

Segmentation-Guided Colorization of Single-Channel Laser Illumination Images for Improved Surgical Navigability

Janmesh Ukey^{1,2*}, Joost Emiel Marcel Calon³, Nils de Koff³, and Clara I. Sánchez^{1,2}

¹ Universiteit van Amsterdam, The Netherlands

² curAI group, Informatics Institute, Universiteit van Amsterdam, The Netherlands

³ LIMIS Development BV, 8934 AD Leeuwarden, The Netherlands

Abstract. Laser speckle contrast imaging (LSCI) enables real-time intraoperative perfusion assessment in laparoscopic surgery, but visible-wavelength LSCI systems require switching to laser illumination mode, producing a single-channel image that strips away the chromatic tissue cues surgeons rely on for anatomical navigation. Existing unpaired image translation methods lack semantic awareness, producing colorizations that are visually inconsistent across tissue types. We propose a segmentation-guided unpaired image-to-image translation framework that colorizes single-channel red-illumination laparoscopic images into perceptually natural RGB views, restoring visual navigability without interfering with the perfusion computation pipeline. We inject anatomical masks at every decoder layer for spatially consistent synthesis and introduce a class-wise contrastive loss that restricts patch sampling to semantically homogeneous regions, yielding harder within-class negatives. The discriminator is additionally conditioned on the mask for class-aware adversarial feedback, supplemented by a Gram matrix style loss. Our method is evaluated on the Dresden Surgical Anatomy Dataset in a fully unpaired setting, with the red channel serving as the source domain. Quantitatively, it achieves SSIM of 0.81, MS-SSIM of 0.87, PSNR of 22.51 dB, and LPIPS of 0.14, outperforming all baselines across all evaluated metrics. This work demonstrates improved colorization fidelity for laser-mode laparoscopic views under a simulated red-channel proxy, a step toward supporting surgical decision-making during perfusion-guided procedures.

Keywords: Image colorization · Unpaired image-to-image translation · Surgical navigation.

1 Introduction

Minimally invasive surgery (MIS) has transformed modern surgical practice, with laparoscopic procedures now representing the standard of care across a wide

* Corresponding author: j.y.ukey@uva.nl

range of abdominal and thoracic interventions [2]. By providing high-resolution visual access to the operative field through a small-caliber camera inserted via trocar ports, laparoscopy reduces patient morbidity, shortens hospital stay, and accelerates postoperative recovery compared to open surgery [4]. However, visual assessment alone offers surgeons only surface-level structural information, leaving critical physiological parameters such as tissue perfusion invisible to the naked eye.

Intraoperative assessment of tissue perfusion is a critical component of many such procedures, as inadequate perfusion is directly associated with life-threatening complications including anastomotic leakage and ischaemic necrosis [18, 13, 14]. Laser speckle contrast imaging (LSCI) has emerged as a compelling solution for intraoperative perfusion assessment [11, 9, 8, 7]. When coherent laser light, typically in the red or near-infrared spectrum, is projected onto biological tissue, the scattered light from moving red blood cells generates a characteristic interference pattern known as speckle contrast. By computing the spatial or temporal contrast of this speckle field, a quantitative perfusion map can be derived in real time [1, 5].

However, despite its functional richness, the integration of LSCI into clinical laparoscopic workflows introduces a perceptual consideration that warrants attention. Visible-wavelength LSCI systems require switching off the white light, producing a single-channel image captured in the red channel when a red laser source is used. This image is devoid of the chromatic tissue cues that surgeons rely upon for anatomical navigation, introducing additional cognitive considerations in interpreting the operative field. In such cases, colorization of these images can provide the anatomical cues while still maintaining the perfusion information.

