

Reference-Free Isotropic 3D EM Reconstruction using Diffusion Models

Kyungryun Lee and Won-Ki Jeong

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

Abstract. Electron microscopy (EM) images exhibit anisotropic axial resolution due to the characteristics inherent to the imaging modality, presenting challenges in analysis and downstream tasks. Recently proposed deep-learning-based isotropic reconstruction methods have addressed this issue; however, training the deep neural networks require either isotropic ground truth volumes, prior knowledge of the degradation process, or point spread function (PSF). Moreover, these methods struggle to generate realistic volumes when confronted with high scaling factors (e.g. $\times 8$, $\times 10$). In this paper, we propose a diffusion-model-based framework that overcomes the limitations of requiring reference data or prior knowledge about the degradation process. Our approach utilizes 2D diffusion models to consistently reconstruct 3D volumes and is well-suited for highly downsampled data. Extensive experiments conducted on two public datasets demonstrate the robustness and superiority of leveraging the generative prior compared to supervised learning methods. Additionally, we demonstrate our method’s feasibility for self-supervised reconstruction, which can restore a single anisotropic volume without any training data. The source code is available on GitHub: <https://github.com/hvcl/diffusion-em-recon>.

Keywords: Diffusion models · Isotropic EM reconstruction · Super-Resolution

1 Introductions

While 3D electron microscopy (EM) provide exceptional lateral resolution of 3 to 5 nanometers per pixel, the prevalent technique of 3D EM imaging involves physically sectioning tissue samples, resulting in a significantly lower axial resolution of approximately 30 to 50 nanometers per pixel (i.e., section thickness). This lower axial resolution poses challenges particularly for small structures such as synaptic clefts that can be smaller than the section thickness. [13] Conventional approaches, such as interpolation and deconvolution, have been used to address this issue, offering fast solutions. However, these methods often produce unsatisfactory results, particularly when dealing with texture-rich EM images.

In recent years, deep learning-based methods have emerged as promising approaches for the isotropic reconstruction of EM images, outperforming the

classical techniques. Heinrich et al. [8] leveraged isotropic FIB-SEM images to generate training data for supervised training of a 3D UNet-based super-resolution model. However, acquiring isotropic data in real-world scenarios is not feasible and there are cases where the downsampling process is unknown. On the other hand, by leveraging the point spread function (PSF), several studies [22, 23] proposed a framework for fluorescence microscopy that does not require isotropic training data. By convolving the PSF with laterally viewed high-resolution image and subsampling, they simulated the anisotropic axial images. Training a 2D-UNet-like architecture using the generated pairs, they achieved superior performance compared to conventional deconvolution algorithms [14]. An interesting aspect of such approaches is the potential for self-supervised learning, as the target data itself can be used for training. Building upon this work, Deng et al. [5] conducted further experiments under more realistic settings. They demonstrated that self-supervised training using an inaccurate PSF could yield poor results. To address the limitation, they adopted a cycle-GAN [26] framework to implicitly learn the degradation process and generate proper low-resolution versions for training. Nevertheless, these works rely on deterministic reconstruction models that aim to minimize the pixel-wise error; hence, when performed on high-scaling factors ($\times 8$, $\times 10$), the results are blurry and fail to preserve fine structures.

