

niiv: Interactive Self-supervised Neural Implicit Isotropic Volume Reconstruction

Jakob Troidl¹, Yiqing Liang², Johanna Beyer¹, Mojtaba Tavakoli³, Johann Danzl³, Markus Hadwiger⁴, Hanspeter Pfister¹, and James Tompkin²

¹Harvard University, ²Brown University, ³IST Austria, ⁴KAUST

Abstract. Three-dimensional (3D) microscopy data is often anisotropic with significantly lower resolution (up to $8\times$) along the z axis than along the xy axes. Computationally generating plausible isotropic resolution from anisotropic imaging data would benefit the visual analysis of large-scale volumes. This paper proposes **niiv**, a self-supervised method for isotropic reconstruction of 3D microscopy data that can quickly produce images at arbitrary output resolutions. The representation embeds a learned latent code within a neural field that describes the implicit higher-resolution isotropic image region. We use an attention-guided latent interpolation approach, which allows flexible information exchange over a local latent neighborhood. Under isotropic volume assumptions, we self-supervise this representation on low-/high-resolution lateral image pairs to reconstruct an isotropic volume from low-resolution axial images. We evaluate our method on simulated and real anisotropic electron (EM) and light microscopy (LM) data. Compared to diffusion-based baselines, **niiv** shows improved reconstruction quality (+1 dB PSNR) and is over three orders of magnitude faster ($1,000\times$) to infer. Specifically, **niiv** reconstructs a 128^3 voxel volume in $2/10$ th of a second, renderable at varying (continuous) high resolutions for display. Our code is available at <https://github.com/jakobtroidl/niiv-miccai>.

Keywords: Volume Reconstruction · Neural Fields · Self-Supervised

1 Introduction

3D imaging data is ubiquitous in scientific domains such as biology or material sciences. However, many imaging modalities like 3D electron microscopy (EM) or light microscopy (LM) have limited axial (z) resolution due to physical sectioning of tissue slices or optical limitations. Thus, resolution is typically much higher in the lateral directions than in the axial direction (Fig. 1a).

Downstream tasks like interactive visual analysis would benefit from high-resolution isotropic volumes, but these can be extremely costly or impossible to obtain. Several computational methods attempt to generate isotropically resolved volumes [31,30,11,21,34,35,8,5]. Some approaches [31,30] require specific shape priors, such as exact point spread functions (PSFs), which are difficult to measure in practice. Machine learning approaches like *diffusion* can model

complex data distributions but require copious training data and are slow to infer [11,21]—requiring multiple minutes or even hours to reconstruct a small isotropic volume.

At the same time, anisotropic imaging volumes grow in size every year, containing terabytes [36,27] or even petabytes [24] of imaging data. Reconstructing isotropic volumes is typically an offline postprocess after image acquisition, but with increasing data sizes this becomes infeasible. For instance, Lee et al. [12] take up to three minutes to reconstruct a 128^3 voxel volume. Instead, isotropic volumes should be reconstructed on-demand and locally at interactive rates from anisotropic data for visual inspection. Faster neural implicit reconstruction approaches are either supervised [5] or rely on bilinear latent interpolation [34] with fixed averaging, thus limiting the field of view per queried latent.