Colorization of single-channel medical images has been explored in several domains, including histology restoration, X-ray synthesis, and depth-to-appearance translation [20, 24]. Supervised methods such as Pix2Pix [10] achieve strong colorization fidelity when pixel-aligned pairs are available. However, for visible-wavelength LSCI systems such as ours, acquiring co-registered paired data is practically infeasible, as the two modalities require different illumination conditions. Unpaired image-to-image translation methods offer a principled alternative by learning a mapping between two image domains without requiring pixel-level correspondence. CycleGAN [25] pioneered this paradigm through the use of cycle-consistency constraints, but is known to suffer from training instability and high memory cost due to its dual-generator architecture. Contrastive Unpaired Translation (CUT) [16] addresses these limitations by replacing cycle-consistency with a patch-wise contrastive learning objective derived from mutual information maximization, enabling high-quality domain translation with a single generator and substantially reduced computational overhead.

In this work, we propose a CUT-based unpaired image-to-image translation framework for the colorization of intraoperative laser speckle images captured in the red channel. Our method learns a mapping from the simulated single-channel red-illumination domain to the full-color white-light domain using unpaired la-




Fig. 1. Overview of the proposed segmentation-guided colorization framework. The generator decoder is conditioned on the segmentation mask S at every upsampling layer. A class-wise contrastive loss restricts patch sampling to semantically homogeneous regions, and the discriminator (D) is additionally conditioned on S for class-aware adversarial feedback.

paroscopic image sets, enabling the generation of perceptually natural, anatomically consistent colorized views that surgeons can readily interpret. By restoring chromatic tissue context to the perfusion imaging view, our approach aims to bridge the perceptual gap between modalities and support more confident, safer intraoperative decision-making. Our key contributions are:

- a **class-wise contrastive loss** that enforces tissue-consistent feature learning by restricting patch sampling to semantically homogeneous regions;
- a **segmentation-conditioned discriminator** that provides class-aware adversarial feedback; and
- a **segmentation-conditioned generator decoder** that injects spatial semantic context at every decoding scale.

2 Methodology

2.1 Problem Formulation

Let $\mathcal{X}$ denote the source domain of simulated single-channel red-illumination laser images and $\mathcal{Y}$ denote the target domain of white-light RGB laparoscopic images. We formulate the colorization task as an unpaired image-to-image translation problem. Given an unpaired dataset $\{x_i\}_{i=1}^N \subset \mathcal{X}$ and $\{y_j\}_{j=1}^M \subset \mathcal{Y}$, our goal is to learn a generator $G : \mathcal{X} \rightarrow \mathcal{Y}$ that produces a perceptually realistic and anatomically faithful colorized image $\hat{y} = G(x)$. In addition to the image data, we assume access to pixel-wise segmentation masks $S \in \{1, \dots, C\}^{H \times W}$, where C denotes the number of anatomical tissue classes. The overall architecture is illustrated in Figure 1.

2.2 Segmentation-Conditioned Generator

The generator follows a ResNet-based encoder-decoder architecture. At each decoder layer l , the segmentation mask is resized to match the current feature resolution and concatenated channel-wise with the incoming feature map F^l , then processed by a lightweight convolutional block ϕ^l :

$$S^l = \text{Resize}(S, H_l \times W_l), \quad \tilde{F}^l = \phi^l([F^l \parallel S^l]). \quad (1)$$

This conditions the decoder on the anatomical layout at every upsampling scale, ensuring colorization decisions remain consistent with the underlying tissue class.

2.3 Loss Functions

Adversarial Loss with Segmentation-Conditioned Discriminator ($\mathcal{L}_{\text{GAN}}$)

We adopt a least-squares GAN objective (Adv-Loss) [15]. The discriminator D is trained to distinguish real target-domain images from generated images, while the generator is trained to fool the discriminator. The discriminator D is additionally conditioned on the ground-truth segmentation mask S , encouraging class-aware evaluation of the synthesized appearance.

$$\mathcal{L}_{\text{GAN}}(G, D) = \mathbb{E}_{y \sim \mathcal{Y}} [(D(y, S_y) - 1)^2] + \mathbb{E}_{x \sim \mathcal{X}} [D(G(x), S_x)^2]. \quad (2)$$

