



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Reference-free Axial Super-resolution of 3D Microscopy Images using Implicit Neural Representation with a 2D Diffusion Prior

Kyungryun Lee and Won-Ki Jeong\*

Korea University, College of Informatics,  
Department of Computer Science and Engineering  
{krlee0000, wkjeong}@korea.ac.kr

**Abstract.** Analysis and visualization of 3D microscopy images pose challenges due to anisotropic axial resolution, demanding volumetric super-resolution along the axial direction. While training a learning-based 3D super-resolution model seems to be a straightforward solution, it requires ground truth isotropic volumes and suffers from the curse of dimensionality. Therefore, existing methods utilize 2D neural networks to reconstruct each axial slice, eventually piecing together the entire volume. However, reconstructing each slice in the pixel domain fails to give consistent reconstruction in all directions leading to misalignment artifacts. In this work, we present a reconstruction framework based on implicit neural representation (INR), which allows 3D coherency even when optimized by independent axial slices in a batch-wise manner. Our method optimizes a continuous volumetric representation from low-resolution axial slices, using a 2D diffusion prior trained on high-resolution lateral slices without requiring isotropic volumes. Through experiments on real and synthetic anisotropic microscopy images, we demonstrate that our method surpasses other state-of-the-art reconstruction methods. The source code is available on GitHub: <https://github.com/hvcl/INR-diffusion>.

**Keywords:** Implicit neural representation · Isotropic reconstruction · Diffusion models

## 1 Introduction

High-resolution 3D microscopy images are widely used in biomedical research to reveal micro- to nano-scale structures in tissue samples. However, their axial resolution is often coarser than their lateral resolution due to physical limitations such as tissue sectioning or light diffraction. This substantial difference in resolution, which can be up to an order of magnitude, results in significant visual artifacts when the samples are viewed from the side, making analysis and downstream tasks difficult. Therefore, there is a growing need for isotropic reconstruction of 3D microscopy images, also known as axial super-resolution (SR).

