

DiffuseIR: Diffusion Models For Isotropic Reconstruction of 3D Microscopic Images

Mingjie Pan^{1*}, Yulu Gan^{1*}, Fangxu Zhou¹,
Jiaming Liu¹, Aimin Wang¹, Shanghang Zhang¹, Dawei Li¹,
¹Peking University

Abstract. Three-dimensional microscopy is often limited by anisotropic spatial resolution, resulting in lower axial resolution than lateral resolution. Current State-of-The-Art (SoTA) isotropic reconstruction methods utilizing deep neural networks can achieve impressive super-resolution performance in fixed imaging settings. However, their generality in practical use is limited by degraded performance caused by artifacts and blurring when facing unseen anisotropic factors. To address these issues, we propose DiffuseIR, an unsupervised method for isotropic reconstruction based on diffusion models. First, we pre-train a diffusion model to learn the structural distribution of biological tissue from lateral microscopic images, resulting in generating naturally high-resolution images. Then we use low-axial-resolution microscopy images to condition the generation process of the diffusion model and generate high-axial-resolution reconstruction results. Since the diffusion model learns the universal structural distribution of biological tissues, which is independent of the axial resolution, DiffuseIR can reconstruct authentic images with unseen low-axial resolutions into a high-axial resolution without requiring re-training. The proposed DiffuseIR achieves SoTA performance in experiments on EM data and can even compete with supervised methods.

Keywords: Isotropic reconstruction · Unsupervised method · Diffusion model

1 Introduction

Three-dimensional (3D) microscopy imaging is crucial in revealing biological information from the nanoscale to the microscale. Isotropic high resolution across all dimensions is desirable for visualizing and analyzing biological structures. However, most three-dimensional imaging techniques often have lower axial (z) resolution than lateral (xy) resolution, due to physical slicing interval limitation (serial section transmission electron microscopy, automated tape-collecting ultra-microtome scanning electron microscopy, etc.) [18] or time-saving consideration (focused ion beam scanning electron microscopy, fluorescence microscopy, etc.) [8,31,28,23]. Therefore, effective isotropic super-resolution algorithms are critical for high-quality 3D image reconstructions, such as electron microscopy and fluorescence microscopy.

* Equal contribution

Recently, deep learning methods have made significant progress in image analysis [9,25,14,13]. To address the isotropic reconstruction problem, [9] employs isotropic EM images to generate HR-LR pairs at axial and train a super-resolution model in a supervised manner, demonstrating the feasibility of inferring HR structures from LR images. [29,30] use 3D point spread function (PSF) as a prior for self-supervised super-resolution. However, isotropic high-resolution images or 3D point spread function (PSF) physical priors are difficult to obtain in practical settings, thus limiting these algorithms. Some methods like [3,21] have skillfully used cycleGAN [32] architecture to train axial super-resolution models without depending on isotropic data or physical priors. They learn from unpaired matching between high-resolution 2D slices in the lateral plane and low-resolution 2D slices in the axial plane, achieving impressive performance. However, these methods train models in fixed imaging settings and suffer from degraded performance caused by artifacts and blurring when facing unseen anisotropic factors. This limits their generality in practice [6]. In conclusion, a more robust paradigm needs to be proposed. Recently, with the success of the diffusion model in the image generation field [26,19,4,17,11], researchers applied the diffusion model to various medical image generation tasks and achieved impressive results [25,22,1,12,20]. Inspired by these works, we attempt to introduce diffusion models to address the isotropic reconstruction problem.

This paper proposes DiffuseIR, an unsupervised method based on diffusion models, to address the isotropic reconstruction problem. Unlike existing methods, DiffuseIR does not train a specific super-resolution model from low-axial-resolution to high-axial-resolution. Instead, we pre-train a diffusion model ϵ_θ to learn the structural distribution $p_\theta(X_{lat})$ of biological tissue from lateral microscopic images X_{lat} , which resolution is naturally high. Then, as shown in Fig. 1, we propose a Sparse Spatial Condition Sampling (SSCS) to condition the reverse-diffusion process of ϵ_θ . SSCS extracts sparse structure context from low-axial-resolution slice x_{axi} and generate reconstruction result $x_0 \sim p_\theta(X_{lat}|x_{axi})$. Since ϵ_θ learns the universal structural distribution p_θ , which is independent of the axial resolution, DiffuseIR can leverage the flexibility of SSCS to reconstruct authentic images with unseen anisotropic factors without requiring re-training. To further improve the quality of reconstruction, we propose a Refine-in-loop strategy to enhance the authenticity of image details with fewer sampling steps.