Standard Patch-wise Contrastive Loss ($\mathcal{L}_{\text{NCE}}$) Following CUT [16], spatial correspondence is enforced via patch-wise contrastive learning (Patch-NCE). For a query feature h^k from $G(x)$ at position k , the feature at the same position in x is the positive h^{k+} , and features at all other positions are negatives:

$$\mathcal{L}_{\text{NCE}}(G, H, x) = - \sum_{k=1}^K \log \frac{\exp(h^k \cdot h^{k+} / \tau)}{\exp(h^k \cdot h^{k+} / \tau) + \sum_{j \neq k} \exp(h^k \cdot h^j / \tau)}, \quad (3)$$

where $h^k = H(G_{\text{enc}}(x))_k$ is the ℓ_2 -normalised feature at position k , h^{k+} is its positive counterpart from the input encoder, and τ is a temperature parameter.

Class-wise Contrastive Loss ($\mathcal{L}_{\text{cls}}$) We extend Patch-NCE by restricting patch sampling to semantically homogeneous regions, which we refer to as Class-NCE. For each class c , we define a binary mask $M_c(k) = \mathbb{I}[S(k) = c]$ and apply Patch-NCE only within M_c , so negatives are drawn from the same tissue class rather than the full image:

$$\mathcal{L}_{\text{cls}}^{(c)} = - \frac{1}{K_c} \sum_{k: M_c(k)=1} \log \frac{\exp(h^k \cdot h^{k+} / \tau)}{\exp(h^k \cdot h^{k+} / \tau) + \sum_{\substack{j \neq k \\ M_c(j)=1}} \exp(h^k \cdot h^j / \tau)}, \quad (4)$$

where K_c is the number of sampled positions within class c . The final class-wise loss ($\mathcal{L}_{\text{cls}}$) is averaged uniformly across all C classes present in the image. This introduces harder within-class negatives while preventing large tissue regions from dominating training.

Style Loss ($\mathcal{L}_{\text{style}}$) To preserve the global chromatic style and textural distribution of the target domain, we incorporate a Gram matrix-based style loss [6]. Let $\{\psi_l\}$ denote feature maps extracted from a set of layers of a pretrained VGG network [19]. The style loss between a generated image $\hat{y}$ and a real target image y is:

$$\mathcal{L}_{\text{style}} = \sum_l \|\mathcal{G}(\psi_l(\hat{y})) - \mathcal{G}(\psi_l(y))\|_F^2, \quad (5)$$

where $\mathcal{G}(F) = \frac{1}{H_l W_l} F F^\top$ denotes the Gram matrix of feature map $F \in \mathbb{R}^{C_l \times H_l W_l}$.

2.4 Total Training Objective

The generator G and discriminator D are trained jointly by optimising the following combined objective:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{GAN}} + \lambda_{\text{NCE}} \mathcal{L}_{\text{NCE}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} + \lambda_{\text{style}} \mathcal{L}_{\text{style}}, \quad (6)$$

where λ_{NCE} , λ_{cls} , and λ_{style} are scalar weights balancing the contribution of each loss term.

3 Experiments and Results

3.1 Experimental Setup

Dataset We evaluate our framework on the *Dresden Surgical Anatomy Dataset* [3], a publicly available, large-scale collection of laparoscopic images designed for abdominal organ segmentation in surgical data science. The dataset comprises white-light laparoscopic images paired with dense, pixel-wise segmentation masks.