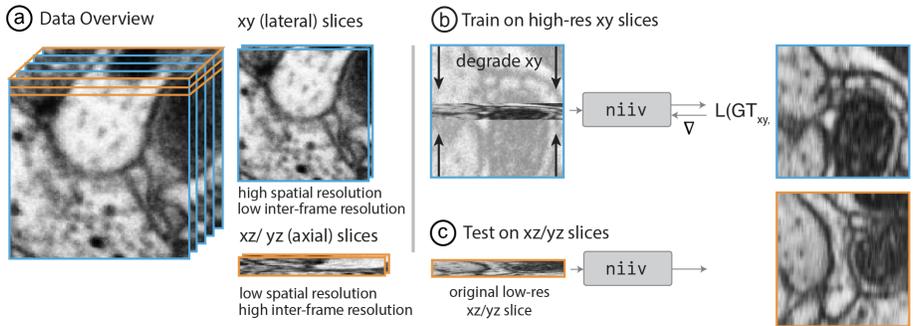


Fig. 1. Workflow. (a) Lateral (xy) slices in anisotropic volumes have a high spatial but low inter-frame resolution, while axial (xz/yz) slices have a low spatial but high inter-frame resolution. (b) Thus, `niiv` is trained to predict high-resolution lateral slices from artificially downsampled input. (c) `niiv` infers isotropic volumes at test time by predicting high-resolution axial (xz/yz) slices from their low-resolution counterparts.

Thus, we propose an interactive self-supervised method to reconstruct isotropic volumes from anisotropic data, which also enables learnt latent interpolation using a local attention mechanism. Building on recent advances in neural field representations [32,10,19,25], our model uses a super-resolution encoder [14] to relate a low-resolution axial slice to a plausible high-resolution image via a latent space [2,3]. A multi-layer perceptron (MLP) decodes a set of latent codes into a high-resolution axial slice sample, with local attention-guided latent interpolation creating an output image at any pixel resolution (Fig. 2). Both encoder and decoder are trained end-to-end on simulated anisotropic slices by downsampling isotropic lateral slices (Fig 1b). Inference is fast, requiring only $2/10$ th of a second to generate a 128^3 voxel volume, allowing interactive isotropic visual inspection.

For validation, we compare `niiv` against bilinear upsampling and three current self-supervised methods, including two neural-field [25,34], and one diffusion-based approach [12]. Our approach shows quality improvements over all and

quantitative improvements by +1 dB over the diffusion model. We demonstrate this through peak-signal-to-noise ratio (PSNR) computations in a sweep of a frequency-clipped Fourier domain, which offers a more robust metric than pixel-wise PSNR to noise in the training data. For computational efficiency, `niiv` is $1,000\times$ faster than the diffusion model, $666\times$ faster than the SIREN baseline considering volume-specific pertaining times and achieves similar inference speeds to Zhang et al. [34].

2 Related Work

Isotropic Volume Reconstruction. Recent self-supervised approaches [21] [12,15,13] train 2D diffusion models to learn the distribution of high-resolution lateral images. During reconstruction, they use low-resolution axial slices as priors for the backward diffusion process to predict missing volume information. While diffusion models achieve high-quality results, their usability is limited by compute-intensive training and time-consuming inference. In contrast, our approach improves reconstruction quality while inferencing three orders of magnitude faster than the diffusion baseline [12]. Recently, Zhang et al. [34] also use neural representations for anisotropic volume reconstruction with bilinear latent interpolation. Our approach additionally enables learning latent interpolation through a local attention mechanism. Deng et al. [4] learns a degradation model that generates realistic low-resolution, high-resolution training pairs. These pairs are used to train a reconstruction model like Iso-Net-2 [31]. Both models are trained independently, leading to complex training setups. On the other hand, supervised approaches [8,5] show high-quality results but are challenging to use in practice since isotropic volumes are required at training time. Other methods use video transformers [7], optical flow field interpolation [1], or standard ConvNets [31,35], which are limited to a specific output pixel resolution, whereas `niiv` can be decoded at any resolution.

Neural Implicit Super-Resolution. Encoding spatial information through implicit neural representations (INR) [20,32] has proven to be useful in areas, such as inverse graphics [19,25], shape representation [26,33,9], video encoding [10], and super resolution [2,3,18,34,35,5]. Our approach builds upon local implicit image functions (LIIF) [2], which allows the sampling of images at arbitrary resolution while retaining high-quality visual details. Critically, our approach differs from related approaches [35,34] by using attention guided latent refinement mechanism that enables learnt feature interpolation over a local latent neighborhood (Fig. 2).

