

# niiv: Fast Self-supervised Neural Implicit Isotropic Volume Reconstruction

Jakob Troidl<sup>1</sup>, Yiqing Liang<sup>2</sup>, Johanna Beyer<sup>1</sup>, Mojtaba Tavakoli<sup>3</sup>, Johann Danzl<sup>3</sup>, Markus Hadwiger<sup>4</sup>, Hanspeter Pfister<sup>1</sup>, and James Tompkin<sup>2</sup>

<sup>1</sup>Harvard University, <sup>2</sup>Brown University, <sup>3</sup>IST Austria, <sup>4</sup>KAUST

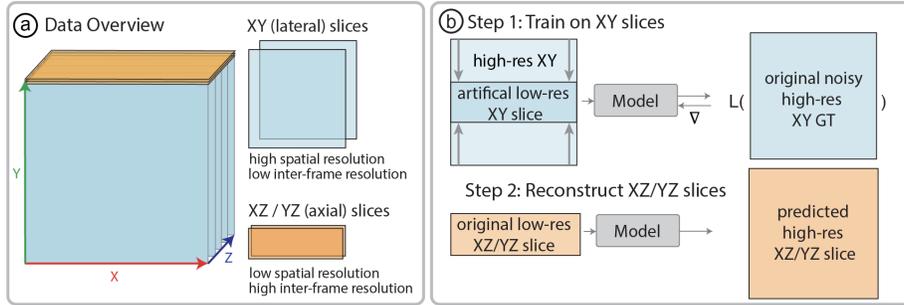
**Abstract.** Three-dimensional (3D) microscopy data often is anisotropic with significantly lower resolution (up to  $8\times$ ) along the  $z$  axis than along the  $xy$  axes. Computationally generating plausible isotropic resolution from anisotropic imaging data would benefit the visual analysis 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 (continuous) output resolutions. Within a neural field, the representation embeds a learned latent code that describes the implicit higher-resolution isotropic image region. 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 a state-of-the-art diffusion-based method, **niiv** shows improved reconstruction quality (+1 dB PSNR) and is over three orders of magnitude faster ( $2,000\times$ ) to infer. Specifically, **niiv** reconstructs a  $128^3$  voxel volume in  $1/10$ th of a second, renderable at varying (continuous) high resolutions for display.

**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 [29,30,11,20]. Some approaches [29,30] 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 [11,20]—requiring multiple minutes or even hours to reconstruct a small isotropic volume.



**Fig. 1. Training and reconstruction.** (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, our model is trained to predict high-resolution lateral slices from artificially downsampled input. Our model infers isotropic volumes at test time by predicting high-resolution axial ( $xz/yz$ ) slices from their low-resolution counterparts.

At the same time, anisotropic imaging volumes grow in size every year, containing terabytes [35,27] or even petabytes [23] 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.

To this end, we propose a self-supervised method to reconstruct isotropic volumes from anisotropic data. Building on recent advances in neural field representations [31,10,18,25], our model uses a super-resolution encoder [13] 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 bilinear 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  $1/10$ th of a second to generate a  $128 \times 128 \times 128$  volume. This allows interactive isotropic visual inspection.

For validation, we compare our approach against bilinear upsampling and two current self-supervised methods: a neural field reconstruction without the super-resolution encoder [25] and a diffusion-based model [12]. Our approach shows quality improvements over both and quantitative improvements by +1 dB over the diffusion model and +3 dB over the neural implicit baseline. 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  $2,000\times$  faster than the diffusion model and  $1,000\times$  faster than the SIREN baseline considering volume-specific pertaining times.

