

HEMIT: H&E to Multiplex-immunohistochemistry Image Translation with Dual-Branch Pix2pix Generator

Chang Bian¹, Beth Philips², Tim Cootes¹, and Martin Fergie¹

¹ Division of Informatics, Imaging and Data Sciences, School of Health Sciences,
University of Manchester

² Division of Cancer Sciences, University of Manchester
`chang.bian@postgrad.manchester.ac.uk`

Abstract. Computational analysis of multiplexed immunofluorescence histology data is emerging as an important method for understanding the tumour micro-environment in cancer. This work presents HEMIT, a dataset designed for translating Hematoxylin and Eosin (H&E) sections to multiplex-immunohistochemistry (mIHC) images, featuring DAPI, CD3, and panCK markers. Distinctively, HEMIT’s mIHC images are multi-component and cellular-level aligned with H&E, enriching supervised stain translation tasks. To our knowledge, HEMIT is the first publicly available cellular-level aligned dataset that enables H&E to multi-target mIHC image translation. This dataset provides the computer vision community with a valuable resource to develop novel computational methods which have the potential to gain new insights from H&E slide archives. We also propose a new dual-branch generator architecture, using residual Convolutional Neural Networks (CNNs) and Swin Transformers which achieves better translation outcomes than other popular algorithms. When evaluated on HEMIT, it outperforms pix2pixHD, pix2pix, U-Net, and ResNet, achieving the highest overall score on key metrics including the Structural Similarity Index Measure (SSIM), Pearson correlation score (R), and Peak signal-to-noise Ratio (PSNR). Additionally, downstream analysis has been used to further validate the quality of the generated mIHC images. These results set a new benchmark in the field of stain translation tasks.

Keywords: HEMIT dataset · Hematoxylin and Eosin (H&E) · Multiplex-immunohistochemistry (mIHC) · stain translation · Swin Transformers · Downstream analysis.

1 Introduction

Multiplex immunohistochemistry/immunofluorescence (mIHC/IF) technologies have revolutionized the study of the tumor microenvironment (TME), providing critical insights for cancer research and immunotherapy [25]. These technologies enable the detection of multiple markers in a single tissue slice, revealing intricate spatial interactions. However, the complexity and cost of mIHC limits its

accessibility [24]. Deep learning offers a promising solution to these challenges, with recent advancements in stain translation and virtual staining allowing for better exploitation of existing H&E slides [3,10,21,6,27].

Capitalizing on the ease and cost-efficiency of producing H&E images, this paper aims to develop an image-to-image translation method that can convert H&E images into their corresponding mIHC counterparts.

Image-to-image translation algorithms are well-established across a range of image analysis domains. Specifically, pix2pix [9] employs a conditional generative adversarial network (cGAN) approach for paired images, adeptly generating high-fidelity translations. Building upon this, pix2pixHD [26] has demonstrated significant outcomes with high-resolution paired images. Notably, a recent study [14] introduced a multi-scale loss term, improving the translation performance for HER2 images. Furthermore, other methods [5,8,12,13,11,19] have been designed specifically for image-to-image translation tasks, underscoring the breadth and depth of research in this domain.

High-quality datasets are essential for effective supervised image translation, yet the realm of pathological image translation lacks comprehensive resources. The BCI dataset [14] is proposed for H&E to HER2 IHC image translation. However, it is limited in use as staining is performed on consecutive tissue sections so there is no cell-to-cell mapping across stains. Further, it is limited to one target stain (HER2) unlike mIHC where multiple stains can be predicted simultaneously. In response, we present HEMIT: A dataset for H&E to mIHC Translation. This is the first publicly-available cellular-level aligned dataset for stain translation. Concurrently, we propose a specialized method tailored for this task which we compare to various state-of-the-art image translation algorithms. This work makes the following contributions:

1. Introduction of HEMIT: a paired dataset for H&E to mIHC image translation. To the best of our knowledge, HEMIT is the first publicly available cellular-level aligned dataset that enables H&E to multi-target mIHC image translation.
2. Development of a SwinTransformer-CNN-based dual-branch pix2pix strategy to convert H&E images into mIHC versions. Our methodology assimilates both global information and spatial details, culminating in superior outcomes, setting benchmark results for the proposed dataset.
3. A thorough empirical evaluation was conducted on the newly introduced HEMIT dataset, setting a benchmark for future investigations within the research community. Further downstream analysis, employing Qupath [1], substantiated the high quality of the images produced. This establishes a solid foundation for subsequent studies and applications, encouraging a deeper exploration into the dataset’s potential.

