

Test-Time Adaptation in Optical Coherence Tomography Using Trajectory-Aligned Time-Independent Flow

Veit Hucke^{*1,2[0009-0002-1519-1252]}, Thomas Pinetz^{*†1,2[0000-0002-6100-2136]},
Gregor Reiter^{3[0000-0001-7661-4015]}, Ursula
Schmidt-Erfurth^{4[0000-0002-7788-7311]}, and Hrvoje
Bogunović^{1,2[0000-0002-9168-0894]}

¹ Institute of Artificial Intelligence, Center for Medical Data Science, Medical University of Vienna, Austria

² Comprehensive Center for Artificial Intelligence in Medicine, Medical University of Vienna, Austria

³ Department of Ophthalmology and Optometry, Medical University of Vienna, Austria

⁴ Laboratory for Ophthalmic Image Analysis, Medical University of Vienna, Austria
`thomas.pinetz@meduniwien.ac.at`, `hrvoje.bogunovic@meduniwien.ac.at`

Abstract. Optical coherence tomography (OCT) is essential in ophthalmology, but inconsistent image quality especially in low-cost devices hinders automated analysis. To address this, we introduce a flow-matching-based test-time adaptation method that generates high-quality surrogate images from noisy inputs. Typically, domain gaps between test and training data cause pixel distribution mismatches during the denoising process. We overcome this by matching the test image’s histogram to synthetic reference trajectories, successfully aligning the input with expected distributions. Additionally, we remove the network’s time conditioning to account for slight deviations in real-world noise distributions. Our approach achieves state-of-the-art performance in segmenting critical biomarkers for two stages of Age-related Macular Degeneration (AMD). Code is available here: <https://github.com/Veit21/tta-flow>.

Keywords: Image Domain Adaptation · Flow Matching · Retina · OCT.

1 Introduction

Deep learning-based medical image segmentation has become a clinical staple for tracking biomarkers and detecting lesions [1]. However, model performance often deteriorates significantly when encountering minor distribution shifts, such as changes in imaging hardware [6]. This vulnerability is especially problematic with low-cost scanners, which introduce more severe noise characteristics and are rarely represented in standard training datasets.

[†] Corresponding Author * Equal contribution.

To address this, researchers have increasingly turned to test-time adaptation (TTA) [4, 22, 15]. Early methods updated pre-trained model weights during inference using target domain data and specialized loss functions [24, 25, 20]. These approaches typically relied on prediction entropy (e.g., TENT [24]) or generating pseudo-labels via augmentations and thresholding (e.g., COTTA [25]). However, developing loss functions to provide consistent improvement across diverse problems without access to the training distribution is difficult, and recent findings indicate that their performance improvements are often inconsistent, especially in segmentation tasks [27].

As an alternative, generative models have been used to adapt the test images directly, effectively projecting them onto the training data manifold for downstream tasks [11]. While earlier attempts utilized GANs [10] and normalizing flows [22], diffusion models are now predominantly used for this task [15, 5, 12, 13]. These methods generally treat the incoming test image as an intermediate point on a diffusion trajectory and reverse it, often incorporating a data-fidelity term to enforce consistency [12, 15, 4]. The primary flaw in this approach is that real-world noise is rarely ideal, fixed-magnitude Gaussian noise. Methods like DDA attempt to compensate by re-noising images to better fit the trajectory and applying data fidelity on downsampled versions [5]. Unfortunately, down-sampling risks blurring or eliminating small lesions, critical for diagnosis. Others have incorporated large language models and CLIP-based method to synthesize realistic surrogates [11], however this requires disease specific models that understand the modality well.

Recently, flow matching [16, 17] has emerged as a robust alternative to classical diffusion models like DDPM [9] improving the state-of-the-art in natural images [18] and in the medical domain [26]. While variational setups have successfully combined flow matching with data-consistency priors for natural image restoration [18], they still assume purely synthetic degradations. To make flow matching viable for real-world low-cost devices, we introduce two key modifications:

1. **Trajectory Alignment:** We sample a reference trajectory histogram and use histogram matching to align the incoming test image, effectively bridging the gap between theoretical and actual noise distributions.
2. **Time-Independent Flow:** Because actual noise levels deviate from the reference trajectory, we intentionally omit time conditioning in the flow network. As demonstrated by He et al. [23], allowing the network to implicitly manage unknown real-world noise levels, rather than forcing a specific time condition, yields superior results.

