

# X-RAY2EM: UNCERTAINTY-AWARE CROSS-MODALITY IMAGE RECONSTRUCTION FROM X-RAY TO ELECTRON MICROSCOPY IN CONNECTOMICS

Yicong Li<sup>1,2,\*</sup>, Yaron Meirovitch<sup>3</sup>, Aaron T. Kuan<sup>4</sup>, Jasper S. Phelps<sup>4</sup>, Alexandra Pacureanu<sup>6</sup>  
Wei-Chung Allen Lee<sup>4,5</sup>, Nir Shavit<sup>1</sup>, Lu Mi<sup>1,\*</sup>

<sup>1</sup> Computer Science and Artificial Intelligence Laboratory, MIT

<sup>2</sup> John A. Paulson School of Engineering and Applied Sciences, Harvard University

<sup>3</sup> Department of Molecular and Cellular Biology, Harvard University

<sup>4</sup> Department of Neurobiology, Harvard Medical School

<sup>5</sup> F.M. Kirby Neurobiology Center, Boston Children's Hospital, Harvard Medical School

<sup>6</sup> ESRF, The European Synchrotron

## ABSTRACT

Comprehensive, synapse-resolution imaging of the brain will be crucial for understanding neuronal computations and function. In connectomics, this has been the sole purview of volume electron microscopy (EM), which entails an excruciatingly difficult process because it requires cutting tissue into many thin, fragile slices that then need to be imaged, aligned, and reconstructed. Unlike EM, hard X-ray imaging is compatible with thick tissues, eliminating the need for thin sectioning, and delivering fast acquisition, intrinsic alignment, and isotropic resolution. Unfortunately, current state-of-the-art X-ray microscopy provides much lower resolution, to the extent that segmenting membranes is very challenging. We propose an uncertainty-aware 3D reconstruction model that translates X-ray images to EM-like images with enhanced membrane segmentation quality, showing its potential for developing simpler, faster, and more accurate X-ray based connectomics pipelines.

**Index Terms**— Connectomics, reconstruction, segmentation, X-ray microscopy, electron microscopy.

## 1. INTRODUCTION

The field of connectomics investigates comprehensive maps of neuronal wiring and connectivity of the nervous system. Mapping the structure of neuronal networks is a key step toward understanding their function. Until recently, electron microscopy (EM) was the only approach that achieves comprehensive, nanometer-resolution imaging [1]. Unfortunately, obtaining three-dimensional (3D) EM volumes can be time and resource prohibitive. Because electrons are easily scattered, EM requires slicing massive numbers of thin sections

and imaging them at nanoscale. An emerging imaging technique, X-ray holographic nano-tomography (XNH), has recently shown its potential to extract neuronal wiring, especially long-range myelinated axons [2]. XNH can provide 3D imaging of whole specimens without nanoscale slicing during the imaging process, and thus significantly reduces the acquisition time and the potential existence of inter-slice alignment errors. Unfortunately, the speed and the scalability of XNH imaging on whole sample acquisition come, at least today, at the expense of delivering lower resolution images compared to EM. A large amount of time and resources could be saved if the reconstruction quality of XNH could be brought closer to that of EM.

Recent advances in deep learning based computer vision techniques [3] have brought forth successful techniques that can assist in the above task. In the field of image-to-image translation, two pioneering works, [4] and [5], develop 2D conditional generative adversarial networks (cGANs) [6] to translate images from one domain to the other. Such methods are further extended to 3D for processing medical data [7, 8, 9]. Several other works have demonstrated the potential to use such techniques to further optimize the imaging process [10, 11, 12].

