$ represents the noise-aware adversarial loss described in Sec 2.3. $\mathcal{L}_{NCE}$ denotes the PatchNCE loss [13], computed between the input H&E image x and the generated IHC image $\hat{y}$ to preserve anatomical morphology. $\mathcal{L}_{GP}$ denotes the gradient penalty term [10], which stabilizes adversarial training and encourages globally coherent texture synthesis.

3 Experiments

Datasets. We collected 127 H&E and IgG4 WSI pairs of orbital masses at 40x magnification. After filtering out slides with blurred staining, severe fading, or significant artifacts, the remaining WSIs were spatially aligned using a WSI-level automated registration pipeline[20]. Finally we obtained 75 well-registered pairs of H&E-IgG4 WSIs. Because the paired slides were from adjacent tissue sections, perfect pixel-level correspondence is not guaranteed due to sectioning gaps. We split these into 60 pairs for training and 15 for testing, generating 22,936 training and 4,246 testing patches, all cropped to 1024 × 1024.

Implementation Details. We adopt CUT [13] as the baseline, with a ResNet-6Blocks generator and a 4-layer PatchGAN discriminator [7]. The cell segmentation network is trained in two stages: first pre-trained using coarse annotations from Cellpose [19] ("cyto2" model) on 1,774 IgG4 image patches of size 512×512, then fine-tuned on pathologist-annotated positive cells. The fine-tuning set consists of 363 patches (512×512, with partial overlap) extracted from 45 ROIs across 6 whole-slide images. These patches naturally contain non-specific staining (NSS) artifacts, but the annotations strictly delineate only IgG4-positive plasma cells and explicitly exclude NSS regions and non-cell structures, ensuring that supervision is free of artifact contamination even in noisy areas. After fine-tuning, on a held-out test set that also contains NSS artifacts, the model achieves a Dice of 0.874 (vs.0.761 pre-trained) and a pixel-level precision of 0.885, confirming its ability to distinguish true positive cells from background artifacts. The segmentation model remains frozen during virtual staining training.

Table 1. Quantitative evaluation on the paired H&E-IgG4 test set. The **best** values are highlighted in bold. The second best values are underlined. KID scores have been multiplied by 1000.

	Image Quality					Pathological Relevance		
Method	PSNR ↑	SSIM ↑	KID ↓	FID ↓	PHV ↓	MAE ↓	PCC ↑	KL ↓
Pix2Pix [7]	21.2897±2.77	0.5244±0.15	39.5±4.6	65.03	0.4555	8.6163	0.0005	9.6003
PyramidP2P [10]	18.5109±3.08	0.4178±0.15	90.8±9.5	61.10	0.5644	9.0790	0.0002	10.3893
CycleGAN [23]	20.2518±3.17	0.4969±0.14	13.8±5.1	36.82	0.4482	19.0494	0.1491	9.7257
CUT+GP [13]	19.8028±2.04	0.4915±0.11	42.6±7.5	69.17	0.4758	12.5336	0.0237	11.8842
ASP [9]	20.6633±2.1	0.4896±0.12	19.6±4.1	51.99	0.4775	5.9501	0.0333	10.967
PPT [22]	20.7801±2.57	0.5195±0.14	39.2±5.6	43.05	0.4476	8.9690	0.0009	6.3865
Ours	**21.774±2.68**	**0.5293±0.13**	**5±2.5**	**20.82**	**0.3913**	**3.4201**	**0.3658**	**4.7570**

All experiments are conducted on a single NVIDIA GeForce RTX 3090 GPU. The model is trained for 40 epochs with a batch size of 1. The loss balancing hyperparameters are: $\lambda_{\text{GAN}} = 1.0$, $\lambda_{\text{NCE}} = 1.0$, $\lambda_{\text{GP}} = 10.0$, $\lambda_{\text{SGIA-feat}} = 10.0$, $\lambda_{\text{SGIA-IDD}} = 1.0$.