To sum up, our contributions are as follows:

- (1) We are the first to introduce diffusion models to isotropic reconstruction and propose DiffuseIR. Benefiting from the flexibility of SSCS, DiffuseIR is naturally robust to unseen anisotropic spatial resolutions.
- (2) We propose a Refine-in-loop strategy, which maintains performance with fewer sampling steps and better preserves the authenticity of the reconstructed image details.
- (3) We perform extensive experiments on EM data with different imaging settings and achieve SOTA performance. Our unsupervised method is competitive with supervised methods and has much stronger robustness.

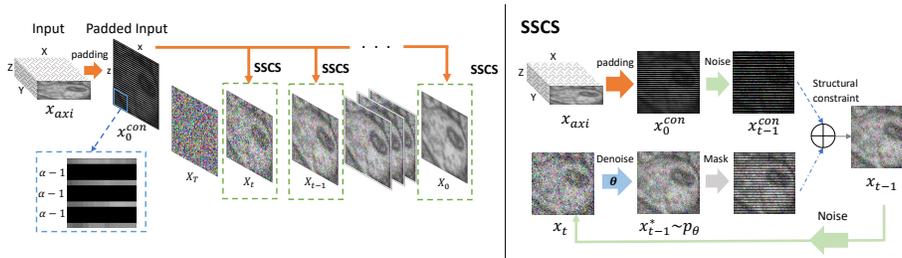


Fig. 1: **Method Pipeline.** DiffuseIR progressively conditions the denoising process with SSCS. For SSCS, we perform intra-row padding on input X_{lat} using the anisotropy factor α to obtain spatially aligned structural context, which is then merged with the diffusion model’s output. Iterative SSCS refines reconstruction.

2 Methodology

As shown in Fig. 1, DiffuseIR address isotropic reconstruction by progressively conditions the denoising process of a pre-trained diffusion model ϵ_θ . Our method consists of three parts: DDPM pre-train, Sparse Spatial Condition Sampling and Refine-in-loop strategy.

DDPM Pretrain on lateral Our method differs from existing approaches that directly train super-resolution models. Instead, we pre-train a diffusion model to learn the distribution of high-resolution images at lateral, avoiding being limited to a specific axial resolution. Diffusion models [10,19] employ a Markov Chain diffusion process to transform a clean image x_0 into a series of progressively noisier images during the forward process. This process can be simplified as:

$$q(x_t|x_0) = N(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I), \quad (1)$$

where $\bar{\alpha}_t$ controls the scale of noises. During inference, the model ϵ_θ predicts x_{t-1} from x_t . A U-Net ϵ_θ is trained for denoising process p_θ , which gradually reverses the diffusion process. This denoising process can be represented as:

$$p_\theta(x_{t-1}|x_t) = N(x_{t-1}; \epsilon_\theta(x_t, t), \sigma_t^2 I), \quad (2)$$

During training, we use 2D lateral slices, which is natural high-resolution to optimize ϵ_θ by mean-matching the noisy image obtained in Eq. 1 using the MSE loss [10]. Only HR slices at lateral plane X_{lat} were used for training, so the training process is unsupervised and independent of the specific axial resolution. So that ϵ_θ learns the universal structural distribution of biological tissues and can generate realistic HR images following $p_\theta(X_{lat})$.