In this study, to tackle a challenging scenario where no training data is available, we propose a novel approach that leverages the denoising diffusion probabilistic model (DDPM) [9] for realistic 3D EM reconstruction. The diffusion model is recently gaining attention due to its high fidelity and diverse generation compared to other generative models [6, 17]. The diffusion model is adopted in various tasks for not only natural image domains [1, 15, 16] but also for medical image modalities [3, 4, 18, 24]. They are also well known for handling inverse problems [2, 10, 20]. In particular, these methods are capable of restoring 2D images without requiring task-specific training or datasets. Consequently, considering that our 3D reconstruction problem can be regarded as a super-resolution (SR) task, which inherently is an inverse problem, we leverage diffusion models to address it. However, in the context of 3D generation, it becomes necessary to employ a generative model that can capture the underlying 3D data. Training such a model is challenging not only due to the significant memory resources required but also acquiring isotropic 3D training data is not feasible. Hence, we adopt a slice-by-slice approach by utilizing 2D diffusion models for the reconstruction of the 3D volume. We train a 2D diffusion model to learn the data distribution from high-resolution lateral images. Subsequently, we leverage the diffusion prior of the laterally trained model in the sequential reconstruction of the low-resolution axial images. In order to maintain coherence across the 3D volume during the independent 2D reconstructions, we propose a sampling strategy where the previously reconstructed slice is encoded and used as a reference for the reconstruction of the next slice. We also introduce a heuristic method that improves the interpolated approximation to handle cases where the PSF is unknown, providing robustness in real-world scenarios. We validate the

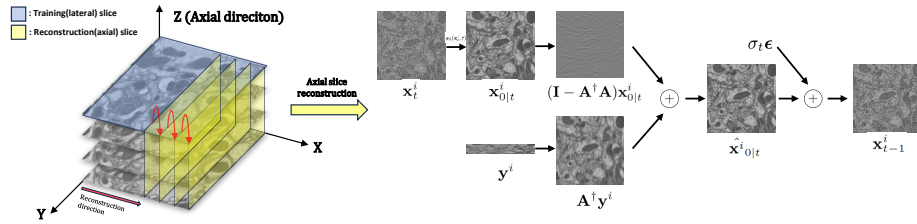


Fig. 1. (left): Reconstruction strategy for 3D EM via 2D diffusion models. The lateral images are used for training a diffusion model. Once trained, consistent sampling can be applied for any kind of 2D degradation(\mathbf{A}). (right): Intuitive illustration of the refinement process. The low-frequency components are replaced with $\mathbf{A}^T \mathbf{y}^i$ to fit the degradation process.

effectiveness and stability of the two proposed strategies via experiments and ablation studies. Our main contributions can be summarized as follows:

- We propose a sampling scheme that allows coherent 3D reconstruction using only 2D diffusion models. This allows smooth and continuous transitions between slices, therefore, eliminating artifacts when viewed in a perpendicular direction.
- We offer a simple but effective heuristic that can be applied without knowing the exact PSF which is often in practice. Moreover, the proposed approach is interpretable, allowing the reconstruction process to be more reliable.
- We demonstrate the superior reconstruction of DDPMs compared to previous auto-encoder-based methods by conducting simulation studies on a public dataset [19] for scenarios with/without prior information of the PSF. We also assess the performance on a real serial-section transmission electron microscopy (ssTEM) [7] volume without reference data or PSF information.

2 Method

As shown in Fig. 1(left), the whole process can be divided into two steps. We initially train a 2D DDPM on the lateral images of our target volume. This allows our generative model to learn "How high resolution images look like". Later, the laterally trained diffusion model is applied to reconstruct anisotropic axial planes slice-by-slice. Especially, we follow the restoring process of diffusion null-space model(DDNM) [20].

Preliminaries Diffusion models first define a T -step forward process that progressively perturbs an image to pure noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ [9]. By defining the noise schedule parameters β_t , $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, the forward process can be marginalized to a simple closed form of

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

The reverse process can be thought of as sampling from the posterior distribution, $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. Therefore, to estimate the true posterior with $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$, a noise predicting model is trained by minimizing the loss:

$$\mathcal{L}_{t-1} = \mathbb{E}_{\mathbf{x}_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1-\alpha_t}\epsilon, t)\|_2] \quad (2)$$

As described in DDIM [17], reparameterization allows inference in two steps by estimating the noise with ϵ_θ as follows:

$$\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \epsilon_\theta(\mathbf{x}_t, t)\sqrt{1-\bar{\alpha}_t}) \quad (3)$$