We demonstrate the effectiveness of our approach, TTA-Flow, on two clinical biomarkers for Age-related Macular Degeneration (AMD) in both 3D-to-3D and 3D-to-2D segmentation tasks using two different low-cost Optical Coherence Tomography (OCT) devices.




Fig. 1. Schematic overview of the TTA-Flow framework. The process begins by (a) generating reference trajectories and (b, c) calculating an average histogram $\bar{H}_s$ for each time step. During inference, (d) an incoming test sample ζ is matched to the corresponding reference histogram before being processed by the flow matching network. Finally, (e) the reconstructed sample z_S is utilized for downstream evaluation.

2 Methods

A comprehensive visual overview of our proposed approach, TTA-Flow, is provided in Fig. 1.

Flow Matching for Image Reconstruction Building on the generative flow-matching (FM) framework introduced by Lipman et al. [16, 17], our approach learns a time-dependent vector field $\mathbf{v}_\theta$. This vector field is designed to transport a known Gaussian source distribution $p_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to the empirical data distribution of our training set, $y_i \in \mathcal{D}$ for $i \in [0, N]$. This transport is defined by a deterministic flow $\{\mathbf{z}_t\}_{t \in [0, 1]}$ governed by the ordinary differential equation (ODE):

$$\frac{d\mathbf{z}_t}{dt} = \mathbf{v}_t(\mathbf{z}_t), \quad \mathbf{z}_0 \sim p_0, \quad (1)$$

which induces a probability path p_t satisfying the continuity equation $\partial_t p_t(\mathbf{z}) + \nabla \cdot (p_t(\mathbf{z}) \mathbf{v}_t(\mathbf{z})) = 0$.

To create a tractable training objective, flow matching requires defining a specific probability path between the source noise $\mathbf{z}_0$ and a target data sample

$\mathbf{y}_i$. We utilize a deterministic linear interpolation for this path:

$$\mathbf{z}_t = (1 - t)\mathbf{z}_0 + t\mathbf{y}_i, \quad t \in [0, 1], \quad (2)$$

This formulation establishes the target velocity as $\mathbf{u}_t(\mathbf{z}_t; \mathbf{z}_0, \mathbf{y}_i) = \mathbf{y}_i - \mathbf{z}_0$. The network is then trained to minimize the regression loss against this target vector field:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t, \mathbf{z}_0, \mathbf{x}_1} \left[\|\mathbf{v}_\theta(\mathbf{z}_t) - \mathbf{u}_t(\mathbf{z}_t; \mathbf{z}_0, \mathbf{y}_i)\|_2^2 \right]. \quad (3)$$

Rather than predicting the velocity directly, we parameterize our model to predict the target sample $\mathbf{x}_\theta(\mathbf{z}_t) \approx \mathbf{y}_i$ [14]. The corresponding velocity can then be recovered via:

$$\mathbf{v}_\theta(\mathbf{z}_t) = \frac{\mathbf{x}_\theta(\mathbf{z}_t) - \mathbf{z}_t}{1 - t}. \quad (4)$$

This recovered velocity $\mathbf{v}_\theta$ is then substituted back into the objective function $\mathcal{L}_{\text{FM}}$. At test time, we conceptualize the noisy input sample ζ as an intermediate point interpolated between pure noise and a theoretical high-quality acquisition y_k (e.g., $\zeta \approx tz_k + (1 - t)y_k$). However, the true noise characteristics of ζ rarely match the exact theoretical noise of a random trajectory at time t . Because of this discrepancy, we intentionally omit explicit noise-level (or time) conditioning from the network. As demonstrated by He et al. [23], removing this constraint prevents the network from forcing inaccurate assumptions and yields better results when the true noise level of the test image is unknown.