Sparse Spatial Condition Sampling on axial We propose Sparse Spatial Condition Sampling (SSCS) to condition the generation process of ϵ_θ and generate high-axial-resolution reconstruction results. SSCS substitutes every reverse-diffusion step Eq. 2. We first transform the input axial LR slice x_{axi} to match the

Algorithm 1: Isotropic reconstruction using basic DiffuseIR

Input: axial slice x_{axi} , anisotropic factor α , refine-in-loop counts K

- 1 $x_0^{con}, M \leftarrow padding(x_{axi}, \alpha)$
- 2 **for** $t = T, \dots, 1$ **do**
- 3 $x_{t-1}^{con} \sim N(\sqrt{\bar{\alpha}_t} x_0^{con}, (1 - \bar{\alpha}_t)I)$
- 4 **for** $i = 1, \dots, K$ **do**
- 5 $x_{t-1}^* \sim N(x_{t-1}; \epsilon_\theta(x_t, t), \sigma_t^2 I)$
- 6 $x_{t-1} = M * x_{t-1}^{con} + (1 - M) * x_{t-1}^*$
- 7 **if** $t > 1$ **and** $i < K$ **then**
- 8 $x_t \sim N(\sqrt{1 - \beta_t} x_{t-1}, \beta_t I)$
- 9 **end**
- 10 **end**
- 11 **end**
- 12 **return** x_0

lateral resolution by intra-row padding: $(\alpha - 1)$ rows of zero pixels are inserted between every two rows of original pixels, where α is the anisotropic spatial factor. We denote M as the mask for original pixels in x_0^{con} , while $(1 - M)$ represents those empty pixels inserted. In this way, we obtain x_0^{con} , which reflects the sparse spatial content at axial, and further apply Eq. 1 to transform noise level:

$$x_{t-1}^{con} \sim N(\sqrt{\bar{\alpha}_t} x_0^{con}, (1 - \bar{\alpha}_t)I) \quad (3)$$

Then, SSCS sample x_{t-1} at any time step t , conditioned on x_{t-1}^{con} . The process can be described as follows:

$$x_{t-1} = M \odot x_{t-1}^{con} + (1 - M) \odot x_{t-1}^* \quad (4)$$

where x_{t-1}^* is obtained by sampling from the model ϵ_θ using Eq. 2, with x_t of the previous iteration. x_{t-1}^* and x_{t-1}^{con} are combined with M . By iterative denoising, we obtain the reconstruction result x_0 . It conforms to the distribution $p_\theta(X_{lat})$ learned by the pre-trained diffusion model and maintains semantic consistency with the input LR axial slice. Since SSCS is parameter-free and decoupled from the model training process, DiffuseIR can adapt to various anisotropic spatial resolutions by modifying the padding factor according to α while other methods require re-training. This makes DiffuseIR a more practical and versatile solution for isotropic reconstruction.

Refine-in-loop Strategy We can directly use SSCS to generate isotropic results, but the reconstruction quality is average. The diffusion model is capable of extracting context from the sparse spatial condition. Still, we have discovered a phenomenon of texture discoordination at the mask boundaries, which reduces the reconstruction quality. For a certain time step t , the content of x_{t-1}^* may be unrelated to x_{t-1}^{con} , resulting in disharmony in x_{t-1} generated by SSCS. During the denoising of the next time step $t - 1$, the model tries to repair the

disharmony of x_{t-1} to conform to p_θ distribution. Meanwhile, this process will introduce new inconsistency and cannot converge on its own. To overcome this problem, we propose the Refine-in-loop strategy: For x_{t-1} generated by SSCS at time step t , we apply noise to it again and obtain a new x_t and then repeat SSCS at time step t . Our discovery showed that this uncomplicated iterative refinement method addresses texture discoordination significantly and enhances semantic precision.

The total number of inference steps in DiffuseIR is given by $T_{total} = T \cdot K$. As T_{total} increases, it leads to a proportional increase in the computation time of our method. However, larger T_{total} means more computational cost. Recent works such as [24,15,16] have accelerated the sampling process of diffusion models by reducing T while maintaining quality. For DiffuseIR, adjusting the sampling strategy is straightforward. Lowering T and raising refinement iterations K improves outcomes with a fixed T_{total} . We introduce and follow the approach presented in DDIM [24] as an example and conducted detailed ablation experiments in Sec. 3 to verify this. Our experiments show that DiffuseIR can benefit from advances in the community and further reduce computational overhead in future work.