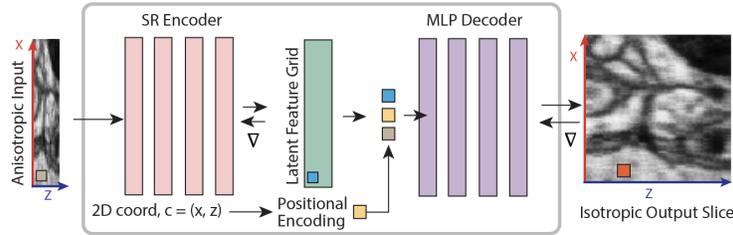
## 2 Related Work

**Isotropic Volume Reconstruction.** Recent self-supervised approaches [20] [12,14] 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. [32] also use arbitrary scale superresolution neural representations for anisotropic volume reconstruction. However, they focus on MRI imaging requiring different degradation models. Additionally, they do not discuss model inference speeds and evaluation strategies in the presence of noisy ground truth data. Deng et al. [4] learn 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 [29]. Both models are trained independently, leading to complex training setups. On the other hand, supervised approaches [9] show high-quality results but are difficult to use in practice since isotropic volumes are required at training time. Other methods use video transformers [8], optical flow field interpolation [1], or standard Conv-Nets [29,34], which are often limited to a specific output pixel resolution, whereas our model can be decoded at any resolution.

**Neural Implicit Super-Resolution.** Encoding spatial information through neural networks [19,31] has proven to be useful in areas, such as inverse graphics [18,25], shape representation [26], video encoding [10], and super resolution [2,3] [17]. 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. Chen et al. [3] extend LIIF for video super-resolution. However, their supervised approach requires high-resolution and high-frame rate training videos and can thus not be directly transferred to self-supervised isotropic reconstruction. Finally, Kim et al. [10] propose a neural video encoding with super-resolution capabilities. However, they optimize a single model per video, making their approach unsuitable for fast reconstruction of unseen volumes.

## 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  $\tilde{V}$  purely by self-supervision. We assume that the volume sampling contains a physical medium whose distribution of material can 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  $\alpha$  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.



**Fig. 2. Model overview.** Inspired by Chen et al. [2,3], a super-resolution encoder [13] first embeds the anisotropic input slice into latent codes that are placed within a 2D spatially-referenced grid  $M$ . We predict a pixel value in the output by querying  $M$  at its coordinate  $c$  and concatenating the resulting latent code (blue) with the positionally encoded  $c$  (yellow) and the value at  $c$  in the input (beige). The decoding function  $f_\theta$ , parametrized by an MLP, predicts the respective pixel value (orange).

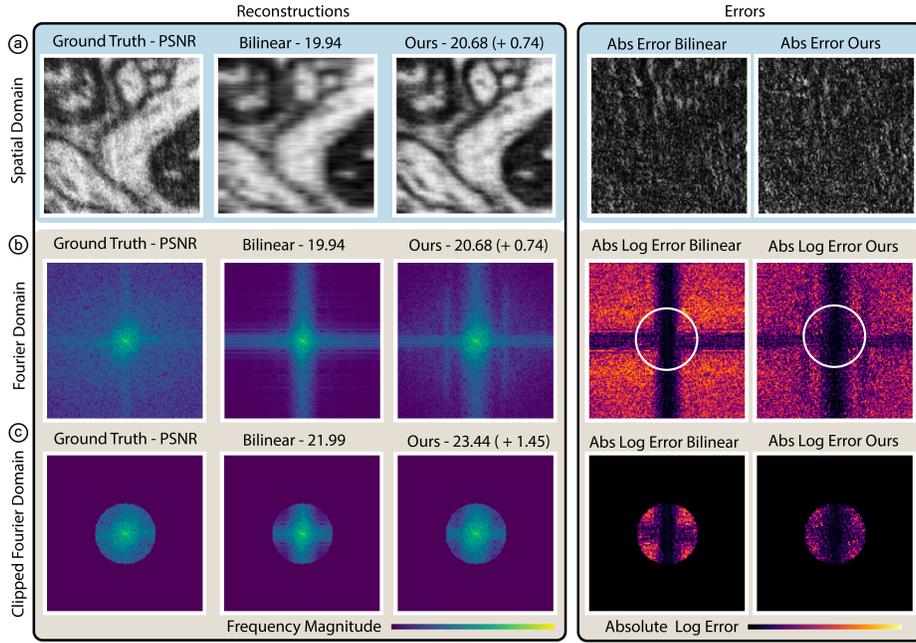
**Model Architecture.** Our hybrid neural field [31] uses latent codes within a 2D space to encode the high-resolution slice. First, a convolutional neural network super-resolution encoder  $E_\phi$  with parameters  $\phi$  embeds a low-resolution axial slice  $I \in \mathbb{R}^{x \times z}$  into a 2D referenced  $d$ -dimensional latent grid  $M \in \mathbb{R}^{x \times z \times d}$ , of latent codes  $l$ . Finally, a MLP decoder  $f_\theta$  with parameters  $\theta$  (Fig. 2) produces reconstructed image intensities at an output pixel coordinate  $c$ :