During inference, we generate novel samples by solving the initial value problem starting from pure noise with an Euler solver for $S \in \mathbb{N}$ steps:

$$\mathbf{z}_{i+1} = \mathbf{z}_i + \frac{1}{S}\mathbf{v}_\theta(\mathbf{z}_i), \quad \mathbf{z}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (5)$$

Reference generation As previously noted, the underlying statistical properties of a test sample ζ will inherently differ from the trajectories observed during training. To account for this discrepancy, we use our trained flow matching model to simulate n sample trajectories according to Eq. 5. From the resulting Sn generated images, we calculate the mean target intensity histograms, denoted as $\{\bar{H}_s\}_{s \in [1, S]}$, for each discrete time step. This sequence of average histograms establishes the reference trajectory for our specific flow model (Fig. 1b, c).

Inference During inference, we utilize this pre-computed sequence of average target histograms $\{\bar{H}_s\}_{s \in [1, S]}$ to normalize the incoming test data. For a given test sample ζ , we apply histogram matching as a preprocessing step. This aligns the sample’s intensity distribution with a specific reference histogram $\bar{H}_{s_{\text{target}}}$, where s_{target} is a predetermined hyperparameter. This alignment produces an intermediate sample that closely mirrors the expected data distribution of the reference trajectories at time s_{target} (Fig. 1d). To obtain the clean reconstructed estimate z_S , we initialize the ODE (Eq. 5) using z_s at step $s = s_{\text{target}}$. We then execute the remaining $S - s_{\text{target}}$ iterations of the flow matching process

(Eq. 5), effectively mapping the image into the target distribution. Notably, our experiments demonstrate that this approach performs robustly even without the inclusion of explicit data-fidelity terms. Finally, the reconstructed image z_S is passed to downstream tasks to produce the final output z_{seg} (Fig. 1e).

3 Experimental Setup

Datasets: We evaluate our approach on two Optical Coherence Tomography (OCT) datasets containing distinct clinically relevant pathologies and multiple scanners. First, the public *RETOUCH* [2] dataset contains annotations for intraretinal fluid (IRF), subretinal fluid (SRF), and pigment epithelial detachment (PED). It comprises 70 total volumes acquired across three devices: Spectralis (24 vols, 49 B-scans/vol), Topcon (22 vols, 128 B-scans/vol), and Cirrus (24 vols, 128 B-scans/vol). Second, we utilize an *in-house paired dataset* [3] of 40 patients collected at Medical University of Vienna diagnosed with geographic atrophy (GA), a late stage of Age-related macular degeneration (AMD). For each patient, paired volumes were acquired using both Topcon (128 B-scans/vol) and Spectralis (49 or 97 B-scans/vol) scanners, with manual annotations of GA provided on every alternating B-scan. We consider Topcon and Cirrus to be low-cost devices and adapt their images towards the Spectralis device, which has notably higher SNR. To establish a controlled baseline and ensure reproducibility, all generative models are trained exclusively on the RETOUCH Spectralis subset. We restrict the quantitative analysis to the annotated B-scans.

Flow Matching Setup: We employ a U-Net backbone [21] ($\sim 103\text{M}$ parameters) for the flow-matching network. Models are trained for 300,000 iterations using the Adam optimizer (learning rate 10^{-4} , weight decay 10^{-2} , batch size 16) with an exponential moving average (EMA) decay of 0.999. All input B-scans are resized to 224×224 pixels, normalized to $[-1, 1]$, and augmented with random horizontal flipping. During inference, the probability-flow ODE (Eq. 5) is integrated over $S = 100$ steps.

Downstream Segmentation Models: For fluid lesion segmentation, we employ a UNet++ architecture [28] featuring a ResNet-18 [7] encoder. To accurately gauge the quality of the reconstruction algorithms, this downstream model is trained solely on the RETOUCH Spectralis subset. We establish device-specific upper performance bounds (Upper) for the Topcon and Cirrus scanners by training dedicated, supervised models and evaluating them via 4-fold cross-validation. For GA segmentation, we use a pre-trained 3D-to-2D CNN based on [19]. Because our in-house dataset contains paired scans, we can directly compute the Spectralis upper bound without requiring further training.