3 Experiments and Discussion

Dataset and implement details. To evaluate the effectiveness of our method, we conducted experiments on two widely used public EM datasets, FIB-25 [27] and Cremi [5]. FIB-25 contains isotropic drosophila medulla connectome data obtained with FIB-SEM. We partitioned it into subvolumes of $256 \times 256 \times 256$ as ground truth and followed [9] to perform average-pooling by factor $\alpha(x2, x4, x8)$ along the axis to obtain downsampled anisotropic data. Cremi consists of drosophila brain data with anisotropic spatial resolution. We followed [3] to generate LR images with a degradation network and conduct experiments on lateral slices. All resulting images were randomly divided into the training (70%), validation (15%) and test (15%) set. For the pre-training of the diffusion model, we follow [19] by using U-Net with multi-head attention and the same training hyperparameters. We use 256×256 resolution images with a batch size of 4 and train the model on $8 \times V100$ GPUs. For our sampling setting, we set $T, K = 25, 40$, which is a choice selected from the ablation experiments in Sec. 3 that balances performance and speed.

Quantitative and Visual Evaluation. To evaluate the effectiveness of our method, we compared DiffuseIR with SoTA methods and presented the quantitative results in Tab. 1. We use cubic interpolation as a basic comparison. 3DSRUNet [9] is a seminal isotropic reconstruction method based on deep learning, which requires high-resolution and low-resolution pairs as ground truth for supervised training. CycleGAN-IR [3] proposed an unsupervised approach using a CycleGAN [32] architecture, learning from unpaired axial and lateral slices. It is worth noting that these methods train specialized models based on a fixed

Table 1: **Quantitative evaluation of DiffuseIR against baselines.** PSNR \uparrow and SSIM \uparrow are used as evaluation metrics. We evaluated the FIB25 and Cremi datasets, considering three anisotropic spatial resolutions, $\alpha = 2, 4, 8$. Unlike other baselines which train a dedicated model for each α , our method only trains a single, generalizable model.

Method	FIB25			Cremi			
	x2	x4	x8	x2	x4	x8	
Interpolation	PSNR	33.21	30.29	29.19	31.44	29.34	28.27
	SSIM	0.854	0.722	0.538	0.782	0.574	0.451
†3DSRUNet	PSNR	33.84	32.31	30.97	32.04	31.12	30.28
	SSIM	0.877	0.824	0.741	0.820	0.761	0.719
CycleGAN-IR	PSNR	33.54	31.77	29.94	31.71	30.47	29.04
	SSIM	0.869	0.798	0.640	0.794	0.721	0.560
DiffuseIR (ours)	PSNR	33.81	32.37	31.09	31.97	31.24	30.24
	SSIM	0.881	0.832	0.774	0.819	0.783	0.726

† Supervised method.

anisotropic spatial setting. In addition, they need to be retrained when facing different anisotropic factors α . Various anisotropic factors α are shown in Tab. 1. Despite the model θ is trained solely for denoising tasks and having no exposure to axial slices during training, DiffuseIR outperforms unsupervised baselines and is even competitive with the supervised method [9]. As shown in Fig. 2, using our proposed refine-in-loop strategy, the results produced by DiffuseIR exhibit more significant visual similarity to the Ground Truth compared to other methods, which may be more prone to causing distortion and blurriness of some details. Notably, the versatility afforded by SSCS allows DiffuseIR to achieve excellent results using only one model, even under different isotropic resolution settings. This indicates that DiffuseIR overcomes the issue of generalization to some extent in practical scenarios, as users no longer need to retrain the model after modifying imaging settings.

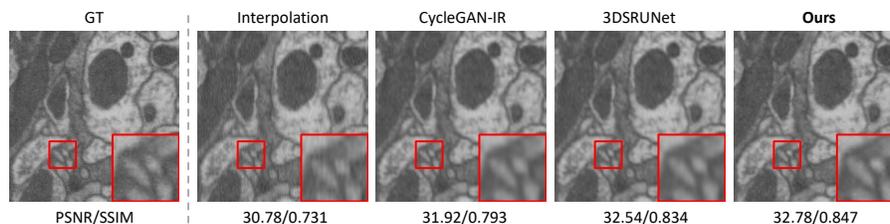


Fig. 2: **Visual comparisons on FIB-25 dataset ($\alpha = 4$).** DiffuseIR can generate competitive results compared to supervised methods, and the results appear more visually realistic.