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\mathbf{x}_{0|t} + \sqrt{1-\bar{\alpha}_t-\sigma_t^2} \cdot \epsilon_\theta(\mathbf{x}_t, t) + \sigma_t\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (4)$$

with $\sigma_t = \sqrt{(1-\bar{\alpha}_{t-1})/(1-\bar{\alpha}_t)}\sqrt{1-\bar{\alpha}_t/\bar{\alpha}_{t-1}}$. Roughly speaking, at every iteration the reverse process is first estimating the clean image $\mathbf{x}_{0|t}$ at time t and again perturbing it with noise level $t-1$, gradually decreasing the noise until $t=0$.

DDNM [20] builds upon this inference process to solve linear inverse problems that are generally defined as $\mathbf{y} = \mathbf{A}\mathbf{x}$, where we aim to restore the data $\mathbf{x} \in \mathbb{R}^{n \times 1}$, given the degradation matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and its observation $\mathbf{y} \in \mathbb{R}^{m \times 1}$. To sample an image that fits the constrain given by \mathbf{A} , range-space replacement is added after Eq. 3 as follows:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})\mathbf{x}_{0|t} \quad (5)$$

where $\mathbf{A}^\dagger \in \mathbb{R}^{m \times n}$ is the pseudo-inverse of \mathbf{A} which can be calculated by the singular value decomposition (SVD) method. This refinement process ensures $\hat{\mathbf{x}}_{0|t}$ to satisfy the linear condition, hence shifting the direction of the reverse process to be consistent with the degradation. By substituting $\mathbf{x}_{0|t}$ with $\hat{\mathbf{x}}_{0|t}$ in Eq. 4, the noise mitigates the discrepancies between the replaced and original components of $\hat{\mathbf{x}}_{0|t}$. Fig. 1(right) illustrates the refinement process for isotropic reconstruction.

Diffusion Models for 3D EM Reconstruction As the degradation process occurs along the Z-axis of the 3D volume, it can be simplified as 2D degradations of contiguous ZY(or ZX) images. The 2D degradation can be represented with a matrix $\mathbf{A} = \mathbf{S}_f \mathbf{P}$, where $\mathbf{S}_f \in \mathbb{R}^{m \times n}$ is the sub-sampling operator choosing every f rows and $\mathbf{P} \in \mathbb{R}^{n \times n}$ is the PSF convolution operator. Assuming that we know the PSF and the downsampling factor, we can construct \mathbf{A} and its pseudo-inverse \mathbf{A}^\dagger , therefore we directly apply the DDNM sampling procedure to the i^{th} low-resolution ZY slice \mathbf{y}^i to reconstruct \mathbf{x}^i . However, it is important to regard that the diffusion model used in the 2D reconstruction does not take into account the continuity between neighboring slices.

Therefore, we propose a consistent sampling strategy where the previous slice is encoded by DDIM and used as a starting point for the subsequent

Algorithm 1 Reconstruction of the i^{th} slice

Require: $\mathbf{y}^i, \mathbf{x}_0^{i-1}, \mathbf{A}, \mathbf{A}^\dagger, \epsilon_\theta$

for $t = 0, \dots, R - 1$ **do** ▷ Encode \mathbf{x}_0^{i-1} deterministically

$$\mathbf{x}_{0|t}^{i-1} = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t^{i-1} - \epsilon_\theta(\mathbf{x}_t^{i-1}, t)\sqrt{1 - \bar{\alpha}_t})$$

$$\mathbf{x}_{t+1}^{i-1} = \sqrt{\bar{\alpha}_{t+1}}\mathbf{x}_{0|t}^{i-1} + \sqrt{1 - \bar{\alpha}_{t+1}} \cdot \epsilon_\theta(\mathbf{x}_t^{i-1}, t)$$

end for

$$\mathbf{x}_R^i = \mathbf{x}_R^{i-1}$$

for $t = R, \dots, 1$ **do** ▷ DDNM reconstruction starting from \mathbf{x}_R^{i-1}

