

Do Not Break the Vessels: Structure-Preserving Mean Flow for Vascular Image Translation

Changjin Sun¹, Zhuo Hu², Kaini Wang³, Baixuan Wu⁴, Shuo Gao¹, Runan Zheng⁵, Cheng Xue¹, Yudong Zhang¹, and Guangquan Zhou^{11(✉)}

¹ School of Biological Science and Medical Engineering, Southeast University, Nanjing, China

guangquan.zhou@seu.edu.cn

² School of Computer Science and Engineering, Southeast University, Nanjing, China

³ Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China

⁴ Zhejiang University, Hangzhou, China

⁵ Suzhou Microclear Medical Instruments Co., Suzhou, China

Abstract. Reconstructing anatomically faithful vascular structures from clinically accessible imaging modalities is of substantial clinical significance. However, existing cross-modal translation methods mainly emphasize pixel-level fidelity or visual realism and treat structure preservation as a property of the final output rather than an invariant of the generative process. This limitation often leads to structural discontinuities and artifacts, compromising anatomical coherence and clinical reliability. In this work, we propose a *Structure-Preserving Mean Flow (SPMF)* framework that formulates vascular image translation as a topology-invariant transport process. From a structural invariance principle, we derive an orthogonality constraint on the flow velocity field that formally separates appearance transport from topological distortion, and implement it as a time-weighted surrogate objective within a Brownian-bridge diffusion model to preserve topology at every diffusion step. Moreover, we propose a *Prototype-Guided Structural Refinement (PGSR)* module to align degraded inference-time structures with reliable training-time structures. Experiments on paired NIRII-2PF and fundus datasets demonstrate consistent improvements over state-of-the-art methods, achieving peak PSNRs of 24.96 dB and 24.83 dB, respectively.

Keywords: Vascular Image Translation · Mean Flow · Diffusion Models

1 Introduction

Microvascular topology is a key biomarker for diagnosing and monitoring vascular diseases [6], making structure-faithful imaging a clinical necessity. Two-photon fluorescence (2PF) microscopy meets this requirement by resolving sub-micron vascular detail through nonlinear optical excitation, but its reliance on specialized pulsed lasers and point-scanning acquisition makes it impractical




Fig. 1. Illustration of the proposed structure-preserving mean flow. By constraining the velocity field to be orthogonal to the structural gradient, the generative dynamics can transform appearance while remaining on the vascular structural manifold, thereby preserving topology along the entire trajectory.

for routine clinical use. Near-infrared window II (NIRII) imaging offers a non-invasive, fast, and clinically compatible alternative [16, 7, 23, 4, 24, 14], yet severe tissue scattering and limited spatial resolution fundamentally degrade vascular topology in the acquired images (Fig. 1). Bridging the NIRII-to-2PF modality gap is therefore a structure-critical translation task that demands topologically faithful vascular recovery, not just visually plausible outputs.

Existing methods for cross-modal medical image translation, whether based on pixel-wise regression [13, 3, 18, 25, 5] or conditional generation [21, 8, 17], treat structure preservation as a property of the final output rather than an invariant of the generative process itself. Pixel-wise methods focus on intensity fidelity, while diffusion-based approaches prioritize visual realism and leave the generative trajectory unconstrained, making topological violations such as vessel breakage and spurious branches admissible outcomes. This problem is further compounded by a distributional gap between training and inference. Structural cues extracted from degraded NIRII images differ substantially from the clean structures available during training, making any output-level structural supervision unreliable even when it is explicitly applied.

Vascular topology is an intrinsic anatomical property that remains invariant across imaging modalities despite substantial appearance variation [22, 10]. We argue that a valid translation must preserve this invariance not only at the output but throughout the generative process. When image translation is modeled as a continuous flow driven by a time-dependent velocity field, the vesselness of the evolving image should remain constant along the trajectory, which in turn implies that the velocity field must be orthogonal to the structural gradient. Standard mean flow formulations impose no such constraint, leaving the velocity field free to alter appearance and topology alike. This distinction motivates our structure-preserving mean flow, in which the generative dynamics may freely transform appearance but are prevented from crossing the boundary of the structural manifold (Fig. 1).