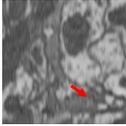
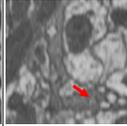
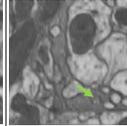
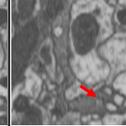
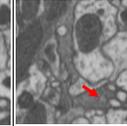
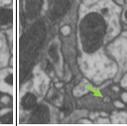
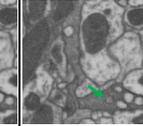
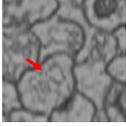
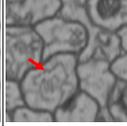
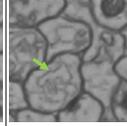
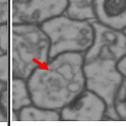
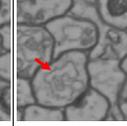
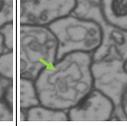
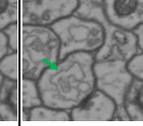
CycleGAN-IR	3DSRUnet	DiffuseIR(Ours)	CycleGAN-IR	3DSRUnet	DiffuseIR(Ours)	GT
FIB25, $\alpha=4 \rightarrow \alpha=8$			Cremi \rightarrow FIB25, $\alpha=4$			
						
30.21/0.677	30.95/0.738	32.10/0.822	32.04/0.821	32.71/0.838	33.17/0.856	
Cremi, $\alpha=4 \rightarrow \alpha=8$			FIB25 \rightarrow Cremi, $\alpha=4$			
						
29.68/0.633	30.66/0.734	31.17/0.768	30.56/0.717	31.34/0.792	31.78/0.820	
(a)			(b)			

Fig. 3: **Analysis on robustness.** (a) Test on unseen anisotropic factor α . (b) Test on different datasets with domain shifts (e.g., train on FIB25, test on Cremi). Our method is robust against various anisotropic factors and domain shifts between two datasets.

Further analysis on robustness. We examined the robustness of our model to variations in both Z-axis resolutions and domain shifts. Specifically, we investigated the following: **(a) Robustness to unseen anisotropic spatial factors.** The algorithm may encounter unseen anisotropic resolution due to the need for different imaging settings in practical applications. To assess the model’s robustness to unseen anisotropic factors, we evaluated the model trained with the anisotropic factor $\alpha = 4$. Then we do inference under the scenario of anisotropic factor $\alpha = 8$. For those methods with a fixed super-resolution factor, we use cubic interpolation to upsample the reconstructed result by 2x along the axis. **(b) Robustness to the domain shifts.** When encountering unseen data in the real world, domain shifts often exist, such as differences in biological structure features and physical resolution, which can impact the model’s performance [2,7]. To evaluate the model’s ability to handle those domain shifts, we trained our model on one dataset and tested it on another dataset.

Analysis: As shown in Fig. 3, DiffuseIR shows greater robustness than other methods. In scenario (a), other methods are trained on specific anisotropic factors for super-resolution of axial LR to lateral HR. This can result in model fragility during testing with unseen anisotropic resolutions. In contrast, DiffuseIR directly learns the universal structural distribution at lateral through generation task, applicable to various axial resolutions. All methods exhibit decreased performance in scenario (b). However, DiffuseIR shows a small performance degradation with the help of the multi-step generation of the diffusion model and sparse spatial constraints imposed by SSCS at each reverse-diffusion step. Further, compared to the previous methods predicting the result by one step, DiffuseIR makes the generating process more robust and controllable by adding constraints at each step to prevent the model from being off-limit.

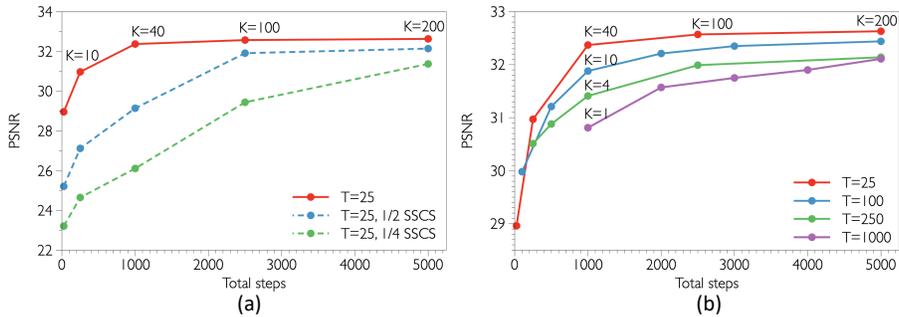


Fig. 4: **Ablation Study: (a)ablation on SSCS frequency.**Experimental results demonstrates the importance of SSCS. When reducing the frequency of SSCS usage, performance will severely decline. **(b)ablation on different refine-in-loop settings.** The results show that when the number of total steps is fixed, increase K will lead to higher PSNR.