$$\epsilon \sim \mathcal{N}(0, \mathbf{I})$$

$$\mathbf{x}_{0|t}^i = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t^i - \epsilon_\theta(\mathbf{x}_t^i, t)\sqrt{1 - \bar{\alpha}_t})$$

$$\hat{\mathbf{x}}_{0|t}^i = \mathbf{A}^\dagger \mathbf{y}^i + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})\mathbf{x}_{0|t}^i$$

$$\mathbf{x}_{t-1}^i = \sqrt{\bar{\alpha}_{t-1}}\hat{\mathbf{x}}_{0|t}^i + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \cdot \epsilon_\theta(\mathbf{x}_t^i, t) + \sigma_t \epsilon$$

end for

slice generation. This approach brings continuity between neighboring slices by leveraging the information encoded in the preceding slice. Moreover, most of the information overlaps between neighboring slices, therefore referencing the previous slice eases the reconstruction process for the diffusion model. Rather than beginning with pure Gaussian noise, we start the generation process of image \mathbf{x}^i by encoding the previously reconstructed images \mathbf{x}_0^{i-1} in a sequence of $[1, \dots, R]$ and use \mathbf{x}_R^{i-1} as a starting point. Specifically, by setting $\sigma_t = 0$ in Eq. 4 the DDIM iteration loses its stochasticity and it is possible to encode an image through a deterministic forward process. Given \mathbf{x}_R^{i-1} and \mathbf{y}^i , the i^{th} slice is reconstructed by the reverse process with $[R - 1, \dots, 0]$, but this time with the introduction of random noise. As there is no previous reference for the first slice, the reverse diffusion process starts from Gaussian noise ($R = 1000$). The overall process is described in Algorithm 1. We also observed that, although our method allows smooth transition along the sampling axis, the perpendicular planes do not directly leverage the diffusion prior, thus showing unrealistic visual results. Therefore, we ensemble the two reconstructions processed along the x-axis and y-axis. Additionally, due to the generative model’s stochastic nature, averaging the two results show a more steady and reliable generation.

In certain scenarios, the exact PSF is unknown. Therefore, we propose a simple approximation where we set \mathbf{A} as linear down-sampling and \mathbf{A}^\dagger as a linear interpolation operator. Despite the fact that linear interpolation is not the exact pseudo-inverse of linear down-sampling, [1] adopts it as a low-frequency guidance to generate an image in a desired direction. This approach enables the diffusion model to fill in the missing high-frequency details on top of the blurry interpolated observation $\mathbf{A}^\dagger \mathbf{y}$. As a result, the reconstructed data preserves the low-frequency structural information of the interpolated observation and remains interpretable without introducing abrupt changes. Although it has limitations that the reconstruction is a heuristic that relies on interpolation, we demonstrate through ablation that it gives better results compared to other assumptions of

A. We note that other kinds of interpolation methods can be used for \mathbf{A} and \mathbf{A}^\dagger instead of linear, e.g. cubic or lanczos.

3 Experiments

We assess the performance of our framework using two widely-used EM datasets: FIB-25 [19] and CREMI [7]. FIB-25 dataset is an isotropic FIB-SEM data commonly used for simulation studies, allowing quantitative evaluation of performance. CREMI is a ssTEM dataset with an anisotropic axial resolution. It serves as a real-world dataset for evaluating the performance of algorithms in handling anisotropic data. We trained the diffusion model with a U-Net backbone following [9] and adapted cosine scheduling [11] where $T = 1000$. The lateral training image size is 512×512 and the batch size is 4 with a learning rate of 0.00002. For sequential sampling, we reconstruct ZY images slice-by-slice along the x-axis, where the encoding/decoding level is $R = 200$. Except for the first axial slice, all slices were encoded by 4 steps and reconstructed(decoding) by 50 steps, where strided steps allow faster sampling. We compare our method with three auto-encoder-based approaches including 3D-SR-UNet [8], IsoNet [22, 23] and the framework proposed by Deng et al. [5]. All the methods were implemented in Pytorch [12] and tested on a single NVIDIA RTX A6000 GPU.