Baselines and Hardware: We compare our approach against classical test-time adaptation strategies (TENT [24] and CoTTA [25]) and two diffusion-based models (CPDM [12] and DDA [5]). Additionally, we include a baseline using standard flow matching [14] without our modifications. Hyperparameters for all methods were chosen by grid search; for our method, this involved selecting the optimal target time point s_{target} . All models were implemented in PyTorch

Table 1. Downstream lesion segmentation performance on Cirrus $\rightarrow$ Spectralis and Topcon $\rightarrow$ Spectralis. Dice similarity coefficient (DSC; higher is better) for IRF, SRF, PED, and GA. FID (lower is better) is also reported for Topcon.* denotes statistical significance (wilcoxon test $p < 0.05$) of Ours_{uncond.} and its closest competitor.

Method	RETOUCH				In-house		
	IRF $\uparrow$	SRF $\uparrow$	PED $\uparrow$	Mean $\uparrow$	GA DSC $\uparrow$	FID $\downarrow$	
Cirrus $\rightarrow$ Spectralis							
Baseline	51.7 $\pm$ 27.0	43.2 $\pm$ 31.0	22.5 $\pm$ 32.1	39.1	–	–	
CPDM [12]	55.2 $\pm$ 27.7	48.5 $\pm$ 29.9	28.9 $\pm$ 33.2	44.2	–	–	
DDA [5]	50.0 $\pm$ 28.3	45.2 $\pm$ 30.5	25.6 $\pm$ 32.5	40.3	–	–	
FM [14]	57.1 $\pm$ 24.6	56.3 $\pm$ 24.6	39.1 $\pm$ 30.7	50.8	–	–	
CoTTA [25]	52.0 $\pm$ 22.3	64.5 $\pm$ 19.0	36.5 $\pm$ 28.6	51.0	–	–	
TENT [24]	51.5 $\pm$ 22.2	64.4 $\pm$ 19.0	35.5 $\pm$ 27.6	50.5	–	–	
Ours _{uncond.}	58.8 $\pm$ 18.6	66.3 $\pm$ 17.2	50.6 $\pm$ 27.0	58.6*	–	–	
Ours _{cond.}	<u>57.4 $\pm$ 18.5</u>	<u>64.7 $\pm$ 17.8</u>	<u>45.8 $\pm$ 27.9</u>	<u>56.0</u>	–	–	
Upper	69.5 $\pm$ 17.0	63.1 $\pm$ 27.6	52.1 $\pm$ 35.3	61.6	–	–	
Topcon $\rightarrow$ Spectralis							
Baseline	49.9 $\pm$ 26.0	42.3 $\pm$ 31.0	28.0 $\pm$ 18.8	40.0	43.7 $\pm$ 24.8	550.5	
CPDM [12]	43.8 $\pm$ 26.4	45.3 $\pm$ 31.0	38.2 $\pm$ 22.0	43.8	52.4 $\pm$ 23.2	100.1	
DDA [5]	38.0 $\pm$ 27.7	33.6 $\pm$ 28.5	30.6 $\pm$ 18.0	34.1	53.2 $\pm$ 21.7	148.4	
FM [14]	53.3 $\pm$ 25.4	49.3 $\pm$ 25.9	42.5 $\pm$ 26.7	48.3	49.4 $\pm$ 25.0	131.2	
CoTTA [25]	<u>53.5 $\pm$ 21.7</u>	47.7 $\pm$ 23.4	32.7 $\pm$ 21.0	44.6	N/A	N/A	
TENT [24]	53.9 $\pm$ 22.4	47.9 $\pm$ 24.5	33.3 $\pm$ 20.4	45.0	48.6 $\pm$ 31.9	N/A	
Ours _{uncond.}	51.6 $\pm$ 23.1	<u>52.2 $\pm$ 16.5</u>	49.2 $\pm$ 20.9	51.0	<u>56.1* $\pm$ 23.3</u>	32.2	
Ours _{cond.}	52.0 $\pm$ 24.5	52.6 $\pm$ 18.6	<u>44.8 $\pm$ 20.3</u>	<u>49.8</u>	57.1 $\pm$ 22.6	<u>85.4</u>	
Upper	60.4 $\pm$ 22.7	53.9 $\pm$ 24.6	56.2 $\pm$ 18.4	56.8	77.3 $\pm$ 13.6	0.0	