---

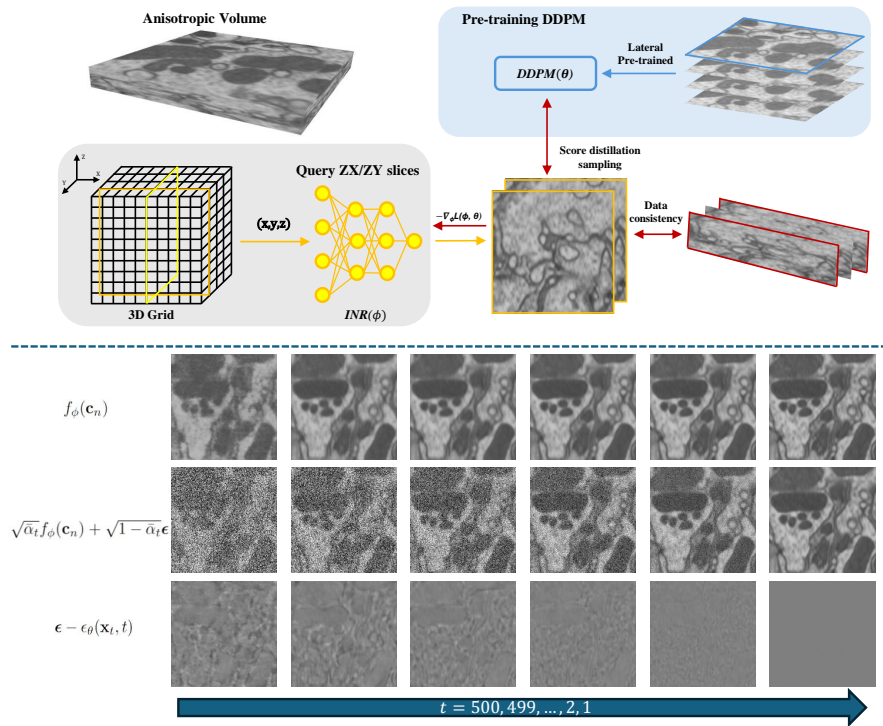
\* Corresponding author

Interpolation and deconvolution algorithms are traditional axial reconstruction methods that offer quick solutions without the need for a training process. However, they often produce suboptimal, blurry results, particularly for data samples rich in texture. More recent deep-learning methods can alleviate this issue, such as supervised training of a 3D neural network using anisotropic/isotropic paired 3D training data [4]. However, training a 3D neural network faces two major challenges: the lack of 3D isotropic training data and the high cost of training a 3D neural network. Since directly using a 3D neural network is challenging, recent studies utilize 2D neural networks and generative models. Weigert et al. [20, 21] proposed training a 2D U-Net[15] with the lateral images of the target 3D volume and then reconstructing the axial slices individually. Deng et al. [1] utilized Cycle-GAN to learn the degradation process between axial and lateral images, then applied this knowledge to train a 2D super-resolution model for reconstructing the axial slices. More recently, advanced generative models, such as diffusion models [5], are actively adopted to improve the axial SR quality, either by reconstructing the 3D volume in a slice-by-slice manner [7, 12] or approximating the 3D data prior with the product of two perpendicular 2D diffusion priors [8]. Nevertheless, as the reconstructed slices are not implicitly connected, approximating the volume with individual slices might fail to capture the 3D coherency of the volume and often cause misalignment artifacts.

To overcome the above limitations, we propose a novel axial SR method based on implicit neural representation(INR), which is a multi-layer perceptron that maps continuous spatial coordinates  $(x, y, z)$  to their intensity value. INR has shown great success in inverse rendering which allows reconstruction in a continuous volumetric representation from discrete measurements [11] and was later applied to diverse medical imaging reconstruction problems [16, 14, 10, 22]. Instead of using a 2D diffusion model to directly reconstruct a 3D isotropic volume in the pixel domain, our idea is to optimize the INR using a 2D diffusion model via score distillation sampling to implicitly learn 3D spatial coherency across the volume. To evaluate the effectiveness and practicality of our approach, we conducted extensive experiments on various datasets with high anisotropic axial scaling factors ( $\times 8$ ,  $\times 10$ ). Our experiments included FIB25 for simulation studies, CREMI for a real anisotropic electron microscopy image of *Drosophila* brain tissue, and a real two-channel fluorescent microscopy image of Zebrafish retina cells. We compared our method with other state-of-the-art techniques to show its reconstruction quality and effectiveness. The results demonstrate that our method produces results that are not only detailed and reliable but also accurate to the low-resolution data.

## 2 Method

The overview of our proposed model is illustrated in Fig. 1(top). As 3D anisotropic volume does not exist, we pre-train a 2D diffusion model with high-resolution lateral slices(XY). Once we have a diffusion model, we train an INR by querying axial slices(ZX or ZY) and fitting it to two loss functions: the measurement con-



**Fig. 1.** (top): The overview of our framework. First, we train the diffusion model using XY slices of the anisotropic volume, where the Z resolution significantly degrades compared to the XY resolution. During the optimization process of the INR, we randomly sample batches of ZX,ZY planes and optimize the loss. (bottom): The evolution of applying the diffusion prior: queried image(first row),  $t$ -step noisy image(second row) and the scaled difference between the original image and the refined output of the diffusion model. The discrepancy between the image before and after passing the diffusion model provides a guiding direction.

sistency loss and a diffusion prior loss. The measurement consistency loss ensures that the queried image matches the low-resolution measurement after passing the forward degradation process. By utilizing the score distillation sampling[13, 9], we incorporate the lateral image pre-trained diffusion model as an image prior for sharp and detailed reconstruction. A more detailed description of each component of our method is provided in the following sections.

## 2.1 Implicit neural representations for isotropic volume reconstruction

Implicit neural representations encode a spatial data field into the weights of a neural network; in other words, the model learns to map spatial coordinates of a pixel( $x, y, z$ ) to its data (e.g., intensity) values. The function is parameterized

by  $\phi$  and defined as  $f_\phi : \mathbf{c} \rightarrow f_\phi(\mathbf{c})$ , where the input  $\mathbf{c} \in \mathbb{R}^d$  is the  $d$  dimension spatial coordinates ( $d = 3$  for  $(x, y, z)$ ) and the output  $f_\phi(\mathbf{c}) \in \mathbb{R}$  is the scalar intensity value for the given the position. The input spatial coordinates are initially projected to a higher-dimensional space, and then processed through a multi-layer perceptron(MLP) architecture, returning an intensity value as an output. As the coordinates are implicitly connected via a neural network, training the MLP can be conducted using separate slice and coordinate pairs.