Evaluation on simulated data A randomly chosen subvolume of size $512 \times 512 \times 512$ from the isotropic FIB-25 data is convolved with a Gaussian filter and downsampled by choosing every $f \in \{4, 8\}$ lateral slices throughout the Z direction to generate synthetic anisotropic data. All deep-learning approaches except 3D-SR-UNet use only the target volume itself for training. 3D-SR-UNet is trained in a fully supervised manner with 26 additional isotropic subvolumes. Experiments are conducted for both cases where we know the PSF or not. For 3D-SR-UNet and IsoNet, we use average downsampling to generate low-resolution pair for blind-PSF. All methods require separate training for different PSF and downsampling factors, whereas our method does not require additional training. Performance is quantitatively evaluated by not only PSNR but also multi-scale structural similarity (MS-SSIM) [21] and LPIPS [25], which are more sensitive to fine patterns and structural details to demonstrate the realistic reconstruction achieved by our approach.

For $\sigma = 4$ and $f = 8$, Table 1 shows that our method outperforms all approaches for blind-PSF scenarios, especially in terms of LPIPS which measures the perceptual similarities between the reference and reconstructed images. In the case where we know the exact PSF, 3D-SR-UNet and IsoNet show better performance in PSNR and MS-SSIM. Nevertheless, the LPIPS score and visual results confirm that they cannot generate high-quality results. Moreover, 3D-SR-UNet is trained with isotropic volumes, thus incomparable. As discussed in Section 2, ‘Ours⁺’ shows that reconstruction along the ZX planes fails to generate realistic images viewed in ZY. Although the pixel-wise metrics preserve, the LPIPS score drops dramatically. ‘Ours⁺’ averages the two reconstructions

Table 1. Quantitative Comparison with other methods. With $f = 8$ and Gaussian filter of $\sigma = 4$, the isotropic FIB25 volume is simulated to an anisotropic resolution of $64 \times 512 \times 512$. The baseline is linear interpolation. \perp and $+$ indicate that the reconstruction was done on ZX slices or is ensembled, respectively. The highest scores are highlighted in **bold**.

PSF	Method	ZY			ZX			XY		
		PSNR \uparrow	MS-SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	MS-SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	MS-SSIM \uparrow	LPIPS \downarrow
	Baseline	26.12	0.842	0.567	26.11	0.840	0.555	26.18	0.848	0.379
Exact	3D-SR-UNet [8]	28.96	0.934	0.486	28.97	0.931	0.479	29.04	0.931	0.412
	IsoNet [22]	28.62	0.928	0.490	28.56	0.924	0.495	28.67	0.922	0.328
	Ours	27.92	0.914	0.375	27.93	0.916	0.434	28.03	0.913	0.296
	Ours $^+$	28.39	0.924	0.426	28.39	0.922	0.425	28.49	0.920	0.264
Blind	3D-SR-UNet	27.52	0.894	0.512	27.52	0.891	0.503	27.57	0.895	0.426
	IsoNet	27.60	0.897	0.503	27.29	0.888	0.515	27.35	0.889	0.363
	Deng <i>et al.</i> [5]	27.65	0.901	0.496	27.65	0.901	0.504	27.75	0.900	0.408
	Ours	27.55	0.901	0.391	27.55	0.903	0.448	27.64	0.901	0.325
	Ours $^+$	27.57	0.905	0.453	27.57	0.900	0.393	27.66	0.901	0.280
	Ours $^+$	27.95	0.911	0.431	27.95	0.909	0.433	28.04	0.908	0.284