2 HEMIT Dataset

We present HEMIT: A cellular-level aligned dataset for H&E to mIHC Image Translation. A schematic of our dataset’s construction pipeline is provided in fig. 1a. Full dataset will be released upon acceptance of this paper.

2.1 Data Collection

HEMIT’s raw data is sourced from ImmunoAIZER [2] which we have adapted to make it suitable for the computer vision community. Notably, HEMIT distinguishes itself from other datasets: the H&E and mIHC slide pairs are derived from the identical tissue section, not consecutive slides. The tissue was first stained with mIHC protocols and then bleached before the H&E staining. This feature leads to higher fidelity and alignment between the matched image pairs and better translation outcomes [4].

In the mIHC images, three pivotal cell type identification markers are incorporated: DAPI to signify cell nuclei, pan-cytokeratin (panCK) to highlight tumor regions, and CD3 to pinpoint T cells—all of which are integral to TME analysis. These specific markers serve as the foundation for our H&E to mIHC stain translation benchmark. We have collated a selection of publicly available datasets that leverage H&E images for predicting IHC expressions. A comparative overview of their characteristics is presented in table 1.

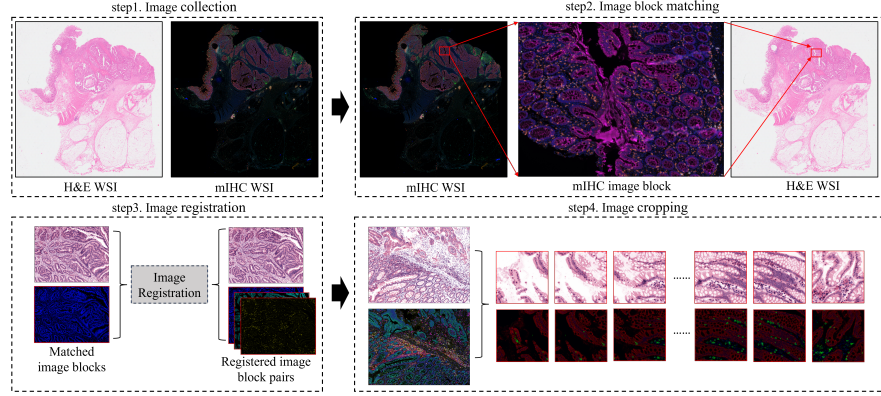
Table 1: Summary of publicly available datasets

Datasets	Staining Types	Sectioning Approach	IHC/mIHC Markers	Ground Truths
HEMIT	H&E & mIHC	same slide	DAPI, CD3, panCK	cellular level
BCI [14]	H&E & IHC	consecutive slides	HER2	structural level
HEROHE [7]	H&E	H&E only	Clinical HER2 status	slide level


2.2 Data Preprocessing

Despite both H&E and mIHC staining being performed on the same tissue slide, the re-staining and scanning processes mean that the captured images do not align perfectly. We employ a 2-step registration process to ensure cellular-wise alignment of the image pairs which is crucial for optimal training performance. Implementation details of the registration process are given in the supplementary material.

Upon concluding the registration, a margin of 50 pixels from each edge is trimmed to account for rotational transformations. Subsequently, the block pairs are cropped into 1024×1024 patches, maintaining a 50% overlap. Color normalization [18] is then applied to all H&E patches to mitigate stain variations. Visualizations of the final registered image pairs are shown in fig. 1b. This process



(a) Processing pipeline of the HEMIT dataset.



(b) Visualization of the registered image pairs, along with constitute channels.

Fig. 1: Overview of dataset processing and visualization in the HEMIT dataset: (a) Processing pipeline. (b) Visualization of registered image pairs.

resulted in 5292 matched image patch pairs. This processed data is distributed into three partitions: *training*, *validation*, and *testing*. The division yields 3717 patches designated for training, 630 for validation, and 945 for the testing phase. Importantly, to prevent data leakage of our evaluation, patches across these subsets are derived from unique patient samples, precluding potential data overlap.

3 Proposed Method

3.1 Architecture

Our dual-branch architecture is based on the original pix2pix [9] model, and we focus on the Swin Transformer-based branch. The overall framework is shown in fig. 2.