3 Methodology

Problem Statement. Given a sampling of a volume V with isotropic xy axes and anisotropic z axis, we aim to learn a model g that reconstructs an isotropic z sampling of a volume \hat{V} purely by self-supervision. We assume that the volume sampling contains a physical medium whose distribution of material can

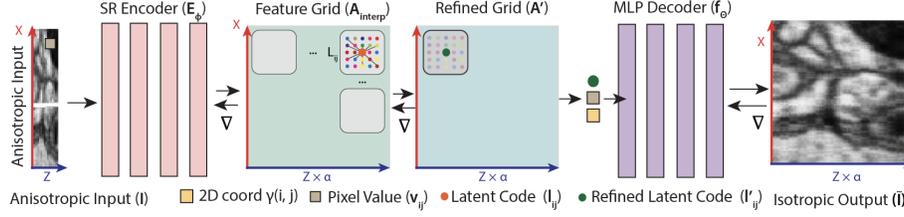


Fig. 2. Model overview. A super-resolution encoder E_ϕ [14] embeds the anisotropic input slice into latent codes placed within a 2D spatially-referenced grid A . A is bilinearly interpolated (A_{interp}) at output coordinates and subsequently refined using local attention (A'). We predict a pixel value in the output by querying A' at its coordinate (i, j) and concatenating the resulting latent code l'_{ij} (green) with the positionally encoded coordinate (i, j) (yellow) and the value at (i, j) in the input (beige). The decoding function f_θ , parametrized by an MLP, predicts the respective pixel value.

be effectively modeled by observing local statistics of the xy samplings. Given the low-resolution anisotropic slice $I_i^{x \times z}$ as input, g must reconstruct a plausible high-resolution isotropic xz slice $\tilde{I}_i^{x \times \alpha z}$, where i denotes a slice from the input volume sampling and α is the axial anisotropy factor (e.g., 8). Then, \tilde{V} is constructed by stacking predicted high-resolution slices $\tilde{V} = \{\tilde{I}_0, \dots, \tilde{I}_n\}$. The same approach applies equally to both xz and yz slices.

INR with Learned Latent Interpolation. Our hybrid neural field [32] uses latent codes within a 2D space to encode the high-resolution slice. First, a convolutional super-resolution encoder E_ϕ with parameters ϕ embeds a low-resolution axial slice $I \in \mathbb{R}^{x \times z}$ into a 2D referenced d -dimensional latent grid $A \in \mathbb{R}^{x \times z \times d}$, of latent codes l . Next, we bilinearly interpolate A at the desired output coordinates, such that $A_{\text{interp}} \in \mathbb{R}^{x \times (z\alpha) \times d}$. Next, we locally refine A_{interp} by attending each latent $l_{ij} \in A_{\text{interp}}$ to its local neighborhood L_{ij} (Fig. 2). We name refined latent codes l'_{ij} , and the refined latent grid A' .

$$L_{ij} = \{l_{i+m, j+n}\}_{m, n \in \{-k, \dots, 0, \dots, k\}, k \in \mathbb{N}} \quad (1)$$

$$l'_{ij} = \text{Attn}(l_{ij}, L_{ij}) = \sigma \left(\frac{l_{ij} W_Q (L_{ij} W_K)^T}{\sqrt{D}} \right) L_{ij} W_V \in \mathbb{R}^d \quad (2)$$

W_Q, W_K, W_V and D define the attention mechanism's trainable weight matrices and dimensionality, respectively [29]. σ represents the softmax function. Finally, a MLP decoder f_θ with parameters θ (Fig. 2) produces reconstructed image intensities \tilde{I} at an output pixel coordinate (x, y) :

$$\tilde{I}(i, j) = f_\theta ([l'_{ij}, v_{ij}, \gamma(i, j)]). \quad (3)$$

(x, y) is encoded using a 2-band frequency basis $\gamma(i, j)$, and v_{ij} is the pixel value obtained by bilinearly interpolating I at (i, j) . Thus, niiv can adapt flexibly to the requirements of interactive display across devices, unlike other ap-

proaches [8,31]. E_ϕ and f_θ are shared between volumes in the training and test dataset. We use the mean absolute error loss function (MAE) during training.