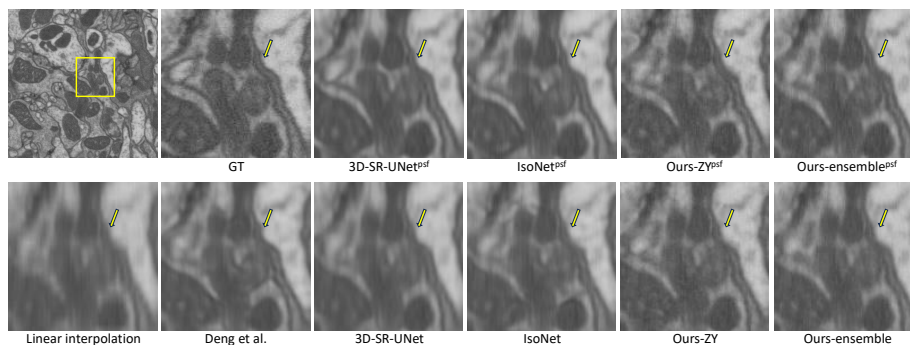


Fig. 2. Qualitative comparison of FIB25 reconstruction viewed in ZY. $f = 8$ and a Gaussian filter of $\sigma = 4$ was used. The superscript "psf" indicates that the exact point spread function was used for reconstruction.

along ZX and ZY, yielding a compromised result viewed in all directions and resulting in higher PSNR and MS-SSIM scores. Fig. 2 shows that the auto-encoder-based approaches tend to produce blurry results. This may be due to the limitation of the deterministic models based on pixel-wise loss functions, which may not fully capture the complexity and intricate details of the data. Furthermore, a noticeable quality gap exists between blind-PSF and exact-PSF cases for all methods except ours, which supports the robustness of our proposed heuristic. Experimental results for $\sigma = 2$ and $f = 4$ are described in Table S1 and Fig S1.

Visual comparison with real ssTEM data In this section, we present reconstruction results of real anisotropic ssTEM data (CREMI) from $52 \times 512 \times 512$ to $512 \times 512 \times 512$, which does not have PSF information nor isotropic reference data. Thus, 3D-SR-UNet cannot be trained. For IsoNet, we use average

downsampling to generate the low-resolution pairs for training. Our method uses the interpolation guidance method introduced in Section 2 for reconstruction. Fig. 3 shows the reconstructed ssTEM volumes viewed in ZY and XY. IsoNet and Deng et al.’s method show blurry results for ZY. Moreover, the XY images show severe artifacts indicating misalignment.

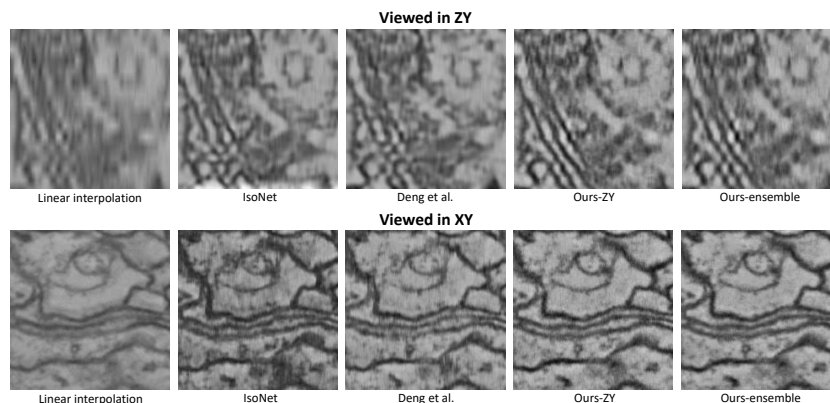


Fig. 3. Visual results of the reconstruction of a CREMI volume viewed in ZY and XY.