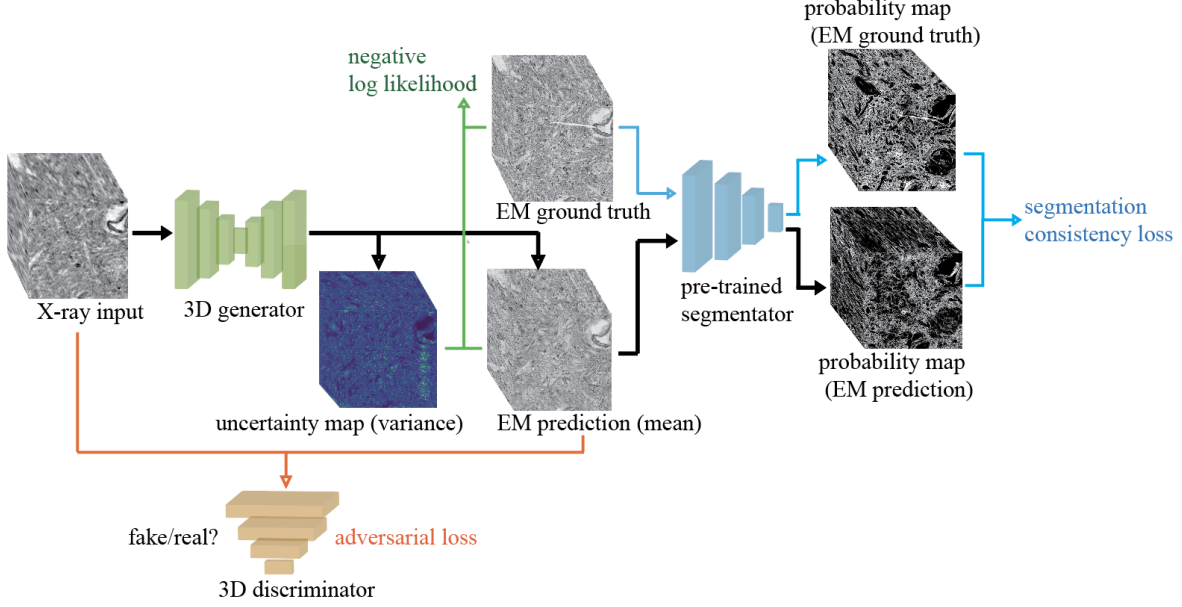
However, to the best of our knowledge, such novel deep learning techniques have not previously been used to bridge the image quality gap between the EM and X-ray modalities in connectomics. We aim to use cross-modality image reconstruction to improve the quality of the X-ray image by leveraging its 3D consistency to achieve enhanced X-ray segmentation quality. Achieving this goal could substantially reduce the time needed for imaging large-scale volumes of densely stained neuronal circuits.

In this work, we make a pioneering first step in this direction by showing that a 3D cGAN can be used to reconstruct high-quality EM-like images from low-quality X-ray images, mitigating the trade-off between imaging time and resolution

\* Corresponding authors: yicong.li@g.harvard.edu, lumi@mit.edu

† Work done during the internship at MIT CSAIL.

Supplementary material is available in this arXiv version.



**Fig. 1.** Overview of the proposed framework. The 3D generator  $G$  processes an image  $x$  from the X-ray image domain  $X$  to look similar to an image  $y$  from the EM image domain  $Y$  and outputs a corresponding uncertainty map. The 3D discriminator  $D_Y$  learns to distinguish between a real EM image  $y$  and a reconstructed output  $G(x)$ . The fixed pre-trained segmentation network  $F_s$  measures the distance between the predicted membrane probability map  $F_s(y)$  from the ground truth EM and the predicted membrane probability map  $F_s(G(x))$  from the EM-like reconstruction. The model is optimized with a weighted combination of an adversarial loss, a negative log-likelihood (NLL), and a segmentation-consistency loss.

in connectomics. Our method allows us to reconstruct the details of neuronal borders (membranes) that are evident at the EM level but hardly visible in the original X-ray images. Major contributions of this work are as follows:

- The first utilization of deep learning techniques to translate between the X-ray and EM modalities in connectomics.
- Our model demonstrates state-of-the-art reconstruction quality with 3D consistency and further improves the membrane segmentation performance.
- Our model also provides interpretable uncertainty maps during 3D reconstruction.