Fig. 2. Overview of the proposed framework. (A) SPMF guides diffusion-based translation for topology-consistent generation, while (B) PGSR aligns degraded input vesselness with reliable training-time anatomical patterns.

Following this formulation, we propose a *Structure-Preserving Mean Flow (SPMF)* and instantiate it within a Brownian-bridge diffusion model [12]. At each diffusion step, the model predicts a clean reconstruction from the current noisy state, and we enforce that the predicted image maintains the same vascular topology as a reference anatomical prior throughout the trajectory. A time-dependent weighting schedule modulates this constraint according to the generative hierarchy of diffusion, imposing strong topology regularization during the early coarse-structure stage and progressively relaxing it as generation shifts to fine-grained appearance refinement. Since this per-step supervision requires a reliable structural reference, and vesselness maps extracted from degraded NIRII inputs are noisy and topologically incomplete, we further propose a *Prototype-Guided Structural Refinement (PGSR)* mechanism that retrieves nearest-neighbor prototypes from a library of training-time vascular patterns and aligns the input vesselness map through frequency-decomposed optimization with topology-sensitive weighting, projecting both training and inference structures onto a shared anatomical manifold.

Our contributions are three-fold. (1) We derive an orthogonality constraint on the flow velocity field from a structural invariance principle, formally separating appearance transport from topological distortion, and implement it as a time-weighted surrogate objective that preserves topology at every diffusion step. (2) We propose a prototype-guided structural refinement that bridges the training–inference structural gap by projecting noisy vesselness estimates onto a shared anatomical manifold learned from training data. (3) Experiments on paired NIRII–2PF and fundus datasets demonstrate consistent improvements over state-of-the-art methods in both structural fidelity and visual quality, with peak PSNRs of 24.96 dB and 24.83 dB, respectively.

2 Methodology

The proposed *Structure-Preserving Mean Flow (SPMF)* framework addresses NIRII-to-2PF translation from a topology-consistent flow perspective (Fig. 2). Section 2.1 formulates the problem setting and highlights the structural discrepancy between training and inference. Section 2.2 introduces the proposed SPMF, which constrains the generative dynamics to follow topology-consistent trajectories. Section 2.3 presents the *Prototype-Guided Structural Refinement (PGSR)* module that constructs a unified anatomical representation to support structure-preserving generation.

2.1 Problem Formulation

Given paired data $\mathcal{D} = \{(y, x)\}$, where y denotes a NIRII image and x the corresponding 2PF image, our goal is to learn a conditional model $p_\theta(x | y)$ that recovers 2PF-like appearance while preserving vascular topology. We extract vesselness maps using the Sato operator $\mathcal{S}(\cdot)$ [19, 9]. During training, reliable structures $s^{2\text{pf}} = \mathcal{S}(x)$ are available, while at inference only degraded estimates $s^{\text{nir}2} = \mathcal{S}(y)$ can be obtained.

Due to severe noise and loss of high-frequency details in NIRII imaging, the distributions of $s^{2\text{pf}}$ and $s^{\text{nir}2}$ differ substantially. Directly enforcing a constraint such as $\mathcal{S}(\hat{x}) \approx \mathcal{S}(y)$ is therefore ill-posed and often leads to unstable optimization and structural artifacts. Our objective is thus to learn a conditional generative model that (i) preserves vascular topology throughout the generative process and (ii) bridges the discrepancy between reliable training-time structures and degraded inference-time estimates via a unified structural representation.

2.2 Structure-Preserving Mean Flow (SPMF)

The proposed SPMF formulates the generative process as a continuous-time transport constrained by anatomical invariance. It models image translation as a flow driven by a time-dependent velocity field while explicitly restricting the dynamics to follow topology-consistent trajectories. Concretely, the generative process is described as a continuous-time flow governed by a time-dependent velocity field and a structure-preserving constraint,

$$\frac{dx(t)}{dt} = v_\theta(t, x(t)). \quad (1)$$

Structure-Preserving Principle. Let $x(t) \in \mathbb{R}^{H \times W}$ denote the generated image at time t . We impose the following principle:

Principle 1 (Structure Preservation). Along a valid generative trajectory, the underlying anatomical structure remains invariant:

$$\mathcal{S}(x(t)) = \mathcal{S}(x(0)), \quad \forall t \in [0, 1]. \quad (2)$$

This reflects the nature of NIRII-to-2PF translation, where appearance changes substantially while vascular topology should remain unchanged.