Ablation studies We perform two ablation studies on the FIB-25 dataset anisotropically simulated with $\sigma = 4$ and $f = 8$, starting with a comparison of different assumptions for the degradation process in blind-PSF scenarios. Imputation refers to the direct filling in of missing information, similar to the process of inpainting. We also compare it with average downsampling and an incorrect Gaussian filter of $\sigma = 2$ (note that $\sigma = 4$ is used to generate the synthetic anisotropic data). Table S2 and Fig S2 suggests that reconstructing on top of the interpolated approximation gives better results than imputing or estimating the degradation using an incorrect filter. Secondly, we investigate the importance of continuous sampling throughout the reconstruction process. As shown in Fig. 4, when the previous slice is not encoded as a reference, the reconstruction exhibits visible artifacts in XY and ZX views.

4 Conclusion

We present a diffusion-model-based approach for reference-free isotropic reconstruction on highly anisotropic 3D EM volumes. We introduced two additional strategies that allow 2D diffusion models for consistent 3D reconstruction where the PSF is unknown. Through quantitative and qualitative results, we demonstrated its superiority compared to SOTA reconstruction methods and showed

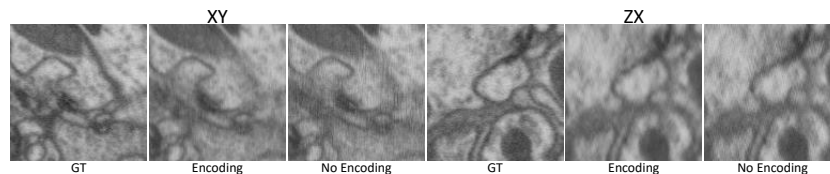


Fig. 4. Blind-PSF reconstructed volume viewed in XY and ZX. Reconstruction without referencing the previous slice shows severe artifacts due to the misalignment of adjacent slices.

the limitations of auto-encoder-based frameworks. In addition to generating high-quality data, it exhibits efficacy in challenging conditions where training data is scarce and prior information is minimal, thereby demonstrating its potential in real-world applications. In the future, we plan to investigate the impact of our methods on various downstream tasks in biomedical domains.

Acknowledgements

This work was partially supported by the National Research Foundation of Korea (NRF-2019M3E5D2A01063819, NRF-2021R1A6A1A 13044830), the Institute for Information & Communications Technology Planning & Evaluation (IITP-2023-2020-0-01819), the Korea Health Industry Development Institute (HI18C0316), the Korea Institute of Science and Technology (KIST) Institutional Program (2E32210 and 2E32211), and a Korea University Grant.

References

1. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14347–14356. IEEE (2021)
2. Chung, H., Kim, J., McCann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687 (2022)
3. Chung, H., Ryu, D., McCann, M.T., Klasky, M.L., Ye, J.C.: Solving 3d inverse problems using pre-trained 2d diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22542–22551 (2023)
4. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. *Medical Image Analysis* **80**, 102479 (2022)
5. 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)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
7. Funke, J., Saalfeld, S., Bock, D., Turaga, S., Perlman, E.: Miccai challenge on circuit reconstruction from electron microscopy images, <https://cremi.org/>

8. 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)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
10. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: *Advances in Neural Information Processing Systems*
11. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
12. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *NeurIPS* p. 8026–8037 (2019)
13. Plaza, S.M., Scheffer, L.K., Chklovskii, D.B.: Toward large-scale connectome reconstructions. *Current opinion in neurobiology* **25**, 201–210 (2014)
14. Richardson, W.H.: Bayesian-based iterative method of image restoration. *JoSA* **62**(1), 55–59 (1972)
15. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–10 (2022)
16. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
17. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: *International Conference on Learning Representations*
18. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. In: *International Conference on Learning Representations* (2022)
19. 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)
20. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490* (2022)
21. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*. vol. 2, pp. 1398–1402. Ieee (2003)
22. Weigert, M., Royer, L., Jug, F., Myers, G.: Isotropic reconstruction of 3d fluorescence microscopy images using convolutional neural networks. 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. 126–134. Springer (2017)
23. 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)
24. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VIII*. pp. 35–45. Springer (2022)

25. 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)
26. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: IEEE International Conference on Computer Vision (ICCV). pp. 2242–2251. IEEE (2017)