Annotations cover **eleven anatomical structures**: eight abdominal organs (colon, liver, pancreas, small intestine, spleen, stomach, ureter, and vesicular glands), the abdominal wall, and two vessel structures (inferior mesenteric artery and intestinal veins). The evaluation set (567 samples) consists of the multi-organ subset, in which all visible structures are annotated simultaneously, enabling consistent class-wise evaluation. The training set (12,628 samples) combines both the single-organ and multi-organ annotated subsets, with the 567 evaluation images excluded from the training pool. We do not explicitly verify procedure-level independence between the training and evaluation partitions and cannot rule out that both sets contain frames drawn from the same underlying surgical procedures. During training, Domain $\mathcal{A}$ (red-channel) and Domain $\mathcal{B}$ (RGB) images are sampled independently, with no enforced pairing.

Table 1. Quantitative comparison against baseline methods on the held-out evaluation set. Best results in **bold**.

Method	SSIM $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$
CycleGAN [25]	0.74	0.83	19.97	0.19
CUT [16]	0.70	0.85	20.02	0.17
Ours	0.81	0.87	22.51	0.14

Implementation Details Our framework is implemented in PyTorch [17]. All models are trained for 50 epochs using the Adam optimizer [12] with $\beta_1 = 0.5$, $\beta_2 = 0.999$, and an initial learning rate of 2×10^{-4} . Input images are resized to 256×256 pixels with a batch size of 1, following the standard CUT training protocol. The loss weights λ_{NCE} , λ_{cls} , and λ_{style} are set to 0.5, 0.5, and 10 respectively. The weights were empirically tuned to balance the contribution of each loss term.

Evaluation Protocol To validate the visual quality of the generated colorized images, we adopt a set of four complementary metrics spanning pixel-level fidelity, structural similarity, and deep perceptual similarity. The four metrics being: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [21], Multi-Scale SSIM (MS-SSIM) [22], and Learned Perceptual Image Patch Similarity (LPIPS) [23]. Since Domain $\mathcal{A}$ images are constructed by extracting the red channel directly from the original RGB images in the dataset, the corresponding full-colour RGB images serve as pixel-aligned ground-truth references during evaluation. This allows quantitative assessment even though the model is trained in a fully unpaired fashion.

Ablation Study To assess the contribution of each proposed component, we conduct an ablation study by incrementally adding each module to the CUT baseline and evaluating on the held-out evaluation set. The components under investigation are: (i) the class-wise contrastive loss ($\mathcal{L}_{\text{cls}}$); (ii) the segmentation-conditioned discriminator; (iii) the segmentation-conditioned generator.

3.2 Results and Discussion

Quantitative Results Using the evaluation protocol described in Section 3.1, we compare our method against two established unpaired image-to-image translation baselines: CycleGAN [25] and CUT [16]. All baselines are trained under identical settings on the same data splits to ensure a fair comparison. Quantitative results are reported in Table 1. Our method consistently outperforms both baselines across all four metrics. Against CycleGAN, we achieve gains in SSIM ($0.74 \rightarrow 0.81$), MS-SSIM ($0.83 \rightarrow 0.87$), PSNR ($19.97 \rightarrow 22.51$ dB), and LPIPS ($0.19 \rightarrow 0.14$). Compared to CUT, improvements are similarly consistent across SSIM ($0.70 \rightarrow 0.81$), MS-SSIM ($0.85 \rightarrow 0.87$), PSNR ($20.02 \rightarrow 22.51$ dB), and LPIPS ($0.17 \rightarrow 0.14$). The consistent gains across both pixel-level metrics (SSIM,




Fig. 2. Qualitative comparison of colorization results. Each row shows a representative surgical scene. Columns from left to right: input red-channel image (Domain $\mathcal{A}$), CycleGAN [25], CUT [16], our method, and ground-truth RGB image (Domain $\mathcal{B}$). Our method produces more uniform and anatomically consistent colorizations, particularly at organ boundaries and within homogeneous tissue regions. Best viewed in colour.

MS-SSIM, PSNR) and the deep perceptual metric (LPIPS) suggest that our method improves both structural fidelity and perceptual quality simultaneously - injecting anatomical context at every decoder layer preserves tissue-boundary coherence, the class-wise contrastive loss encourages consistent feature representations within each tissue class, while the class-aware discriminator and style loss align the generated chromatic distribution with real white-light laparoscopic imagery.