Constraint on Flow Dynamics. Taking the time derivative of Eq. (2) yields

$$\frac{d}{dt}\mathcal{S}(x(t)) = \nabla_x \mathcal{S}(x(t)) \cdot v_\theta(t, x(t)). \quad (3)$$

To preserve structure, the above term must vanish, leading to the constraint

$$\nabla_x \mathcal{S}(x(t)) \cdot v_\theta(t, x(t)) = 0, \quad (4)$$

which implies that the velocity field must lie in the tangent space of the structure manifold, allowing appearance changes while preventing topological distortion.

Surrogate Objective in Diffusion. Directly enforcing Eq. (4) is intractable in practice due to the non-differentiability of the vesselness operator. We therefore adopt a surrogate objective based on structural consistency within a Brownian-bridge diffusion framework [12].

Given a noisy sample x_t , the model predicts a reconstruction $\hat{x}_0(t)$. We enforce structure preservation by minimizing a time-weighted consistency loss:

$$\mathcal{L}_{\text{SP}}(\theta) = \mathbb{E}_t [w(t) \|\mathcal{S}(\hat{x}_0(t)) - \tilde{s}\|_1], \quad w(t) = 1 - \frac{t}{T}. \quad (5)$$

This weighting emphasizes structural consistency during early stages, when global topology is formed, and gradually relaxes the constraint to allow fine-grained appearance refinement. In this way, the diffusion process is guided to follow a structure-preserving mean flow.

2.3 Prototype-Guided Structural Refinement (PGSR)

PGSR provides a stable anatomical reference $\tilde{s}$ by projecting raw vesselness estimates onto a shared anatomical manifold learned from training data [10, 2]. Given an initial vesselness map $s = \mathcal{S}(\cdot)$, a refinement operator is defined as

$$\tilde{s} = \mathcal{R}_{\text{proto}}(s), \quad (6)$$

which produces a structurally consistent representation shared across training and inference stages.

Prototype Library and Retrieval. A prototype set $\mathcal{P} = p_{k=1}^K$ is constructed from vesselness maps of training samples, representing typical vascular patterns in terms of topology, thickness, and connectivity. For a given input structure s , the most similar prototypes are retrieved as references using a similarity measure in the vesselness space. These prototypes serve as anatomical anchors that regularize the refinement toward valid vascular configurations.

Frequency-Decomposed Structural Alignment. Both the input structure s and the retrieved prototype p are decomposed into low-, mid-, and high-frequency components, denoted as $\{s^{(f)}, p^{(f)}\}$ with $f \in \{\text{low}, \text{mid}, \text{high}\}$. While NIRII/2PF observations exhibit large variations in low- and mid-frequency appearance, vascular topology is primarily encoded in high-frequency responses,

such as thin branches and sharp boundaries. We therefore impose frequency-dependent weights and emphasize high-frequency alignment.

The refinement is obtained by minimizing the following objective:

$$\mathcal{L}_{\text{proto}} = \sum_{f \in \{\text{low}, \text{mid}, \text{high}\}} \alpha_f \left\| s^{(f)} - p^{(f)} \right\|_2^2 + \lambda_{\text{tv}} \|\nabla s\|_1, \quad (7)$$

where $(\alpha_{\text{low}}, \alpha_{\text{mid}}, \alpha_{\text{high}}) = (0.5, 1.0, 2.0)$, and the total variation term suppresses noise and encourages spatial continuity.

Unified Refinement for Training and Inference. The above optimization is performed for a small number of iterations, yielding a refined structure $\tilde{s}$ that lies on a valid anatomical manifold. Importantly, the same refinement process is applied during both training and inference:

$$\tilde{s}^{\text{train}} = \mathcal{R}_{\text{proto}}(\mathcal{S}(x)), \quad \tilde{s}^{\text{test}} = \mathcal{R}_{\text{proto}}(\mathcal{S}(y)). \quad (8)$$

This unified formulation reduces the structural distribution gap between training and testing stages and provides a stable anatomical prior for the proposed structure-preserving mean flow.