Consider reconstructing an isotropic volume in the pixel space where it can be approximated as reconstructing  $N$  sequential axial slices(ZX or ZY). Each of the  $N$  2D reconstruction problem is formulated as  $\mathbf{y}_n = \mathbf{A}\mathbf{x}_n + \epsilon$  where  $\mathbf{x}_n$  represents the  $n$ -th unseen isotropic slice,  $\mathbf{y}_n$  is the  $n$ -th downsampled slice(measurement) with additive noise  $\epsilon$  and  $\mathbf{A}$  is the Z direction degradation matrix shared along all the slices. One can reconstruct the  $n$ -th axial slice  $\mathbf{x}_n^*$  by solving the equation

$$\mathbf{x}_n^* = \operatorname{argmin}_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}_n\|_2^2 + R(\mathbf{x}) \quad (1)$$

where the first term ensures data fidelity by maintaining consistency between  $\mathbf{x}$  and  $\mathbf{y}$ , while the second term represents image regularization which incorporates prior knowledge of the image. Reconstructing the  $N$  slices independently and subsequently concatenating them together as a volume can lead to misalignments due to the highly ill-posed nature of each reconstruction problem. Instead, we optimize a single coordinate-based INR to fit the  $n$  equations. This approach enables the INR to implicitly learn 3D spatial coherency.

To train the MLP, we randomly query a  $n$ -th axial plane  $f_\phi(\mathbf{c}_n)$  by forwarding its spatial coordinates  $\mathbf{c}_n$  and minimize a loss that resembles Eq. 1

$$\mathcal{L}(\phi) = \mathbb{E}_n [\|\mathbf{A}f_\phi(\mathbf{c}_n) - \mathbf{y}_n\|_2^2 + \lambda R(f_\phi(\mathbf{c}_n))] \quad (2)$$

where  $\lambda$  balances the strength between data fidelity and the regularization loss. Often choices of  $R(\cdot)$  is the total-variation(TV) loss which gives smoother reconstruction. However, in this study, we apply a diffusion prior for the regularization term which allows realistic and detailed reconstruction. The formulation of this diffusion regularization will be discussed in the next section.

## 2.2 Diffusion models and score distillation sampling

Diffusion models first establish a  $T$ -step Markovian forward process, gradually perturbing an image with gaussian noise until it approaches to  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ [5, 2]. The transition of  $\mathbf{x}_0$  to  $\mathbf{x}_t$  can be expressed in a closed form of

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (3)$$

where  $\beta_t$ ,  $\alpha_t = 1 - \beta_t$  and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$  are the noise scheduling hyper parameters. During the reverse process, it generates an image by sampling from an approximated posterior distribution  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \approx q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ , which can be effectively achieved by training a noise predicting model parameterized by  $\theta$  that minimizes a (simplified unweighted)loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{\epsilon, t} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2], \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (4)$$

This formulation represents unconditional diffusion models.

We pre-train an unconditional diffusion model  $\theta$  with the lateral high-resolution slices to learn the high-resolution image prior. With Score Distillation Sampling(SDS)[13], we can optimize an INR such that the queried images conform to the pre-learned image prior of  $\theta$ . We first query a random axial slice  $\mathbf{x}_n = f_\phi(\mathbf{c}_n)$  and sample a Gaussian noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Following Eq. 3, we add the Gaussian noise to generate a  $t$ -step noisy image  $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x} + \sqrt{1 - \alpha_t}\epsilon$  (subscript  $n$  is omitted from here onwards). By forwarding  $\mathbf{x}_t$  through a frozen diffusion model and minimizing the loss function Eq. 4 with respect to  $\phi(\text{INR})$ ,

$$\nabla_\phi \mathcal{L}(\phi, \theta) = \mathbb{E}_{\epsilon, t} \left[ (\epsilon - \epsilon_\theta(\mathbf{x}_t, t)) \frac{\partial \hat{\epsilon}_\theta(\mathbf{x}_t, t)}{\partial \mathbf{x}_t} \frac{\partial \mathbf{x}_t}{\partial \phi} \right] \approx \mathbb{E}_{\epsilon, t} \left[ (\epsilon - \epsilon_\theta(\mathbf{x}_t, t)) \frac{\partial \mathbf{x}_t}{\partial \phi} \right] \quad (5)$$