Results for the ablation study are reported in Table 2. The class-wise contrastive loss yields the largest single-step improvement, with SSIM rising from 0.70 to 0.79 and LPIPS improving from 0.17 to 0.14, confirming that harder within-class negative sampling has the strongest impact on both structural and perceptual quality. The segmentation-conditioned discriminator improves PSNR (+0.99 dB) but transiently increases LPIPS, suggesting that class-aware adversarial pressure without chromatic regularisation can introduce perceptual artefacts; the style loss resolves this, recovering LPIPS to 0.14 while delivering the largest remaining PSNR gain (+1.18 dB). The segmentation-conditioned generator contributes consistent gains in SSIM and MS-SSIM, reflecting improved

Table 2. Ablation study results. Each row incrementally adds one component to the CUT baseline. Best results in **bold**.

Method	SSIM $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$
CUT [16] (Baseline)	0.70	0.85	20.02	0.17
+ Class-wise Contrastive Loss	0.79	0.85	20.72	0.14
+ Seg.-Cond. Discriminator	0.80	0.85	21.71	0.18
+ Seg.-Cond. Generator	0.81	0.86	21.33	0.17
+ Style Loss (Full)	0.81	0.87	22.51	0.14

structural coherence at tissue boundaries. Each component contributes positively, with the full model achieving the best performance across all metrics.

Qualitative Results Figure 2 presents a visual comparison of colorizations produced by CycleGAN [25], CUT [16], and our method against the ground-truth RGB images. Our method consistently generates more naturalistic and anatomically coherent colorizations, particularly in regions with fine tissue boundaries and distinct organ classes such as the liver, colon, and stomach. Notably, our method produces more uniform colorization within individual tissue regions, with fewer abrupt color shifts across spatially continuous structures. This is most evident in large homogeneous regions such as the stomach surface, where baseline methods exhibit inconsistent hues, while our method maintains a coherent appearance throughout. We attribute this to the combination of mask injection at every decoder layer, which anchors colorization decisions to the anatomical layout. At organ boundaries, our method also produces sharper and more semantically meaningful color transitions, reflecting the class-aware feedback provided by the segmentation-conditioned discriminator.

Limitations and Future Work Our framework requires a segmentation mask as input. We use ground-truth masks here to isolate the effect of segmentation guidance; in deployment, this can be replaced by a pretrained surgical segmentation model, and future work will assess robustness to such predicted masks. Additionally, we approximate the source domain using red-channel images extracted from white-light RGB data rather than real red-laser acquisitions; future work will validate the framework on actual laser speckle imagery. Extending the framework to enforce temporal consistency across video frames is a further step toward real-time deployment in the operating room.

4 Conclusion

We presented a segmentation-guided unpaired image-to-image translation framework for colorizing simulated single-channel red-illumination laparoscopic images, with the goal of restoring visual navigability during laser-mode imaging without interfering with the perfusion assessment pipeline. By extending the

CUT [16] framework with three targeted contributions; mask injection at every decoder layer, a class-wise contrastive loss with harder within-class negatives, and a segmentation-conditioned discriminator paired with a style loss; our method produces anatomically coherent and perceptually natural colorizations that outperform CycleGAN [25] and CUT [16] across all evaluated metrics. Ablation studies confirm that each component contributes meaningfully to the overall performance.

Acknowledgments. This publication is part of the project AI4AI with file number P22.017 of the research programme Perspectief which is (partly) financed by the Dutch Research Council (NWO) under the grant TTW-Perspectief 2022-2023.