3 Experiments


Dataset and Experimental Setup. We evaluate the proposed method on a paired NIRII-2PF dataset (NIR2PF) with 763 image pairs and an external fundus dataset (Fundus) with 1001 patients. For Fundus, paired infrared (IR) and RGB images are used, where the IR image serves as a strong structural reference and the $4\times$ downsampled green channel is treated as a degraded structural observation. The Fundus experiment is used as an external generalization test under degraded structural observations, rather than as a direct replacement for the primary NIRII-to-2PF translation task. All images are normalized and split into training, validation, and test sets with a ratio of 6:2:2. We report PSNR and SSIM, and additionally use the Vesselness Mean Squared Error (V-MSE), defined as the mean squared error between the Sato vesselness responses of the generated and reference images, which serves as a practical proxy for structural accuracy by penalizing missing, blurred, or spurious vessel responses.

Implementation Details. Our method is implemented based on a Brownian Bridge Diffusion Model (BBDM)[12]. All images are resized to 256×256 and normalized to $[0, 1]$. The model is trained for 200 epochs (300K steps) using the Adam optimizer with a learning rate of 1×10^{-4} . Structure preservation is enforced via the proposed SPMF and prototype-guided refinement.

Comparison with State-of-the-Art Methods. We compare the proposed method with several representative image-to-image translation and super-resolution approaches, including SRGAN[11], SRN[5][20], the Brownian Bridge Diffusion Model (BBDM)[12], LDL[15], and SelfRDB[1]. All competing methods are re-trained on both datasets with official implementations and identical splits.

Table 1. Quantitative comparison on the NIR2PF and Fundus datasets. V-MSE is reported in units of 10^{-3} .

Method	NIR2PF			Fundus		
	PSNR $\uparrow$	SSIM $\uparrow$	V-MSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	V-MSE $\downarrow$
SRGAN [11]	24.39	0.5691	5.657	24.23	0.4657	13.19
SRN [5, 20]	24.49	0.6680	5.576	18.04	0.4570	13.63
BBDM [12]	22.39	0.6382	3.568	24.47	0.5204	12.02
LDL [15]	23.81	0.7215	3.634	24.32	0.5412	8.65
SelfRDB [1]	20.88	0.5875	5.565	22.97	0.5325	10.13
Proposed	24.96	0.6904	2.499	24.83	0.5823	7.76


**Fig. 3.** Visual comparison of vascular image translation results on the NIR2PF dataset. Existing methods show vessel breakage, over-smoothing, or spurious structures, whereas the proposed method preserves topologically consistent vascular reconstructions with clearer continuity.

The proposed method achieves the best or highly competitive performance on both the NIR2PF and Fundus datasets (Table 1). On NIR2PF, it attains the highest PSNR of 24.96 dB and the lowest V-MSE of 2.499×10^{-3} , clearly surpassing the BBDM baseline, which achieves 22.39 dB PSNR and 3.568×10^{-3} V-MSE. Although LDL reports a slightly higher SSIM of 0.7215, this improvement is accompanied by noticeably worse structural accuracy, as reflected by its higher V-MSE of 3.634×10^{-3} . On the Fundus dataset, the proposed method achieves the best PSNR of 24.83 dB, the best SSIM of 0.5823, and the lowest V-MSE of 7.76×10^{-3} , demonstrating strong robustness under severely degraded structural observations.

Visual Comparison. The proposed method consistently produces more topologically consistent and visually faithful vascular reconstructions (Fig. 3). Regression-based methods tend to over-smooth thin vessels and suppress fine branches, while diffusion-based methods without explicit structural constraints often introduce distorted or hallucinated patterns. Although LDL achieves a relatively high SSIM, its results exhibit evident over-smoothing and loss of fine branches. In

Table 2. Ablation study of different components.

BBDM	PGSR	SPMF	PSNR (dB) $\uparrow$	SSIM $\uparrow$	V-MSE($\times 10^{-3}$) $\downarrow$
$\checkmark$			22.39	0.6382	3.568
$\checkmark$	$\checkmark$		24.59	0.6769	3.124
$\checkmark$		$\checkmark$	23.31	0.6534	2.539
$\checkmark$	$\checkmark$	$\checkmark$	24.96	0.6904	2.499

**Fig. 4.** Component-wise analysis of the proposed framework. (a) PGSR and the training-inference structural gap. (b) Sensitivity to frequency weights. (c) Frequency-wise energy distribution. (d) Structural drift along the generative trajectory.

contrast, the proposed method better preserves vessel continuity and bifurcation geometry, yielding results closer to the ground truth.