we optimize the MLP so that it follows the diffusion prior. Note that the Jacobian of the diffusion model,  $\frac{\partial \hat{\epsilon}_\theta(\mathbf{x}_t, t)}{\partial \mathbf{x}_t}$ , is omitted which was proven to show comparable performance with much lower computational cost. Intuitively speaking, we are optimizing the INR to minimize the difference between the original queried image and the refined image after passing the diffusion model. By applying the SDS as a regularization in Eq. 2, we are able to train an INR representing an isotropic volume that accurately fits the measurements while preserving high levels of detail:

$$\mathcal{L}(\phi, \theta) = \mathbb{E}_{n, \epsilon, t} \left[ \|\mathbf{A}f_\phi(\mathbf{c}_n) - \mathbf{y}_n\|_2^2 + \lambda \text{stopgrad}(\epsilon - \epsilon_\theta(\mathbf{x}_t, t))^\top f_\phi(\mathbf{c}_n) \right] \quad (6)$$

where  $\mathbf{x}_t = \sqrt{\alpha_t}f_\phi(\mathbf{c}_n) + \sqrt{1 - \alpha_t}\epsilon$  is the  $t$ -level noisy image and **stopgrad** is used to represent the gradient Eq. 5 in the form of a loss function.

We linearly decrease the level of noise starting from 500(or 250) as the optimization progresses,  $t = 500, 499, \dots, 1$ . At the initial steps of the optimization, higher regularization of the diffusion model is applied, while towards the end, only fine details are tuned as shown in Fig. 1(bottom).

## 3 Experiments

### 3.1 Datasets and implementation details

We evaluate our framework by conducting a performance evaluation on three volumetric microscopy datasets. We use FIB-25 which is an isotropic FIB-SEM[18] dataset to conduct simulation studies (i.e., making synthetic downsampled data) for quantitative evaluation. For our real-world data applications, we use CREMI[3], which is a ssTEM data with an axial downsampling factor of 10, and a two-channel fluorescence microscopy image of a zebrafish retina sampled at a downsampling factor of 10.2 [6, 21]. As there is no isotropic ground-truth data available for those two datasets, we only conducted qualitative assessments on them. We use linear interpolation as our baseline method.

We compare our framework with state-of-the-art isotropic reconstruction methods which are all based on diffusion models. We use a same diffusion

**Table 1.** Quantitative Comparison with other methods. With a Gaussian filter ( $\sigma_z/\sigma_{xy} = 4/0.5$ ) and a downsampling rate of 8, the isotropic FIB25 volume is simulated to an anisotropic resolution of  $32 \times 256 \times 256$ . The baseline is linear interpolation. The highest scores are highlighted in **bold** and the second is underlined.

Method	ZX			ZY			XY		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Baseline	26.38	0.544	0.535	26.38	0.541	0.539	26.51	0.585	0.398
Lee et al.[7]	28.62	0.665	0.423	28.63	0.664	0.416	28.78	0.688	<b>0.278</b>
DiffuseIR[12]	26.16	0.524	0.456	26.16	0.517	0.527	26.34	0.548	0.411
TPDM[8]	28.57	0.663	<b>0.357</b>	28.58	0.661	<b>0.356</b>	28.65	0.687	<u>0.296</u>
Ours	<b>28.78</b>	<b>0.667</b>	<u>0.410</u>	<b>28.78</b>	<u>0.664</u>	<u>0.406</u>	<b>28.88</b>	<b>0.688</b>	0.359

model for all the methods including ours. Before starting the reconstruction(or optimization) the shared diffusion model is pre-trained with only 2D lateral images of the target volume where there is no isotropic volume to use as a reference. The diffusion model is trained using the backbone and setting of [5] where  $T = 1000$  and the learning rate is 0.00002. For the methods[7, 12] that take a slice-by-slice approach we reconstruct ZX slices along the y axis.