Simulating axial degradation. During training, we use artificially degraded xy images I_d^{xy} as model inputs and supervise outputs with the respective high-resolution xy slices I_{gt}^{xy} . We apply a function d that aims to simulate degradation along the z axis such that $d(I_{gt}^{xy}) = I_d^{xy}$. Here, we define d as an average pooling operator. In principle, other degradation models [4,12,6] can be applied based on the specific application and imaging domain.

4 Evaluating Reconstructions with Noisy Ground Truth

Evaluating niiv brings challenges, as only noisy ground truth exists, leading to uninformative PSNR values (Fig. 3a). Thus, a method that perfectly reconstructs lateral slices is overfitting to the noise; this is especially problematic in low-data regimes with powerful data-fitting models [12]. We wish to assess the reconstruction of biological structures despite the noise. Prior research has addressed this problem by downscaling the data to diminish noise [8] at the cost of sacrificing resolution. We propose calculating the PSNR in the Fourier domain where it is easier to separate high-frequency components of signals [6] such as noise (Fig. 3b), where we vary a cutoff frequency f_{cutoff} (Fig. 3c) across a range of values to observe the quality across frequencies. For example, if a method only achieves a greater PSNR than another at high f_{cutoff} but not at low f_{cutoff} , then it is likely overfitting the noise. Given Parseval’s theorem [22] and the unitary nature of the Fourier transform \mathcal{F} , we can directly compute the PSNR in the frequency domain, sidestepping the inverse transformation to the spatial domain.

$$PSNR(I, I_{\text{pred}}) = PSNR(\mathcal{F}I, \mathcal{F}I_{\text{pred}}). \quad (4)$$

We now incorporate the clipping operation in the Fourier domain, denoted by a low-pass operator L that discards frequencies above f_{cutoff} . Given the above relationship, we receive

$$PSNR(\mathcal{F}^{-1}L\mathcal{F}I, \mathcal{F}^{-1}L\mathcal{F}I_{\text{pred}}) = PSNR(L\mathcal{F}I, L\mathcal{F}I_{\text{pred}}). \quad (5)$$

That is, we can transform both images into the Fourier domain, apply the clipping operation via L , and then compute the PSNR directly in Fourier space. For more details, please refer to the supplement.

5 Experiments

5.1 Data and Implementation Details

We demonstrate the effectiveness of our approach on the publicly available FlyEM Hemibrain [23], FAFB [36] EM datasets and also ablate against LM approaches like LICONN [28] (Fig 5b). While the Hemibrain contains the central