$$f_\theta([l_c, \gamma(c), v]) = \tilde{I}(c). \quad (1)$$

Here,  $c$  is encoded using a 2-band frequency basis  $\gamma(c)$ , and  $v$  is the pixel value obtained by simply bilinearly interpolating  $I$  at  $c$ . As  $M$  accepts continuous input, we can query  $M$  at any 2D coordinate  $c = (h_c, w_c)$  to retrieve a latent code  $l_c \in \mathbb{R}^d$ , making it possible to sample arbitrary pixel resolutions. Thus, `niiv` can adapt flexibly to the requirements of interactive display across devices, unlike other approaches [9,29]. Note that  $E_\phi$  and  $f_\theta$  are shared between all volumes in the respective training and test dataset.

**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 [4,12,7] can be applied based on the specific application and imaging domain.

**Model Training.** We minimize a loss function  $\mathcal{L}$  by end-to-end optimization of the encoder and decoder parameters  $\phi$  and  $\theta$ . Studies show that perceptual image quality metrics are superior to simple pixel-wise metrics for image super-resolution tasks [6]. Perceptual metrics [6,5,33] use features of deep pretrained models like VGG [24] as a basis for comparison. Yet, the *mean absolute error* (MAE) excels at image denoising tasks [6]. Thus, we combine the perceptual *Deep Image Structure and Texture Similarity* (DISTS) [5] metric and the MAE between predicted  $I_{pred}^{xy}$  and isotropic slices  $I_{gt}^{xy}$ . Our loss function  $\mathcal{L}$  is therefore



**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.

$$\mathcal{L} = w_1 \text{DISTS}(I_{\text{pred}}^{\text{xy}}, I_{\text{gt}}^{\text{xy}}) + w_2 \text{MAE}(I_{\text{pred}}^{\text{xy}}, I_{\text{gt}}^{\text{xy}}), \quad (2)$$

where  $w_i$  are the respective weights. We find  $w_1 = 1$  and  $w_2 = 30$  to consistently achieve good results.

## 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 [9] 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 [7] such as noise (Fig. 3b), where we vary a cutoff frequency  $f_{\text{cutoff}}$  (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 [21] and the unitary nature of the Fourier transform  $\mathcal{F}$ , we can directly compute the PSNR in the frequency domain (see derivation in supplement), sidestepping the inverse transformation to the spatial domain.

$$PSNR(I, I_{\text{pred}}) = PSNR(\mathcal{F}I, \mathcal{F}I_{\text{pred}}). \quad (3)$$

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

$$PSNR(\mathcal{F}^{-1}L\mathcal{F}I, \mathcal{F}^{-1}L\mathcal{F}I_{\text{pred}}) = PSNR(L\mathcal{F}I, L\mathcal{F}I_{\text{pred}}). \quad (4)$$

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.