Considering that our microscopy images contain lots of high-frequency details, we implement the INR with an MLP consisting of sine activation [17] layers and a Gaussian Fourier feature embedding at the beginning [19] to capture the fine details. The MLP comprises a fully connected layer with a width of 768 and a depth of 8, along with a 512-dimensional random Gaussian Fourier feature embedding sampled from a standard deviation of 16. For the CREMI and FIB-25 datasets, we set the strength of the diffusion regularizer as  $\lambda = 0.25$ , with a decreasing level of noise starting from  $t = 500$  to  $t = 1$ . The MLP is trained with a learning rate of 0.00001 for 500 epochs, allowing the model to observe the volume 500 times. The volume is divided into batches of 8 random axial slices, and the reconstruction direction (ZX or ZY) is alternated with each iteration.

The training protocol remains consistent for the two-channel fluorescence microscopy volume, except that the MLP returns two outputs for each channel and is trained for 200 epochs.

### 3.2 Simulation studies on FIB25

We randomly extracted a  $512 \times 512 \times 512$  sub-volume, convolved it with a 3D Gaussian filter stretched in the Z-direction with a standard deviation of 4, and then downsampled it by selecting every 8 lateral slices along the Z-direction to simulate a  $64 \times 512 \times 512$  anisotropic volume. The 64 lateral images are used to train the diffusion model, wherein they are randomly cropped to a size of  $256 \times 256$ . We reconstruct a subpart of the training volume of size  $32 \times 256 \times 256$  and evaluate ZX, ZY, XY planes with PSNR, Structural similarity index measure(SSIM), and LPIPS [23] to assess human perceptual quality.

Table 1 shows that our framework offers the highest PSNR/SSIM score. TPDM [8], which is a diffusion model based SOTA 3D reconstruction method




**Fig. 2.** Visual comparison of reconstructing the simulated FIB25 volume with a down-sampling factor of 8. The volume is viewed in the ZY direction. All methods except ours show disconnection in the intersection.

originally used for MRI Z-axis SR, shows the highest LPIPS score showing texture-rich reconstruction, however it falls behind in PSNR/SSIM. This indicates that our method produces more reliable results with less hallucination which is common in generative models. Fig.2 also demonstrates that our method yields more reasonable results for uncertain regions. This can be attributed to the implicit 3D connectivity of the INR. During the alternating optimization with ZX and ZY slices, information exchange occurs, complementing the uncertainty of each region. Moreover, the implicit 3D coherency prevents misalignments, unlike DiffuseIR, which suffers from severe artifacts due to separate reconstruction of ZX slices. By indirectly applying the diffusion model, our method is capable of capturing both high-frequency details and structural information while preserving data consistency.

### 3.3 Real world anisotropic volumes

We reconstruct two real anisotropic microscopy images to demonstrate the practical usage of our method. Both datasets have dimensions of  $26 \times 256 \times 256$ , and our goal is to restore them to dimensions of  $256 \times 256 \times 256$ . Since we do not have prior knowledge of the degradation process, we assume linear downsampling for all methods.

In the CREMI data, severe downsampling rates lead to disconnections between horizontal structures. While other methods fail to restore these disconnections, our method successfully addresses this issue. As shown in Fig.3(row1), boundary disconnection is commonly shown in other methods, whereas in our case they are smoothly connected. Furthermore, while hallucinations may occur when the



**Fig. 3.** (row1): Visual comparison of reconstructing real anisotropic electron microscopy image with a downsampling factor of 10. Inside the red circle shows an example of hallucination. As there is no ground truth image, we use the lateral slices(XY) as a reference. (row2): A merged fluorescence microscopy image of a zebrafish retina cell composed of two channels: the red channel representing the nuclei and the green channel representing the nuclear envelope.

diffusion model directly generates the reconstructed result, our method utilizes the diffusion model only as a regularizer indirectly, thereby avoiding such cases.

The zebrafish retina image consists of two channels where the nuclei and nuclear envelope is separately imaged. In such situations, the INR shows strong adaptability because it can easily share information between different channels. By leveraging the correlation between the channels, it allows not only proper channel alignment but also mutual complementation. The second row of the figure proves that our approach effectively eliminates the disconnection artifacts on the nuclear envelope and accurately aligns it with the inner side.