We devised a dual-branch generator architecture. The main branch with residual blocks of the generator is to extract spatial nuances from the input H&E images, and an auxiliary branch powered by Swin Transformers [16] is incorporated to integrate information across multiple scales. The features extracted by each stage of the Swin Transformer branch are fused with the feature

maps of the CNN branch by the Feature Map Fusion (FMF) module. This architecture caters to the multi-scale nature of pathological images. This design captures information from the structure of individual cells to the interaction of cell clusters, vital for interpreting complex tissue environments.

The auxiliary branch, operating in parallel with the main branch, leverages Swin Transformer modules [16] to adeptly capture multi-scale features. Notably, the Swin Transformer introduces both Window-based Multi-head Self Attention (W-MSA) and Shifted Window-based MSA (SW-MSA). Prior to each MSA module and MLP, a Layer Normalization (LN) is applied. \hat{s}_l and s_l denote the outputs of W-MSA and MLP respectively, at the l^{th} layer. These are computed as:

$$\begin{aligned} \hat{s}_l &= \text{W-MSA}(\text{LN}(s_{l-1})) + s_{l-1}; s_l = \text{MLP}(\text{LN}(\hat{s}_l)) + \hat{s}_l, \\ \hat{s}_{l+1} &= \text{SW-MSA}(\text{LN}(\hat{s}_l)) + s_l; s_{l+1} = \text{MLP}(\text{LN}(\hat{s}_{l+1})) + \hat{s}_{l+1} \end{aligned} \quad (1)$$

The FMF module integrates features from the two branches by utilizing Gated Cross Attention, which emphasizes spatially significant regions. The Swin Transformer output S_i is aligned with the CNN feature map F_i in both dimensions and resolution to form S'_i . A gating layer then processes F_i through convolution and sigmoid activation to produce a gating map G , pinpointing areas of importance within F_i . The model then selects the top k pixels in G , indicative of the most informative regions. Subsets $F_{i,k}$ and $S'_{i,k}$ are extracted from F_i and S'_i correspondingly for cross attention operation in equation 2. Focusing on these k points allows the cross-attention to exchange information primarily in regions crucial for the model's performance, thereby improving efficiency. The choice of k is based on empirical analysis, details in the supplementary materials.

$$A = \text{Attention}(F_{i,k}, S'_{i,k}, S'_{i,k}) \quad (2)$$

where:

- $F_{i,k}$ serves as the query, representing the top k elements from the CNN feature maps.
- $S'_{i,k}$ serves as the key and value, derived from the corresponding elements in the Swin Transformer output.

The adversarial loss is adopted from the pix2pix framework:

$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}_{x,y}[\log D(x, y)] + \mathbb{E}_{x,z}[\log(1 - D(x, G(x, z)))]. \quad (3)$$

The terms x , y , and z denote the input H&E image, the ground truth mIHC image, and random noise, respectively. Additionally, the L1 loss from pix2pix is employed to preserve structural consistency:

$$\mathcal{L}_1 = \mathbb{E}_{x,y,z}[\|y - G(x, z)\|_1]. \quad (4)$$

Thus, our overarching objective function becomes:

$$G^* = \arg \min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_1(G), \quad (5)$$

where λ balances the contributions of the adversarial loss and the L_1 loss.

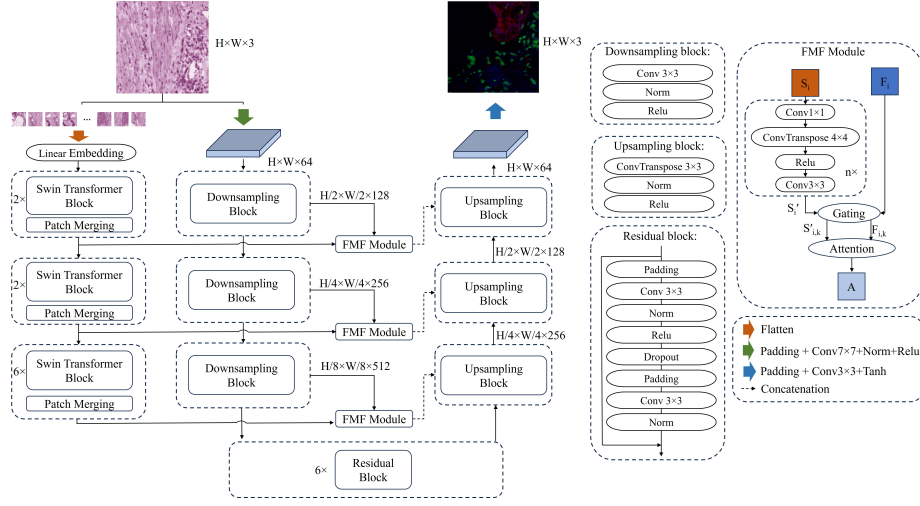


Fig. 2: Overall structure of the proposed generator.