Ablation Study. We conducted extensive ablation experiments Fig. 4. First, to demonstrate the effectiveness of SSCS, we use it only in partially alternate reverse-diffusion steps, such as $1/4$ or $1/2$ steps. As shown in Fig. 4 (a), increasing the frequency of SSCS significantly improves PSNR while bringing negligible additional computational costs. This indicates that SSCS have a vital effect on the model’s performance. Second, for the Refine-in-loop strategy, results show that keeping the total number of steps unchanged (reducing the number of time steps T while increasing the refine iterations K) can markedly improve performance. Fig. 4 (b) have the following settings: $T = \{25, 100, 250, 1000\}$ with $K\{40, 10, 4, 1\}$ to achieve a total of 5000 steps. The results show that the model performs best when $T = 25$ and PSNR gradually increases with the increase of K . A balanced choice is $\{T = 25, K = 40\}$, which improves PSNR by 1.56dB compared to $\{T = 1000, K = 1\}$ without using the Refine-in-loop strategy.

4 Conclusion

We introduce DiffuseIR, an unsupervised method for isotropic reconstruction based on diffusion models. To the best of our knowledge, We are the first to introduce diffusion models to solve this problem. Our approach employs Sparse Spatial Condition Sampling (SSCS) and a Refine-in-loop strategy to generate results robustly and efficiently that can handle unseen anisotropic resolutions. We evaluate DiffuseIR on EM data. Experiments results show our methods achieve SoTA methods and yield comparable performance to supervised methods. Additionally, our approach offers a novel perspective for addressing Isotropic Reconstruction problems and has impressive robustness and generalization abilities.

References

1. Chung, H., chul Ye, J.: Score-based diffusion models for accelerated mri (2023) [2](#)
2. Csurka, G.: Domain adaptation for visual applications: A comprehensive survey. arXiv: Computer Vision and Pattern Recognition (2017) [7](#)
3. Deng, S., Fu, X., Xiong, Z., Chen, C., Liu, D., Chen, X., Ling, Q., Wu, F.: Isotropic reconstruction of 3d em images with unsupervised degradation learning. medical image computing and computer assisted intervention (2020) [2](#), [5](#)
4. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Neural Information Processing Systems (2021) [2](#)
5. Funke, J., S.: cremi.org, <http://cremi.org/> [5](#)
6. González-Ruiz, V., García-Ortiz, J., Fernández-Fernández, M., Fernández, J.J.: Optical flow driven interpolation for isotropic fib-sem reconstructions. Computer Methods and Programs in Biomedicine (2022) [2](#)
7. Guan, H., Liu, M.: Domain adaptation for medical image analysis: A survey (2021) [7](#)
8. Hayworth, K.J., Xu, C.S., Lu, Z., Knott, G., Fetter, R.D., Tapia, J.C., Lichtman, J.W., Hess, H.F.: Ultrastructurally smooth thick partitioning and volume stitching for large-scale connectomics. Nature Methods (2015) [1](#)
9. Heinrich, L., Bogovic, J.A., Saalfeld, S.: Deep learning for isotropic super-resolution from non-isotropic 3d electron microscopy. medical image computing and computer assisted intervention (2017) [2](#), [5](#), [6](#)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Neural Information Processing Systems (2020) [3](#)
11. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models (2023) [2](#)
12. Kim, B., Chul, J.: Diffusion deformable model for 4d temporal medical image generation (2023) [2](#)
13. Li, X., Liu, J., Wang, S., Lyu, C., Lu, M., Chen, Y., Yao, A., Guo, Y., Zhang, S.: Efficient meta-tuning for content-aware neural video delivery (Jul 2022) [2](#)
14. Liu, J., Lu, M., Chen, K., Li, X., Wang, S., Wang, Z., Wu, E., Chen, Y., Zhang, C., Wu, M.: Overfitting the data: Compact neural video delivery via content-aware feature modulation. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV) (Oct 2021). <https://doi.org/10.1109/iccv48922.2021.00459>, <http://dx.doi.org/10.1109/iccv48922.2021.00459> [2](#)
15. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps (2022) [5](#)
16. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models (2022) [5](#)
17. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Gool, L.V.: Repaint: Inpainting using denoising diffusion probabilistic models (2023) [2](#)
18. Mikula, S.: Progress towards mammalian whole-brain cellular connectomics. Frontiers in Neuroanatomy (2016) [1](#)
19. Nichol, A., Dhariwal, P.: Improved denoising diffusion probabilistic models. arXiv: Learning (2021) [2](#), [3](#), [5](#)
20. Özbey, M., Dar, S.U., Bedel, H.A., Dalmaz, O., Şaban Öztürk, Güngör, A., Çukur, T.: Unsupervised medical image translation with adversarial diffusion models (2022) [2](#)
21. Park, H., Na, M., Kim, B., Park, S., Kim, K.H., Chang, S., Ye, J.C.: Deep learning enables reference-free isotropic super-resolution for volumetric fluorescence microscopy. Nature Communications (2021) [2](#)