## 4 Conclusion and Future work

We introduce a microscopy axial super-resolution framework that utilizes 2D diffusion models, eliminating the need for an isotropic 3D volume. By employing score distillation to enforce a 2D diffusion prior on an implicit neural representation, we can efficiently reconstruct the volume without the loss of 3D coherency. Through evaluation on three datasets, we demonstrate the superior reconstruction capabilities of our framework compared to other state-of-the-art diffusion model based methods. Results show that our method not only generates detailed reconstructions but also ensures reliability with reduced risk of hallucination.

While our method utilizes SDS, which does not require back-propagation through the diffusion model, optimizing the INR involves fitting the volume hundreds of times, resulting in a bottleneck in training time. Future research will focus on achieving faster optimization by adapting advanced techniques for both the diffusion model and the INR.



**Acknowledgements.** This work was partially supported by the National Research Foundation of Korea (RS-2024-00349697, NRF-2021R1A6A1A13044830), the Institute for Information Communications Technology Planning Evaluation (IITP-2024-2020-0-01819), the Technology development Program(RS-2024-00437796) funded by the Ministry of SMEs and Startups(MSS, Korea), and a Korea University Grant.

**Disclosure of Interests.** There are no conflicts of interest to declare.

## References

1. 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. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part V 23. pp. 163–173. Springer (2020)
2. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
3. Funke, J., Saalfeld, S., Bock, D., Turaga, S., Perlman, E.: Miccai challenge on circuit reconstruction from electron microscopy images, <https://cremi.org/>
4. Heinrich, L., Bogovic, J.A., Saalfeld, S.: Deep learning for isotropic super-resolution from non-isotropic 3d electron microscopy. In: Medical Image Computing and Computer-Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part II 20. pp. 135–143. Springer (2017)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
6. Icha, J., Kunath, C., Rocha-Martins, M., Norden, C.: Independent modes of ganglion cell translocation ensure correct lamination of the zebrafish retina. *Journal of Cell Biology* **215**(2), 259–275 (2016)
7. Lee, K., Jeong, W.K.: Reference-free isotropic 3d em reconstruction using diffusion models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention Workshops. pp. 235–245. Springer (2023)
8. Lee, S., Chung, H., Park, M., Park, J., Ryu, W.S., Ye, J.C.: Improving 3d imaging with pre-trained perpendicular 2d diffusion models. *arXiv preprint arXiv:2303.08440* (2023)
9. Mardani, M., Song, J., Kautz, J., Vahdat, A.: A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391* (2023)
10. McGinnis, J., Shit, S., Li, H.B., Sideri-Lampretsa, V., Graf, R., Dannecker, M., Pan, J., Stolt-Ansó, N., Mühlau, M., Kirschke, J.S., et al.: Single-subject multi-contrast mri super-resolution via implicit neural representations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 173–183. Springer (2023)
11. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
12. Pan, M., Gan, Y., Zhou, F., Liu, J., Zhang, Y., Wang, A., Zhang, S., Li, D.: Diffuseir: Diffusion models for isotropic reconstruction of 3d microscopic images.

- In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 323–332. Springer (2023)
13. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: The Eleventh International Conference on Learning Representations (2022)
  14. Reed, A.W., Kim, H., Anirudh, R., Mohan, K.A., Champley, K., Kang, J., Jayasuriya, S.: Dynamic ct reconstruction from limited views with implicit neural representations and parametric motion fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2258–2268 (2021)
  15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
  16. Shen, L., Pauly, J., Xing, L.: Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
  17. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
  18. Takemura, S.y., Xu, C.S., Lu, Z., Rivlin, P.K., Parag, T., Olbris, D.J., Plaza, S., Zhao, T., Katz, W.T., Umayam, L., et al.: Synaptic circuits and their variations within different columns in the visual system of drosophila. *Proceedings of the National Academy of Sciences* **112**(44), 13711–13716 (2015)
  19. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems* **33**, 7537–7547 (2020)
  20. Weigert, M., Royer, L., Jug, F., Myers, G.: Isotropic reconstruction of 3d fluorescence microscopy images using convolutional neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 126–134. Springer (2017)
  21. Weigert, M., Schmidt, U., Boothe, T., Müller, A., Dibrov, A., Jain, A., Wilhelm, B., Schmidt, D., Broaddus, C., Culley, S., et al.: Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods* **15**(12), 1090–1097 (2018)
  22. Zang, G., Idoughi, R., Li, R., Wonka, P., Heidrich, W.: Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1960–1970 (2021)
  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)