4 Experiments

We used the Adam optimizer over 100 epochs, with an initial learning rate of 0.00003 for 50 epochs, and then linearly decayed to zero over the remaining epochs, on an NVIDIA Tesla V100 16GB GPU.

4.1 Benchmark Results

Following established benchmarks in image-to-image translation [14,3,6,15], we adopt the Structural Similarity Index Measure (SSIM), Pearson correlation score (R), and Peak Signal to Noise Ratio (PSNR) as evaluation metrics to gauge the quality of the synthesized images.

Our method demonstrated superior performance on the HEMIT dataset, achieving the highest scores in SSIM (0.875), Pearson correlation (0.746), and PSNR (29.886), as shown in table 2. This indicates its effectiveness in image translation quality compared to other models.

Visually, as seen in fig. 3, while U-Net and ResNet scored well in SSIM and PSNR, they fell short in Pearson correlation due to insufficient emphasis on less dominant markers, such as CD3. This highlights the importance of Pearson correlation for evaluating multi-stain translations and demonstrates that GAN-based frameworks significantly improve the realism of image translations, particularly for low-expression markers.

4.2 Downstream Analysis

Our evaluation, incorporating downstream analysis with Qupath [1,17], compared the fidelity of our generated images and other methods against real im-

Table 2: Comparison of evaluation metrics across different methods.

Methods	SSIM				R				PSNR (dB)			
	DAPI	CD3	panCK	Average	DAPI	CD3	panCK	Average	DAPI	CD3	panCK	Average
U-Net	0.791	0.907	0.901	0.866	0.659	0.005	0.949	0.538	27.929	24.733	33.773	28.695
ResNet	0.790	0.907	0.898	0.865	0.677	0.005	0.949	0.544	27.531	24.733	33.212	28.268
pix2pix_UNet	0.775	0.879	0.903	0.852	0.652	0.455	0.943	0.683	27.691	25.659	34.306	29.219
pix2pix_ResNet	0.723	0.906	0.914	0.848	0.668	0.553	0.946	0.723	27.152	26.225	34.359	29.189
pix2pixHD	0.786	0.900	0.888	0.858	0.721	0.530	0.943	0.731	28.156	26.281	34.033	29.471
ours	0.815	0.898	0.913	0.875	0.716	0.571	0.951	0.746	28.610	26.349	34.875	29.886




Fig. 3: Visualization of different methods on HEMIT dataset. DAPI is shown in blue, panCK in red and CD3 in green.