and trained on a single NVIDIA RTX A6000 (48GB) GPU. Training the flow-matching model took approximately 80 hours. Inference time per B-scan averages 2.24 and 2.20 seconds with and without histogram matching. Generating the initial $n = 100$ reference trajectories is a one-time process that takes roughly 231 seconds.

Evaluation Protocol: To evaluate the quality of the image reconstructions (Cirrus $\rightarrow$ Spectralis and Topcon $\rightarrow$ Spectralis), we measure the performance of the downstream segmentation models. Specifically, we report the DSC between the networks’ predictions and the ground-truth annotations. We assess the segmentation of IRF, SRF, and PED lesions for the RETOUCH dataset, and we quantify the segmentation of GA for the paired in-house dataset. For the paired in-house dataset we also report the Frechet Inception Distance [8] (FID) to circumvent the problem of pixel-wise correspondences.

4 Results and Discussion

Image reconstruction: Table 1 summarizes the downstream segmentation performance for the Cirrus $\rightarrow$ Spectralis and Topcon $\rightarrow$ Spectralis reconstruction tasks,




Fig. 2. Comparison to the state-of-the-art on downstream fluid and GA segmentation on the **Cirrus** (first two rows) and **Topcon** (remaining rows) data. For the GA segmentation, a white horizontal line denotes the position of the corresponding bscan (In-house, top row). Note that only every second bscan is annotated in the GA GT.

respectively. On the Cirrus data, our unconditional method achieves a state-of-the-art mean DSC of 58.6, significantly outperforming the baseline methods, and approaching the supervised upper bound.

The Topcon $\rightarrow$ Spectralis shift appears to be more challenging, indicated by systematically lower scores. Despite this difficulty, our approach maintains the highest mean DSC (51.0) on the RETOUCH dataset. While TENT slightly outperforms on IRF lesions, it fails to generalize to other structures like PED, where our method remains stable. Fig. 2 provides a qualitative comparison of our reconstructions against baseline generative methods. For both Cirrus and Topcon data, our approach yields high-quality reconstructions that preserve the original anatomy while visibly reducing noise and maintaining high contrast. The improved reconstruction translates to the downstream task, resulting in




Fig. 3. Comparison of downstream lesion segmentation performance (DSC) for different reference time points s_{target} . For each target, the results of the conditional and unconditional networks are shown. *denotes statistical significance.

fluid segmentations that align well with the ground-truth annotations. Similarly, for the GA task (Fig. 2, bottom), our approach notably reduces segmentation errors compared to all baselines, demonstrating a significant improvement over its closest competitor DDA.

Ablation Study: To validate our architectural design, we evaluate the impact of the histogram matching module and the time-conditioning parameterization. First, integrating the histogram matching module provides a significant performance gain over plain Flow Matching (FM), improving the mean DSC from 50.8 to 58.6 and from 49.4 to 56.1 on the Cirrus and the in-house data, respectively. Second, comparing the noise-conditional (Ours_{cond.}) and unconditional (Ours_{uncond.}) networks reveals a trade-off in Table 1. Dropping the time-conditioning forces the network to infer the noise state implicitly, yielding better overall fluid segmentation. As illustrated in Fig. 3, this unconditional approach consistently outperforms the conditional version and exhibits low sensitivity to the starting point s_{target} . In fact, oracle selection (choosing the absolute best hyperparameter per individual patient) improves overall results by only 1 percentage point. Additionally, the unconditioned model produces superior perceptual quality, reducing the FID from 85.4 to 32.2. However, the conditional model offers a non-significant benefit on the specific GA task (DSC 57.1 vs. 56.1).