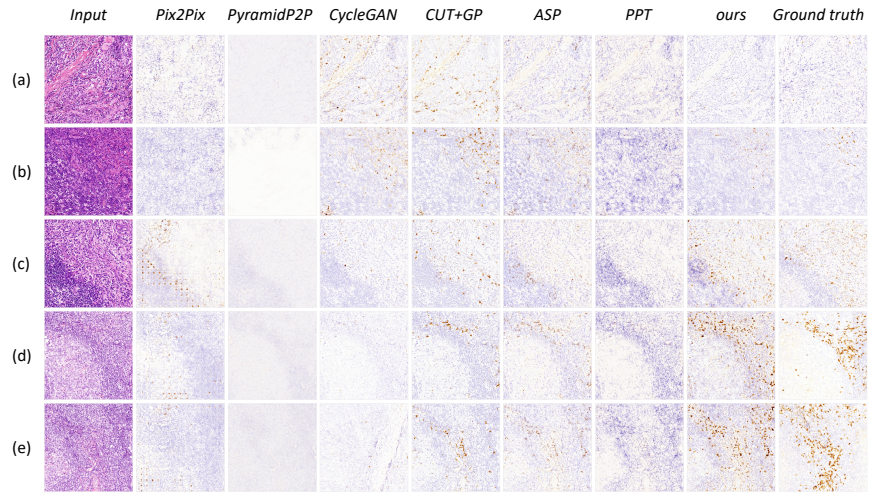

Fig. 3. Qualitative comparison between our model and other methods.

Evaluation Metrics. To comprehensively evaluate the quality and clinical value of the generated images, we employed metrics covering both image quality and pathological relevance. For image quality, PSNR and SSIM were used to assess low-level pixel consistency, while FID [6]and KID [2]were utilized to measure

the realism of the generated distribution, complemented by PHV [11]for perceptual similarity. To assess instance-level quantitative pathological relevance, we use the frozen HoVerNet [5]segmentation model to detect IgG4-positive plasma cell. We calculated the Mean Absolute Error (MAE) between real and generated positive cell counts to evaluate density consistency (abundance), and computed the Pearson Correlation Coefficient (PCC) and KL Divergence on normalized distribution maps to measure the accuracy of spatial cellular distribution.

Quantitative Evaluation. Table 1 reports the quantitative comparison on the paired H&E-IgG4 test set. In terms of image-level quality metrics, our method achieves competitive or improved performance compared to baseline methods, indicating that the introduction of instance-level constraints does not compromise overall visual fidelity. More importantly, for pathology-oriented metrics, the proposed method consistently reduces the error in IgG4-positive cell density(MAE), improves the correlation of spatial density distributions (PCC), and yields lower distributional divergence (KL). These results demonstrate that explicitly modeling discrete cell instances and suppressing background noise leads to more reliable instance-level pathological consistency.

Qualitative Evaluation. Fig. 3 visualizes representative results. Most baseline methods mainly preserve the global staining style and coarse tissue appearance, but often fail to accurately generate sparse IgG4-positive plasma cells. Notably, some approaches (e.g., CUT, ASP) generate false-positive cells in negative regions. In contrast, our approach improves the localization and clarity of positive cell instances while maintaining a clean background in negative regions.

Table 2. Ablation study of proposed components. SGIA: Segmentation-Guided Instance-Level Alignment; NA: Noise-Aware Learning.

	Image Quality					Pathological Relevance		
Method	PSNR ↑	SSIM ↑	KID ↓	FID ↓	PHV ↓	MAE ↓	PCC ↑	KL ↓
Baseline (CUT)	19.8028±2.04	0.4915±0.11	42.6±7.5	69.17	0.4758	12.5336	0.0237	11.8842
w/o SGIA	20.2791±2.63	0.5159±0.13	21.8±5.4	44.98	0.4394	6.0022	0.2012	8.1587
w/o NA	21.4862±2.57	**0.5295±0.14**	9.6±2.8	24.69	0.4251	**3.2318**	0.3524	5.2666
Ours	**21.774±2.68**	0.5293±0.13	**5±2.5**	**20.82**	**0.3913**	3.4201	**0.3658**	**4.7570**

Ablation Study. We conduct ablation studies to validate the effectiveness of the proposed SGIA module and noise-aware adversarial learning strategy. As shown in Table 2, the complete model consistently achieves the best or second-best performance across both image quality and pathology-oriented metrics, demonstrating that SGIA and noise-aware learning are complementary. Together, they enable more reliable instance-level quantitative generation of diagnostically relevant positive plasma cells.