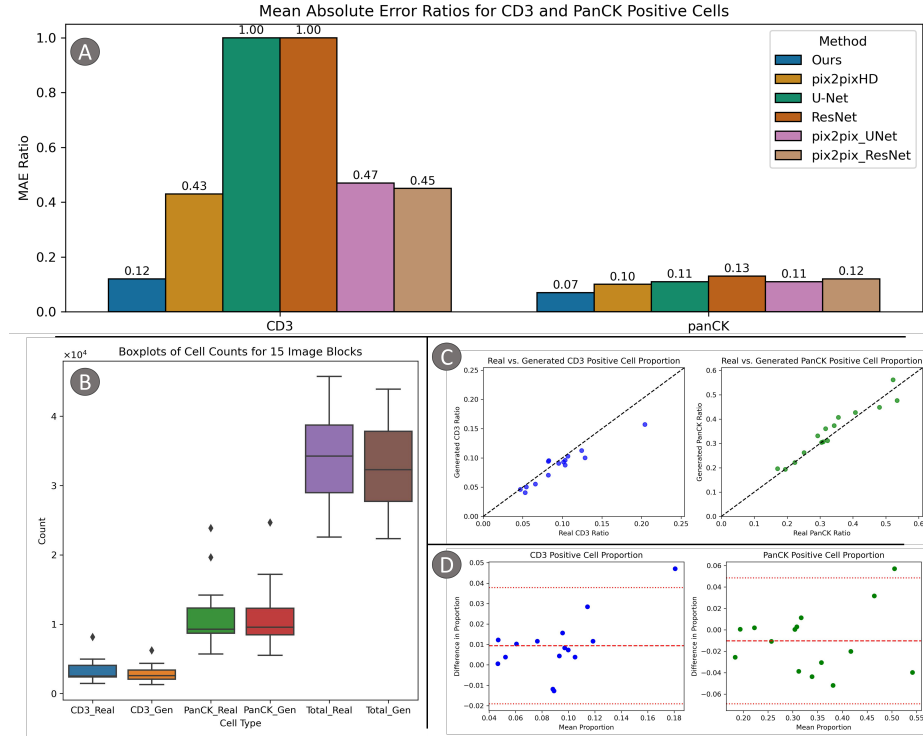


Fig. 4: Downstream Analysis Results: (A) Comparative MAE Ratios for CD3 and PanCK; (B) shows boxplots of cell counts for CD3, panCK, and total of the generated images of our method and real images; (C) and (D) provide scatter plots and Bland-Altman plots for CD3 and panCK positive cell proportions of the generated images of our method and the real images.

ages. We analyzed 15 image blocks, using Stardist [22] for cell detection and the Otsu method [20] for marker positivity, to validate the translation accuracy by comparing the positive cell proportions.

Mean absolute error ratios of positive cell proportions are shown in fig. 4A. Our method shows considerable improvement over the existing state-of-the-art methods, especially on the challenging CD3 marker. Further comparative results of our method are shown in fig. 4B-D. C and D showcase the close correspondence between real and generated images, evidenced by the tight distribution around the identity line in scatter plots and marginal deviations in Bland-Altman plots.

5 Conclusion

In this study, we introduced HEMIT, a pioneering dataset tailored for pathology image translation tasks. HEMIT is specifically designed to bridge the gap between H&E stained sections and their mIHC counterparts, showcasing multiple

markers including DAPI, CD3, and panCK. Notably, the mIHC images in our dataset are multi-component, providing a cellular-level registration with the associated H&E images. This unique alignment presents an invaluable opportunity for supervised stain translation research.

Complementing our dataset contribution, we proposed a dual-branch generator architecture optimized for pathology image analysis. Through extensive experimentation, we contrasted our approach with several SOTA algorithms. The empirical outcomes with downstream analysis underscore the efficacy of our method, establishing a new benchmark for the H&E to mIHC translation domain. We hypothesise that the significant performance improvement is due to the spatial context provided by the Swin Transformer where the tissue compartment can inform the cell protein expression. This is in line with other works [23] which has demonstrated that tissue context can inform cell classification leading to improved accuracy.

This work provides an opportunity for further exploration in the use of predicting biomarkers from H&E pathology images to accelerate research in biomarker prediction.

References

1. Bankhead, P., Loughrey, M.B., Fernández, J.A., Dombrowski, Y., McArt, D.G., Dunne, P.D., McQuaid, S., Gray, R.T., Murray, L.J., Coleman, H.G., et al.: Qupath: Open source software for digital pathology image analysis. *Scientific reports* **7**(1), 1–7 (2017)
2. Bian, C., Wang, Y., Lu, Z., An, Y., Wang, H., Kong, L., Du, Y., Tian, J.: Immunoaizer: a deep learning-based computational framework to characterize cell distribution and gene mutation in tumor microenvironment. *Cancers* **13**(7), 1659 (2021)
3. Burlingame, E.A., Margolin, A.A., Gray, J.W., Chang, Y.H.: Shift: speedy histopathological-to-immunofluorescent translation of whole slide images using conditional generative adversarial networks. In: *Medical Imaging 2018: Digital Pathology*. vol. 10581, p. 1058105. International Society for Optics and Photonics (2018)
4. Burlingame, E.A., McDonnell, M., Schau, G.F., Thibault, G., Lanciault, C., Morgan, T., Johnson, B.E., Corless, C., Gray, J.W., Chang, Y.H.: Shift: speedy histological-to-immunofluorescent translation of a tumor signature enabled by deep learning. *Scientific reports* **10**(1), 17507 (2020)
5. Chong, M.J., Forsyth, D.: Gans n’roses: Stable, controllable, diverse image to image translation (works for videos too!). *arXiv preprint arXiv:2106.06561* (2021)
6. Christiansen, E.M., Yang, S.J., Ando, D.M., Javaherian, A., Skibinski, G., Lipnick, S., Mount, E., O’Neil, A., Shah, K., Lee, A.K., et al.: In silico labeling: predicting fluorescent labels in unlabeled images. *Cell* **173**(3), 792–803 (2018)
7. Conde-Sousa, E., Vale, J., Feng, M., Xu, K., Wang, Y., Della Mea, V., La Barbera, D., Montahaei, E., Baghshah, M., Turzynski, A., et al.: Herohe challenge: Predicting her2 status in breast cancer from hematoxylin–eosin whole-slide imaging. *Journal of Imaging* **8**(8), 213 (2022)

8. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European conference on computer vision (ECCV). pp. 172–189 (2018)
9. 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)
10. Jen, K.Y., Murali, L.K., Lutnick, B., Ginley, B., Govind, D., Mori, H., Gao, G., Sarder, P.: In silico multi-compartment detection based on multiplex immunohistochemical staining in renal pathology. In: Medical Imaging 2021: Digital Pathology. vol. 11603, pp. 294–301. SPIE (2021)
11. Kong, L., Lian, C., Huang, D., Hu, Y., Zhou, Q., et al.: Breaking the dilemma of medical image-to-image translation. *Advances in Neural Information Processing Systems* **34**, 1964–1978 (2021)
12. Lee, H.Y., Tseng, H.Y., Mao, Q., Huang, J.B., Lu, Y.D., Singh, M., Yang, M.H.: Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision* **128**, 2402–2417 (2020)
13. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. *Advances in neural information processing systems* **30** (2017)
14. 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)
15. Liu, Y., Yuan, H., Wang, Z., Ji, S.: Global pixel transformers for virtual staining of microscopy images. *IEEE Transactions on Medical Imaging* **39**(6), 2256–2266 (2020)
16. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
17. Loughrey, M.B., Bankhead, P., Coleman, H.G., Hagan, R.S., Craig, S., McCorry, A.M., Gray, R.T., McQuaid, S., Dunne, P.D., Hamilton, P.W., et al.: Validation of the systematic scoring of immunohistochemically stained tumour tissue microarrays using qupath digital image analysis. *Histopathology* **73**(2), 327–338 (2018)
18. Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., Thomas, N.E.: A method for normalizing histology slides for quantitative analysis. In: 2009 IEEE international symposium on biomedical imaging: from nano to macro. pp. 1107–1110. IEEE (2009)
19. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging* **34**(10), 1993–2024 (2014)
20. Ostu, N.: A threshold selection method from gray-level histograms. *IEEE Trans SMC* **9**, 62 (1979)
21. Rivenson, Y., Wang, H., Wei, Z., de Haan, K., Zhang, Y., Wu, Y., Günaydin, H., Zuckerman, J.E., Chong, T., Sisk, A.E., et al.: Virtual histological staining of unlabelled tissue-autofluorescence images via deep learning. *Nature biomedical engineering* **3**(6), 466–477 (2019)
22. Schmidt, U., Weigert, M., Broaddus, C., Myers, G.: Cell detection with star-convex polygons. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11. pp. 265–273. Springer (2018)

23. Shephard, A.J., Graham, S., Bashir, S., Jahanifar, M., Mahmood, H., Khurram, A., Rajpoot, N.M.: Simultaneous nuclear instance and layer segmentation in oral epithelial dysplasia. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 552–561 (2021)
24. Stack, E.C., Wang, C., Roman, K.A., Hoyt, C.C.: Multiplexed immunohistochemistry, imaging, and quantitation: a review, with an assessment of tyramide signal amplification, multispectral imaging and multiplex analysis. *Methods* **70**(1), 46–58 (2014)
25. Tan, W.C.C., Nerurkar, S.N., Cai, H.Y., Ng, H.H.M., Wu, D., Wee, Y.T.F., Lim, J.C.T., Yeong, J., Lim, T.K.H.: Overview of multiplex immunohistochemistry/immunofluorescence techniques in the era of cancer immunotherapy. *Cancer Communications* **40**(4), 135–153 (2020)
26. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
27. Weinstein, J.N., Collisson, E.A., Mills, G.B., Shaw, K.R., Ozenberger, B.A., Ellrott, K., Shmulevich, I., Sander, C., Stuart, J.M.: The cancer genome atlas pan-cancer analysis project. *Nature genetics* **45**(10), 1113–1120 (2013)