Ablation Study. The ablation results on NIR2PF show that both PGSR and SPMF contribute to the overall performance (Table 2). Introducing PGSR yields a substantial improvement, raising PSNR from 22.39 dB to 24.59 dB and reducing V-MSE, which confirms the importance of a reliable anatomical prior aligned to the training manifold. Using SPMF alone also improves structural accuracy by constraining the generative dynamics, indicating the benefit of trajectory-level regularization even without explicit refinement. When both components are combined, the model achieves the best overall performance with 24.96 dB PSNR, 0.6904 SSIM, and the lowest V-MSE, clearly demonstrating their complementary effects.

Component-wise Analysis. The following analysis examines the roles of PGSR and SPMF from both structural and dynamical perspectives (Fig. 4). Fig. 4(a) shows that refined structures cluster near the training manifold, whereas raw inference-time structures remain scattered, indicating that PGSR reduces the structural distribution gap. Fig. 4(b) evaluates the sensitivity to frequency weights and shows that balanced or low-frequency-dominated settings lead to higher errors, while emphasizing high-frequency components reduces the V-MSE to approximately 2.499×10^{-3} , supporting the frequency-weighted refinement design. Fig. 4(c) reports the frequency-wise vesselness energy distribution, where NIRII is dominated by low-frequency responses while 2PF retains non-negligible mid- and high-frequency components, confirming that fine vascular structures are mainly encoded at higher frequencies. Fig. 4(d) analyzes the structural drift along the generative trajectory and shows that SPMF maintains a more sta-

ble drift than the baseline, indicating that structural consistency is enforced throughout the diffusion process rather than only at the final output.

4 Conclusion

NIRII-to-2PF image translation is important for clinically accessible vascular visualization. We propose a structure-preserving mean flow framework that encourages topology-consistent diffusion dynamics by combining prototype-guided refinement with trajectory-level structural constraints. While the current evaluation mainly relies on vesselness-based structural errors, future work will incorporate explicit skeleton- or centerline-level metrics, larger datasets, and expert assessment to further validate clinical utility.

Acknowledgments. This work was supported in part by the National Key Research and Development Program of China (2024YFF1206700, 2022YFC2404300).

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

References

1. Arslan, F., Kabas, B., Dalmaz, O., Ozbey, M., Çukur, T.: Self-consistent recursive diffusion bridge for medical image translation. *Medical Image Analysis* **106**, 103747 (2025)
2. Capitão, N.M.F., Zhao, Y., Zhang, Y., Geerts, N., Lopes, J.V., Tao, Q.: Anatomy-compliant medical image synthesis by latent diffusion models. In: *Medical Imaging with Deep Learning* (2024)
3. Chaudhary, S., Moon, S., Lu, H.: Fast, efficient, and accurate neuro-imaging denoising via supervised deep learning. *Nature communications* **13**(1), 5165 (2022)
4. Chen, M., Feng, Z., Fan, X., Sun, J., Geng, W., Wu, T., Sheng, J., Qian, J., Xu, Z.: Long-term monitoring of intravital biological processes using fluorescent protein-assisted nir-ii imaging. *Nature Communications* **13**(1), 6643 (2022)
5. Chen, R., Peng, S., Zhu, L., Meng, J., Fan, X., Feng, Z., Zhang, H., Qian, J.: Enhancing total optical throughput of microscopy with deep learning for intravital observation. *Small Methods* **7**(9), 2300172 (2023)
6. Cheung, C.Y.I., Ikram, M.K., Klein, R., Wong, T.Y.: The clinical implications of recent studies on the structure and function of the retinal microvasculature in diabetes. *Diabetologia* **58**(5), 871–885 (2015)
7. Choe, K., Hontani, Y., Wang, T., Hebert, E., Ouzounov, D.G., Lai, K., Singh, A., Béguelin, W., Melnick, A.M., Xu, C.: Intravital three-photon microscopy allows visualization over the entire depth of mouse lymph nodes. *Nature immunology* **23**(2), 330–340 (2022)
8. Fei, B., Li, Y., Yang, W., Gao, H., Xu, J., Ma, L., Yang, Y., Zhou, P.: A diffusion model for universal medical image enhancement. *Communications Medicine* **5**(1), 294 (2025)