4 Conclusion

In this work, we address the challenge of quantitative IgG4 virtual staining under noisy supervision, where pathological assessment critically depends on diagnostic discrete cell instances rather than continuous tissue appearance. We proposed IgG4-QPC, a framework integrating Segmentation-Guided Instance-Level Alignment (SGIA) and Noise-Aware Adversarial Learning. Experimental results show that, compared with existing virtual staining methods, our approach more effectively preserves pathology-relevant cell density and spatial distribution, leading to improved instance-level pathological consistency. Despite the inherent difficulty of the task and the lack of pixel-wise alignment, our method provides a more reliable foundation for quantitative, cell-centric pathological analysis. Future work will extend validation to multi-center, multi-organ cohorts and incorporate reader studies with pathologists to further establish clinical utility and generalizability. Extending the framework to other immunohistochemical markers characterized by sparse, small objects also remains an important direction.

Acknowledgments. We thank for the support of National Natural Science Foundation of China under Grant 62476183 and the National Key Research and Development Program of China under Grant 2025YFE0216500-1.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bai, B., Yang, X., Li, Y., Zhang, Y., Pillar, N., Ozcan, A.: Deep learning-enabled virtual histological staining of biological samples. Light: Science & Applications **12**(1), 57 (2023)
2. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. arXiv preprint arXiv:1801.01401 (2018)
3. Buchwalow, I., Samoilova, V., Boecker, W., Tiemann, M.: Non-specific binding of antibodies in immunohistochemistry: Fakes and facts. Nature Precedings pp. 1–1 (2011)
4. Chen, F., Zhang, R., Zheng, B., Sun, Y., He, J., Qin, W.: Pathological semantics-preserving learning for h&e-to-ihc virtual staining. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 384–394. Springer (2024)
5. Graham, S., Vu, Q.D., Raza, S.E.A., Azam, A., Tsang, Y.W., Kwak, J.T., Rajpoot, N.: Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. Medical image analysis **58**, 101563 (2019)
6. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
7. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)

8. Lanzillotta, M., Mancuso, G., Della-Torre, E.: Advances in the diagnosis and management of igg4 related disease. Bmj **369** (2020)
9. 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: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 632–641. Springer (2023)
10. 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)
11. 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)
12. Liu, Y., Li, X., Zheng, A., Zhu, X., Liu, S., Hu, M., Luo, Q., Liao, H., Liu, M., He, Y., et al.: Predict ki-67 positive cells in h&e-stained images using deep learning independently from ihc-stained images. Frontiers in Molecular Biosciences **7**, 183 (2020)
13. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: European conference on computer vision. pp. 319–345. Springer (2020)
14. Pati, P., Karkampouna, S., Bonollo, F., Compérat, E., Radić, M., Spahn, M., Martinelli, A., Wartenberg, M., Kruithof-de Julio, M., Rapsomaniki, M.: Accelerating histopathology workflows with generative ai-based virtually multiplexed tumour profiling. Nature machine intelligence **6**(9), 1077–1093 (2024)
15. Peng, Q., Lin, W., Hu, Y., Bao, A., Lian, C., Wei, W., Yue, M., Liu, J., Yu, L., Wang, L.: Advancing h&e-to-ihc virtual staining with task-specific domain knowledge for her2 scoring. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–13. Springer (2024)
16. Pillar, N., Ozcan, A.: Virtual tissue staining in pathology using machine learning. Expert Review of Molecular Diagnostics **22**(11), 987–989 (2022)
17. Ruifrok, A.: Quantification of histochemical staining by color deconvolution. Analytical and quantitative cytology and histology/the International Academy of Cytology [and] American Society of Cytology (2001)
18. Shi, L.l., Xiong, P., Yang, M., Ardicli, O., Schneider, S.R., Funch, A.B., Kiykim, A., Lopez, J., Akdis, C.A., Akdis, M.: Role of igg4 antibodies in human health and disease. Cells **14**(9), 639 (2025)
19. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. Nature methods **18**(1), 100–106 (2021)
20. Wodzinski, M., Marini, N., Atzori, M., Müller, H.: Regwsi: Whole slide image registration using combined deep feature-and intensity-based methods: Winner of the acrobat 2023 challenge. Computer methods and programs in biomedicine **250**, 108187 (2024)
21. Xu, Q., Ma, Z., Duan, W., et al.: Dcsau-net: A deeper and more compact split-attention u-net for medical image segmentation. Computers in Biology and Medicine **154**, 106626 (2023)
22. Zhang, W., Hui, T.H., Tse, P.Y., Hill, F., Lau, C., Li, X.: High-resolution medical image translation via patch alignment-based bidirectional contrastive learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 178–188. Springer (2024)

23. 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. pp. 2223–2232 (2017)