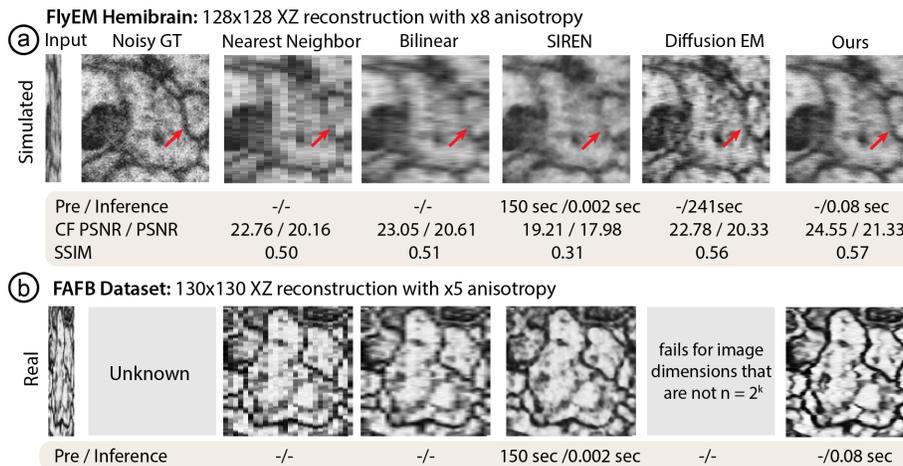
## 5 Experiments

### 5.1 Data and Implementation Details

We demonstrate the effectiveness of our approach on the publicly available FlyEM Hemibrain [22], FAFB [35] EM datasets and also ablate against LM approaches like LICONN [28] (Fig 5b). While the Hemibrain contains the central 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

**Table 1. Quantitative comparison on the Hemibrain data.** 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). We consistently outperform baselines for  $8 \times$  reconstruction of  $128^3$  volumes. The highest scores are highlighted in red, the second and third scores in orange and yellow, respectively.

<b>Hemibrain Data</b>	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
Ours	24.91	21.56	0.56	(-/0.113s)	3.4

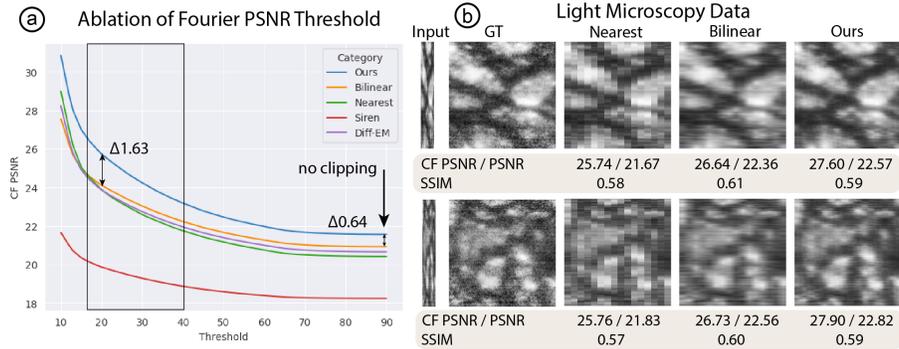


**Fig. 4. Qualitative comparison of simulated (a) and real (b) anisotropic data.** We achieve up to three orders of magnitude faster inference compared to diffusion baselines [12] and other neural implicit approaches like SIREN [25] while also reconstructing smaller structures with higher fidelity (see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

are reported on entire volumes rather than individual images. Our method is implemented in PyTorch, and all experiments were performed on a single NVIDIA RTX 3090 Ti GPU. All experiments use the EDSR [13] super-resolution encoder without upsampling modules, 16 residual blocks, and 64-dimensional output features. The MLP is 5 layers deep, each 256 neurons wide. We train our model for 1500 epochs using the Adam optimizer and a learning rate of  $5 \times 10^{-5}$ . Note that we train separate models for the Hemibrain, FAFB, and LICONN data.