9. Garret, G., Vacavant, A., Frindel, C.: Deep vessel segmentation based on a new combination of vesselness filters. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
10. Konz, N., Chen, Y., Dong, H., Mazurowski, M.A.: Anatomically-controllable medical image generation with segmentation-guided diffusion models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 88–98. Springer (2024)
11. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4681–4690 (2017)
12. Li, B., Xue, K., Liu, B., Lai, Y.K.: Bbdm: Image-to-image translation with brownian bridge diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition. pp. 1952–1961 (2023)
13. Li, X., Zhang, G., Wu, J., Zhang, Y., Zhao, Z., Lin, X., Qiao, H., Xie, H., Wang, H., Fang, L., et al.: Reinforcing neuron extraction and spike inference in calcium imaging using deep self-supervised denoising. *Nature methods* **18**(11), 1395–1400 (2021)
14. Li, Y., Cai, Z., Liu, S., Zhang, H., Wong, S.T., Lam, J.W., Kwok, R.T., Qian, J., Tang, B.Z.: Design of aiegens for near-infrared iib imaging through structural modulation at molecular and morphological levels. *Nature Communications* **11**(1), 1255 (2020)
15. Liang, J., Zeng, H., Zhang, L.: Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5657–5666 (2022)
16. Lu, R., Sun, W., Liang, Y., Kerlin, A., Bierfeld, J., Seelig, J.D., Wilson, D.E., Scholl, B., Mohar, B., Tanimoto, M., et al.: Video-rate volumetric functional imaging of the brain at synaptic resolution. *Nature neuroscience* **20**(4), 620–628 (2017)
17. Luo, Y., Yang, Q., Fan, Y., Qi, H., Xia, M.: Measurement guidance in diffusion models: Insight from medical image synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(12), 7983–7997 (2024)
18. Qiao, C., Li, D., Liu, Y., Zhang, S., Liu, K., Liu, C., Guo, Y., Jiang, T., Fang, C., Li, N., et al.: Rationalized deep learning super-resolution microscopy for sustained live imaging of rapid subcellular processes. *Nature biotechnology* **41**(3), 367–377 (2023)
19. Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., Kikinis, R.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical image analysis* **2**(2), 143–168 (1998)
20. Tao, X., Gao, H., Shen, X., Wang, J., Jia, J.: Scale-recurrent network for deep image deblurring. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8174–8182 (2018)
21. Xing, Z., Yang, S., Chen, S., Ye, T., Yang, Y., Qin, J., Zhu, L.: Cross-conditioned diffusion model for medical image to image translation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 201–211. Springer (2024)
22. Xu, M., Gupta, S., Hu, X., Li, C., Abousamra, S., Samaras, D., Prasanna, P., Chen, C.: Topocellgen: Generating histopathology cell topology with a diffusion model. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 20979–20989 (2025)

23. Yildirim, M., Sugihara, H., So, P.T., Sur, M.: Functional imaging of visual cortical layers and subplate in awake mice with optimized three-photon microscopy. *Nature communications* **10**(1), 177 (2019)
24. Zhang, M., Yue, J., Cui, R., Ma, Z., Wan, H., Wang, F., Zhu, S., Zhou, Y., Kuang, Y., Zhong, Y., et al.: Bright quantum dots emitting at 1,600 nm in the nir-iib window for deep tissue fluorescence imaging. *Proceedings of the National Academy of Sciences* **115**(26), 6590–6595 (2018)
25. Zhao, Y., Zhang, M., Zhang, W., Zhou, Y., Chen, L., Liu, Q., Wang, P., Chen, R., Duan, X., Chen, F., et al.: Isotropic super-resolution light-sheet microscopy of dynamic intracellular structures at subsecond timescales. *Nature Methods* **19**(3), 359–369 (2022)