5 Conclusion

In this study, we introduced a novel, Flow Matching-based framework TTA-Flow for high-quality OCT image reconstruction. We demonstrated that histogram matching is a simple yet highly effective preprocessing technique for aligning the intensity distribution of test data with the network’s expectations. Because our reconstruction process integrates the probability-flow ODE of a generative

model, the required reference histograms can be extracted at virtually no additional computational cost from synthetic trajectories. Furthermore, our results indicate that removing time-conditioning from the flow-matching network enhances denoising performance, leading to superior fluid segmentation and visual fidelity compared to standard conditional setups. While our setup without the usage of data-fidelity terms can theoretically hallucinate structures, we did not observe this in practice and data-fidelity terms are very sensitive to the hyperparameters. Notably, our method outperforms multiple baselines that explicitly use data-fidelity terms, including CPDM [13], which utilizes a fidelity term strictly designed for OCT. Currently, the optimal reference time point for histogram matching is set globally as a fixed hyperparameter. However, our future work will explore dynamically estimating the noise level directly from the input image as well as speeding up the adaptation process. This advancement will enable the framework to adaptively choose the ideal histogram-matching target for each volume, fully automating the pipeline and optimizing reconstruction fidelity.

Acknowledgements : Funded by the European Union (I-SCREEN, grant no. 101130093), EIC-2023-PATHFINDEROPEN-01 . Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Innovation Council and SMEs Executive Agency (EISMEA). Neither the European Union nor the granting authority can be held responsible for them.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Azad, R., Aghdam, E.K., Rauland, A., Jia, Y., Avval, A.H., Bozorgpour, A., Karim-ijafarbigloo, S., Cohen, J.P., Adeli, E., Merhof, D.: Medical image segmentation review: The success of U-Net. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(12), 10076–10095 (2024)
2. Bogunović, H., Venhuizen, F., Klimscha, S., Apostolopoulos, S., Bab-Hadiashar, A., Bagci, U., Beg, M.F., Bekalo, L., Chen, Q., Ciller, C., Gopinath, K., Gostar, A.K., Jeon, K., Ji, Z., Kang, S.H., Koozekanani, D.D., Lu, D., Morley, D., Parhi, K.K., Park, H.S., Rashno, A., Sarunic, M., Shaikh, S., Sivaswamy, J., Tennakoon, R., Yadav, S., De Zanet, S., Waldstein, S.M., Gerendas, B.S., Klaver, C., Sánchez, C.I., Schmidt-Erfurth, U.: RETOUCH: The retinal oct fluid detection and segmentation benchmark and challenge. *IEEE Transactions on Medical Imaging* **38**(8), 1858–1874 (2019). <https://doi.org/10.1109/TMI.2019.2901398>
3. Eidenberger, A., Birner, K., Frank-Publig, S., Schrittwieser, J., Tratnig-Frankl, M., Gumpinger, M., Schmidt-Erfurth, U.: Comparison of choroidal hypertransmission and retinal pigment epithelium loss for quantification of geographic atrophy across commonly used SD-OCT devices. *Scientific Reports* p. 7240 (2026). <https://doi.org/10.1038/s41598-026-38182-7>
4. Fazeakas, B., Pinetz, T., Aresta, G., Emre, T., Bogunović, H.: GARD: Gamma-based anatomical restoration and denoising for retinal OCT. In: *International Workshop on Ophthalmic Medical Image Analysis*. pp. 32–42. Springer (2025)