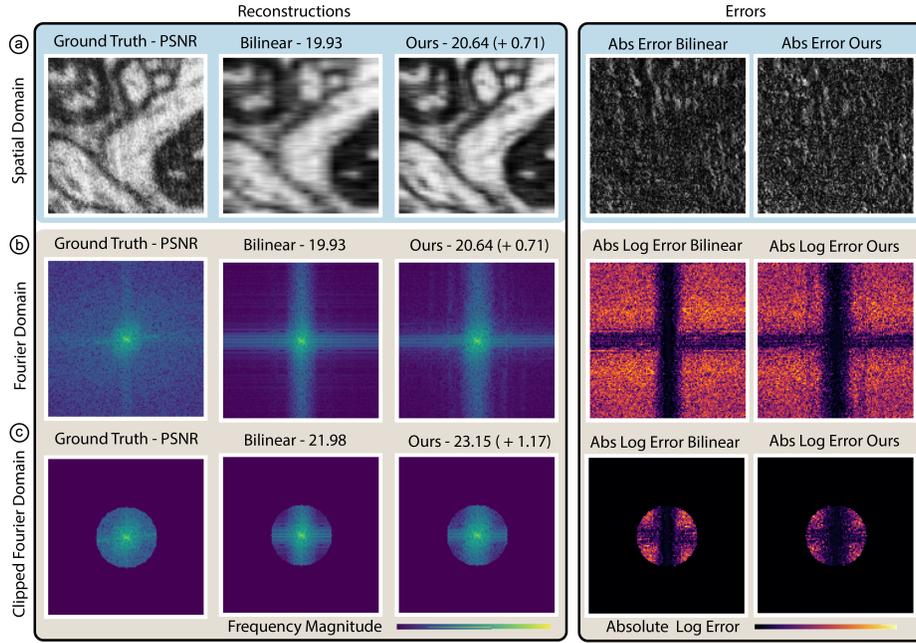


Fig. 3. Fourier space separates incorrect from correctly reconstructed frequencies. (a) In the spatial domain, reconstruction errors are randomly distributed for the baseline and our reconstruction. (b) Visualizing errors in the Fourier domain (right column) separates erroneous (bright) and correctly reconstructed frequencies (dark). (c) By cropping the Fourier domain’s high frequencies, the PSNR is not perturbed by frequencies corresponding to noise, yielding more informative PSNR values.

brain region of *Drosophila melanogaster* imaged at isotropic $8 \times 8 \times 8$ nm pixel resolution, we downsample the data to $8\times$ anisotropy along the z -axis through average pooling. FAFB shows the entire brain of a female adult fruit fly at naturally $5\times$ anisotropic $8 \times 8 \times 40$ nm pixel resolution. We randomly sample 400 subvolumes (128^3 pixels in the Hemibrain and 130^3 pixels in FAFB) and separate them into training ($N = 350$) and test datasets ($N = 50$). All metrics are reported on entire volumes rather than individual images. Our method is implemented in PyTorch, and all tests were performed on a single NVIDIA RTX 3090 Ti GPU. All experiments use the EDSR [14] super-resolution encoder without upsampling modules, six residual blocks, and 64-dimensional output features. The MLP is five layers deep, each 256 neurons wide. We train our model for 900 epochs using the Adam optimizer and a learning rate of 5×10^{-5} . We will make all code and data publicly available upon acceptance of this work.

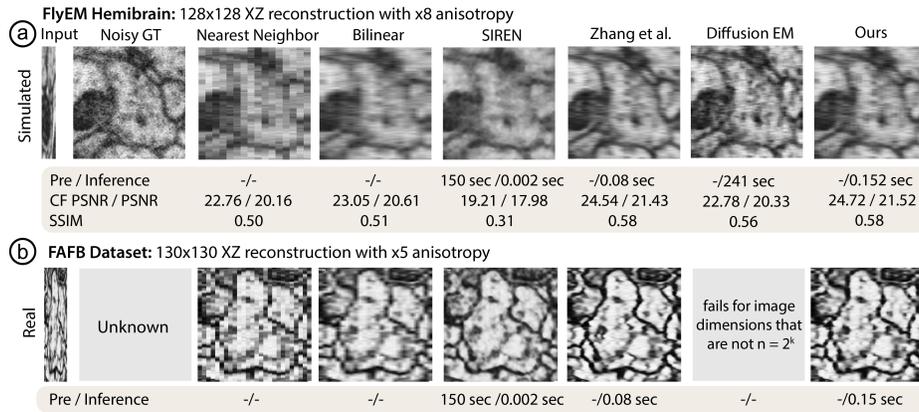


Fig. 4. Qualitative comparison of simulated (a) and real (b) anisotropic data. We achieve up to three orders of magnitude faster inference compared to diffusion baselines [12] and other neural implicit approaches like SIREN [25] while also reconstructing smaller structures with higher fidelity.