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

References

1. Briers, J.D.: Laser doppler, speckle and related techniques for blood perfusion mapping and imaging. *Physiological measurement* **22**(4), R35–R66 (2001)
2. Buia, A., Stockhausen, F., Hanisch, E.: Laparoscopic surgery: a qualified systematic review. *World journal of methodology* **5**(4), 238 (2015)
3. Carstens, M., Rinner, F.M., Bodenstedt, S., Jenke, A.C., Weitz, J., Distler, M., Speidel, S., Kolbinger, F.R.: The dresden surgical anatomy dataset for abdominal organ segmentation in surgical data science. *Scientific Data* **10**(1), 3 (2023)
4. Czarnota, M., Adamowicz, M., Górnaiak, A.: Minimally invasive surgery in modern medicine: A review of laparoscopic techniques, indications, and postoperative outcomes. *Journal of Education, Health and Sport* **83**, 61761–61761 (2025)
5. Draijer, M., Hondebrink, E., Van Leeuwen, T., Steenbergen, W.: Review of laser speckle contrast techniques for visualizing tissue perfusion. *Lasers in medical science* **24**(4), 639–651 (2009)
6. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2414–2423 (2016)
7. Heeman, W., Calon, J., van der Bilt, A., Pierie, J.P.E., Pereboom, I., van Dam, G.M., Boerma, E.C.: Dye-free visualisation of intestinal perfusion using laser speckle contrast imaging in laparoscopic surgery: a prospective, observational multi-centre study. *Surgical Endoscopy* **37**(12), 9139–9146 (2023)
8. Heeman, W., Wildeboer, A.C., Al-Taher, M., Calon, J.E., Stassen, L.P., Diana, M., Derikx, J.P., van Dam, G.M., Boerma, E.C., Bouvy, N.D.: Experimental evaluation of laparoscopic laser speckle contrast imaging to visualize perfusion deficits during intestinal surgery. *Surgical Endoscopy* **37**(2), 950–957 (2023)
9. Hren, R., Brezar, S.K., Marhl, U., Sersa, G.: Laser speckle contrast imaging of perfusion in oncological clinical applications: a literature review. *Radiology and Oncology* **58**(3), 326–334 (2024)
10. 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)

11. Kaneko, T., Funahashi, K., Ushigome, M., Kagami, S., Yoshida, K., Koda, T., Nagashima, Y., Miura, Y., Kurihara, A., Murakami, Y.: Noninvasive assessment of bowel blood perfusion using intraoperative laser speckle flowgraphy. *Langenbeck's Archives of Surgery* **405**(6), 817–826 (2020)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
13. Klevebro, F., Boshier, P.R., Low, D.E.: Application of standardized hemodynamic protocols within enhanced recovery after surgery programs to improve outcomes associated with anastomotic leak and conduit necrosis in patients undergoing esophagectomy. *Journal of Thoracic Disease* **11**(Suppl 5), S692 (2019)
14. Liu, D., Liang, L., Liu, L., Zhu, Z.: Does intraoperative indocyanine green fluorescence angiography decrease the incidence of anastomotic leakage in colorectal surgery? a systematic review and meta-analysis. *International journal of colorectal disease* **36**(1), 57–66 (2021)
15. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2794–2802 (2017)
16. 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)
17. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
18. Renna, M.S., Grzeda, M.T., Bailey, J., Hainsworth, A., Ourselin, S., Ebner, M., Vercauteren, T., Schizas, A., Shapey, J.: Intraoperative bowel perfusion assessment methods and their effects on anastomotic leak rates: meta-analysis. *British Journal of Surgery* **110**(9), 1131–1142 (2023)
19. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
20. Wang, Y.: Colorizing grayscale CT images of human lung using deep learning. Ph.D. thesis, Auckland University of Technology (2021)
21. 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)
22. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: *The thirty-seventh asilomar conference on signals, systems & computers, 2003. vol. 2*, pp. 1398–1402. Ieee (2003)
23. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
24. Zhang, Z., Li, Y., Shin, B.S.: Robust medical image colorization with spatial mask-guided generative adversarial network. *Bioengineering* **9**(12), 721 (2022)
25. 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)