5. Gao, J., Zhang, J., Liu, X., Darrell, T., Shelhamer, E., Wang, D.: Back to the source: Diffusion-driven adaptation to test-time corruption. In: CVPR. pp. 11786–11796 (2023)
6. Gomariz, A., Kikuchi, Y., Li, Y.Y., Albrecht, T., Maunz, A., Ferrara, D., Lu, H., Goksel, O.: Joint semi-supervised and contrastive learning enables domain generalization and multi-domain segmentation. *Medical Image Analysis* **103**, 103575 (2025)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
8. 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. *NeurIPS* **30** (2017)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* **33**, 6840–6851 (2020)
10. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: ICML. vol. 80, pp. 1989–1998. PMLR (2018)
11. Li, M., Li, S., Su, T., Yuan, L., Liang, J., Li, W.: Exploring structured semantic priors underlying diffusion score for test-time adaptation. In: *NeurIPS*. vol. 37, pp. 12164–12194 (2024)
12. Li, S., Higashita, R., Fu, H., Li, H., Niu, J., Liu, J.: Content-preserving diffusion model for unsupervised AS-OCT image despeckling. In: MICCAI. pp. 660–670. Springer (2023)
13. Li, S., Higashita, R., Fu, H., Yang, B., Liu, J.: Score prior guided iterative solver for speckles removal in optical coherent tomography images. *IEEE Journal of Biomedical and Health Informatics* **29**(1), 248–258 (2025)
14. Li, T., He, K.: Back to basics: Let denoising generative models denoise. In: CVPR. pp. 36115–36125 (2026)
15. Lin, X., Du, C., Wu, Q., Tian, X., Yu, J., Zhang, Y., Wei, H.: Zero-shot low-field MRI enhancement via denoising diffusion driven neural representation. In: MICCAI. pp. 775–785. Springer (2024)
16. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: ICLR (2023), <https://openreview.net/forum?id=PqvMRDCJT9t>
17. Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R.T.Q., Lopez-Paz, D., Ben-Hamu, H., Gat, I.: Flow matching guide and code (2024), <https://arxiv.org/abs/2412.06264>
18. Martin, S., Gagneux, A., Hagemann, P., Steidl, G.: PnP-Flow: Plug-and-play image restoration with flow matching. In: ICLR. pp. 45466–45492 (2025)
19. Morano, J., Aresta, G., Lachinov, D., Mai, J., Schmidt-Erfurth, U., Bogunović, H.: Self-supervised learning via inter-modal reconstruction and feature projection networks for label-efficient 3D-to-2D segmentation. In: MICCAI. pp. 589–599. Springer (2023). https://doi.org/10.1007/978-3-031-43901-8_56
20. Pinetz, T., Hucke, V., Bogunović, H.: Exploiting intermediate reconstructions in optical coherence tomography for test-time adaption of medical image segmentation. In: MIDL. vol. 315, pp. 1081–1094. PMLR (2026)
21. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241. Springer (2015). https://doi.org/10.1007/978-3-319-24574-4_28

22. Safdari, R., Nikouei Mahani, M.A., Koohi-Moghadam, M., Bae, K.T.: MixStyle-Flow: Domain generalization in medical image segmentation using normalizing flows. In: MICCAI. pp. 376–385. Springer (2025)
23. Sun, Q., Jiang, Z., Zhao, H., He, K.: Is noise conditioning necessary for denoising generative models? In: ICML (2025), <https://openreview.net/forum?id=pTSWi6RTtJ>
24. Wang, D., Shelhamer, E., Liu, S., Bruno, O., Darrell, T.: TENT: Fully test-time adaptation by entropy minimization. In: ICLR (2021), <https://openreview.net/forum?id=uXl3bZLkr3c>
25. Wang, Q., Fink, O., Van Gool, L., Dai, D.: Continual test-time domain adaptation. In: CVPR. pp. 7191–7201 (2022). <https://doi.org/10.1109/CVPR52688.2022.00706>
26. Yazdani, M., Medghalchi, Y., Ashrafian, P., Hacıhaliloglu, I., Shahriari, D.: Flow matching for medical image synthesis: Bridging the gap between speed and quality. In: MICCAI. pp. 216–226. Springer (2025)
27. Yi, C., Chen, H., Zhang, Y., Xu, Y., Zhou, Y., Cui, L.: From question to exploration: Can classic test-time adaptation strategies be effectively applied in semantic segmentation? In: ACMML. pp. 10085–10094 (2024). <https://doi.org/10.1145/3664647.3680910>
28. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: UNet++: A nested U-Net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. pp. 3–11. Springer (2018)