5.2 Qualitative and Quantitative Comparison

To showcase our method’s suitability for interactive reconstruction, we capped the GPU memory usage at 4 GB for all methods, reflecting a mid-tier laptop’s capacity. Within this constraint, our approach significantly outperforms the diffusion baseline, delivering inference speeds up to three orders of magnitude faster (0.225 vs. 264 seconds) for an anisotropy reconstruction task with anisotropy $\alpha = 8$ on 128^3 volumes (Hemibrain). The advantage is due to Diffusion-EM’s slow iterative inference process and the need to enforce frame-by-frame consistency for the probabilistic reconstruction process by conditioning each slice inference on a latent code retrieved from the previous slice, prohibiting batch processing. We also outperform SIREN [25] as it requires separate pretraining for each subvolume, leading to costly inference on unseen data (Table 1). Additionally, we achieve better qualitative results (CF PSNR, PSNR, SSIM) than Zhang

Table 1. Hemibrain Evaluation. We report clipped Fourier PSNR (CF PSNR) with $f_{cutoff} = 25$, regular PSNR, and SSIM. We differentiate between volume-specific pre-training time and inference time (Pre/Infer). The best scores are shown in bold.

| Hemibrain Data | CF PSNR \uparrow | PSNR \uparrow | SSIM \uparrow | Time (Pre/Infer) \downarrow | VRAM (GB) |
|-------------------|--------------------|-----------------|-----------------|-------------------------------|-----------|
| Nearest | 23.16 | 20.42 | 0.50 | - | - |
| Bilinear | 23.53 | 20.92 | 0.51 | - | - |
| SIREN [25] | 19.55 | 18.24 | 0.34 | (150s/0.003s) | 3.9 |
| Diffusion EM [12] | 23.24 | 20.65 | 0.56 | (-/264s) | 3.9 |
| Zhang et al. [34] | 24.86 | 21.63 | 0.57 | (-/0.113s) | 3.4 |
| Ours | 25.06 | 21.72 | 0.57 | (-/0.225s) | 3.9 |

et al. [34], while being slightly slower due to the additional computations required to compute attention-based latent interpolation. Comparing reconstruction quality, in contrast to the baselines, our model can reconstruct fine details (Fig. 4a) with sharp edges (Fig. 4b). While the diffusion results visually look sharp, small details are often reconstructed incorrectly, explaining the lower metric scores. Diffusion EM also fails for volume sizes, not in $\{2^i\}$ (Fig. 4b).

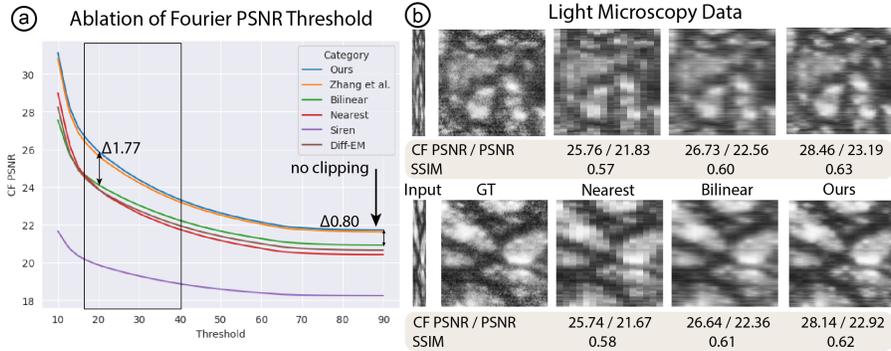


Fig. 5. Ablation studies. (a) While the unclipped PSNR in the Fourier domain shows no significant difference between our approach and the baselines ($\Delta 0.6$), the difference becomes more evident for thresholds $f_{cutoff} \in [14, 40]$ ($\Delta 1.63$). (b) We also test `niiv` on recent expansion light microscopy data [28] with simulated $8\times$ anisotropy. $f_{cutoff} = 25$.

5.3 Ablation Studies

Fourier PSNR. We tested the effect of the Fourier clipping threshold on the PSNR (Fig. 5a). If no clipping threshold is applied, the PSNR of our method and the baselines are low and close together due to the random image noise in the ground truth data. However, the black box (Fig. 5a) indicates a clipping window in Fourier space where image quality differences for the SIREN, Diff-EM, and Bilinear baselines are more accurately represented through the PSNR.