22. Peng, C., Guo, P., Zhou, S.K., Patel, V., Chellappa, R.: Towards performant and reliable undersampled mr reconstruction via diffusion model sampling (2023) [2](#)
23. Schrödel, T., Prevedel, R., Aumayr, K., Zimmer, M., Vaziri, A.: Brain-wide 3d imaging of neuronal activity in caenorhabditis elegans with sculpted light. *Nature Methods* (2013) [1](#)
24. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv: Learning* (2020) [5](#)
25. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. *Cornell University - arXiv* (2021) [2](#)
26. Su, X., Song, J., Meng, C., Ermon, S.: Dual diffusion implicit bridges for image-to-image translation (2023) [2](#)
27. ya Takemura, S., Xu, C.S., Lu, Z., Rivlin, P.K., Parag, T., Olbris, D.J., Plaza, S.M., Zhao, T., Katz, W.T., Umayam, L., Weaver, C., Hess, H.F., Horne, J.A., Nunez-Iglesias, J., Aniceto, R., Chang, L.A., Lauchie, S., Nasca, A., Ogundeyi, O., Sigmund, C., Takemura, S., Tran, J., Langille, C., Lacheur, K.L., McLin, S., Shinomiya, A., Chklovskii, D.B., Meinertzhagen, I.A., Scheffer, L.K.: Synaptic circuits and their variations within different columns in the visual system of drosophila. *Proceedings of the National Academy of Sciences of the United States of America* (2015) [5](#)
28. Verveer, P.J., Swoger, J., Pampaloni, F., Greger, K., Marcello, M., Stelzer, E.H.K.: High-resolution three-dimensional imaging of large specimens with light sheet-based microscopy. *Nature Methods* (2007) [1](#)
29. Weigert, M., Royer, L., Jug, F., Myers, G.: Isotropic reconstruction of 3d fluorescence microscopy images using convolutional neural networks. *arXiv: Computer Vision and Pattern Recognition* (2017) [2](#)
30. Weigert, M., Schmidt, U., Boothe, T., Müller, A., Dibrov, A., Jain, A., Wilhelm, B., Schmidt, D., Broaddus, C., Culley, S., Rocha-Martins, M., Segovia-Miranda, F., Norden, C., Henriques, R., Zerial, M., Solimena, M., Rink, J.C., Tomancak, P., Royer, L., Jug, F., Myers, E.W.: Content-aware image restoration: Pushing the limits of fluorescence microscopy. *bioRxiv* (2018) [2](#)
31. Wu, Y., Rivenson, Y., Wang, H., Luo, Y., Ben-David, E., Bentolila, L.A., Pritz, C., Ozcan, A.: Three-dimensional virtual refocusing of fluorescence microscopy images using deep learning. *Nature Methods* (2019) [1](#)
32. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. *International Conference on Computer Vision* (2017) [2](#), [5](#)