## 2. METHODS

### 2.1. Cross-modality image reconstruction framework

As shown in Fig. 1, we utilize a cGAN combined with segmentation constraints to perform this image reconstruction task. Given a target EM image  $y$  and an input X-ray image  $x$ , we build a 3D generator  $G$ , a 3D discriminator  $D_Y$ , and a fixed pre-trained segmentation network  $F_s$ . The 3D generator learns a mapping  $X \rightarrow Y$  to reconstruct EM-like images from X-ray images as well as generates corresponding uncertainty maps and the 3D discriminator learns to distinguish the reconstructed EM-like images from the ground truth EM im-

ages. Afterward, the fixed pre-trained segmentation network will output segmentation probability maps of the ground truth and the reconstruction, the distance between which is minimized in order to provide additional structural constraints.

### 2.2. Training objectives


#### 2.2.1. Adversarial loss

The adversarial loss is optimized as follows:

$$\min_G \max_{D_Y} \mathcal{L}_{\text{GAN}}(G, D_Y, X, Y) = \mathbb{E}_{y \sim p_{\text{data}}(y)} [\log D_Y(y)] + \mathbb{E}_{y \sim p_{\text{data}}(x)} [\log(1 - D_Y(G(x)))] \quad (1)$$

#### 2.2.2. Negative log-likelihood

The reconstruction model (generator) is optimized on aligned training pairs of X-ray and ground truth EM images  $(x_i, y_i)$ . Typically, we could have an additional  $L_1$  loss term which further constrains the learning. Instead, to provide robustness analysis and interpretability via uncertainty estimation, we use the negative log-likelihood (NLL) [13] over an output Gaussian distribution  $(G(x_i), \sigma_{y_i}^2)$  for every pixel  $y_i$ , where  $G(x_i)$  indicates the reconstructed image,  $\sigma_{y_i}^2$  indicates the measured uncertainty,  $N^2$  is the number of pixels, and we



**Fig. 2.** Reconstruction results: visualization of images reconstructed by different methods along X-Y, X-Z Y-Z directions, in comparison with input X-ray and ground truth EM.

have:

$$\mathcal{L}_{\text{NLL}}(G, X, Y) = \frac{1}{N^2} \sum_{i=1}^{N^2} \frac{\|y_i - G(x_i)\|_2^2}{2\sigma_{y_i}^2} + \frac{1}{2} \log(2\pi\sigma_{y_i}) \quad (2)$$

### 2.2.3. Segmentation-consistency loss

We utilize a pre-trained membrane segmentation network obtained from a related EM data domain [1]. It constrains the predicted probability maps from the ground truth EM image and the reconstructed EM-like image to be consistent in terms of membranes. The constraint is defined as:

$$\mathcal{L}_{\text{Seg}}(G, X, Y) = \|F_s(G(X)) - F_s(Y)\|_2^2 \quad (3)$$

### 2.2.4. Final objective

Our final objective is a weighted combination of the adversarial loss  $\mathcal{L}_{\text{GAN}}$ , the negative log-likelihood  $\mathcal{L}_{\text{NLL}}$ , and the segmentation-consistency loss  $\mathcal{L}_{\text{Seg}}$ :

$$\mathcal{L}(G, D_Y, X, Y) = w_{\text{GAN}} \mathcal{L}_{\text{GAN}}(G, D_Y, X, Y) + w_{\text{NLL}} \mathcal{L}_{\text{NLL}}(G, X, Y) + w_{\text{Seg}} \mathcal{L}_{\text{Seg}}(G, X, Y) \quad (4)$$

## 3. EXPERIMENTS AND RESULTS

### 3.1. Datasets

#### 3.1.1. X-Ray2EM dataset

For the EM modality, we have an EM dataset of mouse somatosensory cortex near the layer I/II boundary. The original