Data Modality. We use a recent, near isotropic voxel-size ($9.7 \times 9.7 \times 13$ nm) expansion LM dataset [28] of a mammalian hippocampus and simulate $8\times$ anisotropy using average pooling as a degradation model (Fig. 5b). Next, we train on 350 randomly sampled volumes and reconstruct 50 unseen 128^3 voxel test volumes at isotropic voxel size. Fig. 5b shows input and GT images and also compares our results with nearest- and bilinear interpolation. We find that `niiv` produces sharper images compared to bilinear interpolation also for LM data.

6 Conclusions and Future Work

Interactive isotropic rendering of anisotropic data is useful for large-scale data visual inspection tasks. Thus, we demonstrate that neural fields and encoder-

based superresolution representations are promising for fast and flexible self-supervised volume reconstruction. We propose three avenues for future work. First, integrating machine-learning elements like our approach into low-power Web-based image-rendering tools such as neuroglancer [16] or Viv [17] would rapidly deploy these advances. Second, future work should investigate whether latent image representations express higher-level, semantically interpretable morphological features, as these could be useful in downstream tasks such as tissue classification. Third, developing more accurate physics-based axial degradation models will help to improve the simulation of anisotropic training slices.

Acknowledgements. This work was supported by NIH grants 1U01NS132158 and R01HD104969. We thank the reviewers for their constructive feedback.

References

1. Carata, L., Shao, D., Hadwiger, M., Groeller, E.: Improving the visualization of electron-microscopy data through optical flow interpolation. In: Proceedings of the 27th spring conference on computer graphics. pp. 103–110 (2011)
2. Chen, Y., Liu, S., Wang, X.: Learning Continuous Image Representation with Local Implicit Image Function (2020)
3. Chen, Z., Chen, Y., Liu, J., Xu, X., Goel, V., Wang, Z., Shi, H., Wang, X.: VideoINR: Learning video implicit neural representation for continuous space-time super-resolution. CVPR (2022)
4. 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: MICCAI 2020. pp. 163–173. Springer (2020)
5. Fang, W., Tang, Y., Guo, H., Yuan, M., Mok, T.C.W., Yan, K., Yao, J., Chen, X., Liu, Z., Lu, L., Zhang, L., Xu, M.: CycleINR: Cycle Implicit Neural Representation for Arbitrary-Scale Volumetric Super-Resolution of Medical Data. In: CVPR. pp. 11631–11641 (Jun 2024)
6. Fritsche, M., Gu, S., Timofte, R.: Frequency separation for real-world super-resolution. In: ICCVW. pp. 3599–3608. IEEE (2019)
7. He, J., et al.: IsoVEM: Isotropic Reconstruction for Volume Electron Microscopy Based on Transformer. Tech. rep., Biophysics (2023)
8. Heinrich, L., et al.: Deep learning for isotropic super-resolution from non-isotropic 3D electron microscopy. In: MICCAI. pp. 135–143. Springer (2017)
9. Hendriks, T., Vilanova, A., Chamberland, M.: Implicit Neural Representation of Multi-shell Constrained Spherical Deconvolution for Continuous Modeling of Diffusion MRI. *Imaging Neuroscience* (Feb 2025)
10. Kim, S., Yu, S., Lee, J., Shin, J.: Scalable neural video representations with learnable positional features. *Advances in Neural Information Processing Systems* **35**, 12718–12731 (2022)
11. Lee, J., Jin, K.H.: Local Texture Estimator for Implicit Representation Function (2021), publisher: arXiv Version Number: 6
12. Lee, K., Jeong, W.K.: Reference-Free Isotropic 3D EM Reconstruction using Diffusion Models (2023)
13. Lee, K., Jeong, W.K.: Reference-Free Axial Super-Resolution of 3D Microscopy Images Using Implicit Neural Representation with a 2D Diffusion Prior. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 593–602. Springer (2024)

14. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 136–144 (2017)
15. Lu, C., Chen, K., Qiu, H., Chen, X., Chen, G., Qi, X., Jiang, H.: EMDiffuse: a diffusion-based deep learning method augmenting ultrastructural imaging and volume electron microscopy. preprint, Cell Biology (Jul 2023)
16. Maitin-Shepard, J., Baden, A., Silversmith, W., Perlman, E., et al.: google/neuroglancer: (Oct 2021)
17. Manz, T., Gold, I., et, a.: Viv: multiscale visualization of high-resolution multiplexed bioimaging data on the web. *Nature Methods* **19**(5), 515–516 (2022)
18. McGinnis, J., Shit, S., Li, H.B., Sideri-Lampretsa, V., Graf, R., Dannecker, M., Pan, J., Stolt-Ansó, N., Mühlau, M., Kirschke, J.S., others: Single-subject multi-contrast MRI super-resolution via implicit neural representations. In: MICCAI. pp. 173–183. Springer (2023)
19. 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)
20. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit Neural Representation in Medical Imaging: A Comparative Survey. In: ICCVW. pp. 2381–2391 (Oct 2023)
21. 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: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) MICCAI, vol. 14229, pp. 323–332 (2023)
22. Parseval, M.A.: Mémoire sur les séries et sur l’intégration complète d’une équation aux différences partielles linéaires du second ordre, à coefficients constants. *Mém. prés. par divers savants, Acad. des Sciences, Paris*,(1) **1**, 638–648 (1806)
23. Scheffer, L.K., Xu, C.S., Januszewski, M., et. al: A connectome and analysis of the adult *Drosophila* central brain. *eLife* **9**, e57443 (Sep 2020)
24. Shapson-Coe, A., Januszewski, M., Berger, D.R., Pope, A., Wu, Y., others: A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science* **384** (2024)
25. 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)
26. Sitzmann, V., Rezkikov, S., Freeman, B., Tenenbaum, J., Durand, F.: Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering. In: *Advances in Neural Information Processing Systems*. vol. 34, pp. 19313–19325 (2021)
27. Takemura, S.y., Hayworth, K.J., Huang, G.B., others: A connectome of the male *drosophila* ventral nerve cord pp. 2023–06 (2023)
28. Tavakoli, M.R., Lyudchik, J., et al.: Light-microscopy based dense connectomic reconstruction of mammalian brain tissue. *Tech. rep., Neuroscience* (Mar 2024)
29. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, , Polosukhin, I.: Attention is all you need. *NeurIPS* **30** (2017)
30. Weigert, M., Müller, A., others: Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods* **15**(12), 1090–1097 (2018)
31. Weigert, M., Royer, L., Jug, F., Myers, G.: Isotropic reconstruction of 3D fluorescence microscopy images using convolutional neural networks. In: MICCAI. pp. 126–134. Springer (2017)

32. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *Computer Graphics Forum* (2022)
33. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume Rendering of Neural Implicit Surfaces. In: *Advances in Neural Information Processing Systems*. vol. 34, pp. 4805–4815. Curran Associates, Inc. (2021)
34. Zhang, H., Zhang, Y., Wu, Q., Wu, J., Zhen, Z., Shi, F., Yuan, J., Wei, H., Liu, C., Zhang, Y.: Self-supervised arbitrary scale super-resolution framework for anisotropic MRI. In: *ISBI*. pp. 1–5. IEEE (2023)
35. Zhao, C., Dewey, B.E., Pham, D.L., Calabresi, P.A., Reich, D.S., Prince, J.L.: SMORE: A Self-Supervised Anti-Aliasing and Super-Resolution Algorithm for MRI Using Deep Learning. *IEEE TMI* **40**(3), 805–817 (Mar 2021)
36. Zheng, Z., Lauritzen, J.S., Perlman, E., Robinson, C.G., Nichols, M., Milkie, D., Torrens, O., Price, J., Fisher, C.B., Sharifi, N., others: A complete electron microscopy volume of the brain of adult *Drosophila melanogaster*. *Cell* **174**(3), 730–743 (2018)