## 5.2 Qualitative and Quantitative Comparison

To showcase our method’s suitability for interactive isotropic reconstruction, we capped the GPU memory usage at 4 GB for all methods, reflecting a mid-tier laptop’s typical capacity. Within this constraint, our approach significantly outperforms the diffusion baseline, delivering inference speeds up to three orders of magnitude faster (0.11 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 the neural implicit SIREN [25] baseline as it requires separate pretraining for each subvolume, leading to costly inference on unseen data (Table 1). Comparing reconstruction quality, in contrast to the baselines, our model can reconstruct fine details (Fig. 4a) with slice-by-slice consistency (see supplementary video) and sharper edges (Fig. 4b). While



**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.  $f_{cutoff} = 25$ .

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). Also, SIREN reconstruction results look blurry, making it unsuitable for isotropic reconstruction.

### 5.3 Ablation Studies

**Fourier PSNR Threshold.** We tested the effect of the Fourier clipping threshold on the PSNR (Fig. 5a). If no clipping threshold is applied, the PSNR values 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 are more accurately represented through the PSNR given noisy GT images.

**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. To this end, 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, developing more accurate physics-based axial degradation models (Sec. 3) to improve the simulation of anisotropic  $xy$  slices during training. Deng et al. [4] take a first step in that direction. Second, integrating machine-learning elements like our approach into low-power Web-based image-rendering tools such as neuroglancer [15] or Viv [16] would rapidly deploy these advances. Third, future work should investigate if latent image representations express higher-level semantically-interpretable morphological features, as these could be useful in downstream tasks like tissue classification (e.g., neuron typing).

## 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. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (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. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence* **44**(5), 2567–2581 (2020)
6. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Comparison of full-reference image quality models for optimization of image processing systems. *International Journal of Computer Vision* **129**, 1258–1281 (2021)
7. Fritsche, M., Gu, S., Timofte, R.: Frequency separation for real-world super-resolution. In: 2019 IEEE/CVF international conference on computer vision workshop (ICCVW). pp. 3599–3608. IEEE (2019)
8. He, J., Zhang, Y., Sun, W., Yang, G., Sun, F.: IsoVEM: Isotropic Reconstruction for Volume Electron Microscopy Based on Transformer. Tech. rep., Biophysics (Nov 2023)
9. Heinrich, L., Bogovic, J.A., Saalfeld, S.: Deep learning for isotropic super-resolution from non-isotropic 3D electron microscopy. In: MICCAI 2017. pp. 135–143. Springer (2017)
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. 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)

14. 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)
15. Maitin-Shepard, J., Baden, A., Silversmith, W., Perlman, E., et al.: google/neuroglancer: (Oct 2021)
16. Manz, T., Gold, I., Patterson, N.H., McCallum, C., Keller, M.S., Herr, B.W., Börner, K., Spraggins, J.M., Gehlenborg, N.: Viv: multiscale visualization of high-resolution multiplexed bioimaging data on the web. *Nature Methods* **19**(5), 515–516 (2022)
17. 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)
18. 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)
19. Molaie, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit Neural Representation in Medical Imaging: A Comparative Survey (2023)
20. 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)
21. 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)
22. 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)
23. Shapson-Coe, A., Januszewski, M., Berger, D.R., Pope, A., Wu, Y., Blakely, T., Schalek, R.L., Li, P.H., et al.: A connectomic study of a petascale fragment of human cerebral cortex. preprint, Neuroscience (May 2021)
24. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2014)
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., Rezhikov, 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., Januszewski, M., et al.: Light-microscopy based dense connectomic reconstruction of mammalian brain tissue. preprint, Neuroscience (Mar 2024)
29. 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)
30. Weigert, M., Schmidt, U., Boothe, T., Müller, A., others: Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods* **15**(12), 1090–1097 (2018)

31. 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)
32. 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: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5. IEEE (2023)
33. 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)
34. 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 Transactions on Medical Imaging* **40**(3), 805–817 (Mar 2021)
35. 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)