Model	Dir.	PSNR	SSIM	Dir.	PSNR	SSIM	Dir.	PSNR	SSIM
Full 2D		14.670	0.140		14.490	0.124		14.500	0.106
Hybrid 3D		14.701	0.143		14.566	0.129		14.702	0.108
Full 3D w/ $L_1$	X-Y	14.707	<b>0.153</b>	X-Z	14.702	0.133	Y-Z	14.734	0.115
Full 3D		<b>14.774</b>	<u>0.150</u>		<b>14.763</b>	<b>0.134</b>		<b>14.793</b>	0.114
Full 3D + Seg		14.738	0.149		14.722	0.132		14.745	<b>0.116</b>

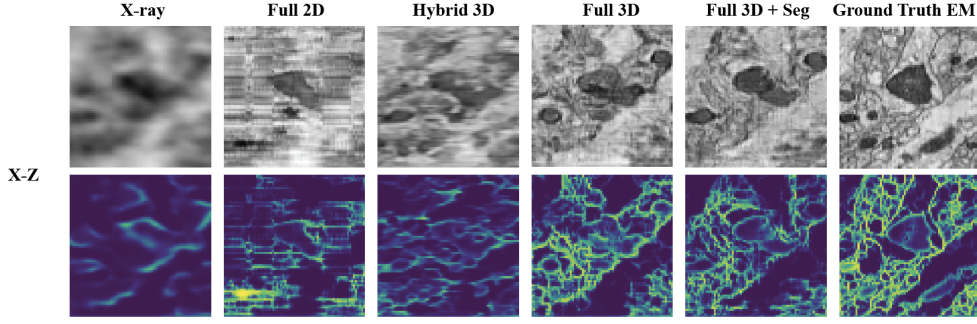
**Table 1.** Quantitative evaluation on reconstruction quality. PSNR: peak signal-to-noise ratio; SSIM: structural similarity index measure; Dir.: direction; Bold Number: best score; Underlined Number: second best score.

Model	Dir.	JS	Dice	Dir.	JS	Dice	Dir.	JS	Dice	Dir.	JS	Dice
X-ray		0.366	0.439		0.350	0.424		0.361	0.429		0.167	0.625
Full 2D		<b>0.413</b>	<b>0.546</b>		0.382	0.495		0.388	0.495		0.313	0.838
Hybrid 3D	X-Y	0.403	0.521	X-Z	0.369	0.466	Y-Z	0.377	0.470	3D	0.279	0.824
Full 3D		0.399	0.506		0.383	0.498		0.392	0.505		0.424	0.919
Full 3D + Seg		0.400	0.507		<b>0.386</b>	<b>0.502</b>		<b>0.393</b>	<b>0.506</b>		<b>0.433</b>	<b>0.922</b>

**Table 2.** Quantitative evaluation on membrane segmentation. JS: Jaccard Score; Dice: Dice score; Dir.: direction; Bold Number: best score; Underlined Number: second best score.

voxel size of EM data is  $4 \times 4 \times 45$  nm, while the volumes we used here have been downsampled by a factor of 16 in both X and Y directions to a voxel size of  $64 \times 64 \times 45$  nm. The overall volume size is  $2700 \times 2700 \times 243$  and is split into training/validation/test volumes with size of  $2700 \times 1620 \times 243$ ,  $2560 \times 512 \times 243$ , and  $2560 \times 512 \times 243$ , respectively.

For the X-ray modality, we have an XNH dataset with a voxel size of  $100 \times 100 \times 100$  nm. However, the true resolution is closer to  $200 \times 200 \times 200$  nm [2]. We warp this X-ray data into the EM volume space ( $64 \times 64 \times 45$  nm) using an elastic alignment algorithm [14, 15]. The overall volume size is also  $2700 \times 2700 \times 243$ . The training/validation/test splits are the same as the aforementioned EM dataset. It is important to note that there are larger resolution gaps and misalignment



**Fig. 3.** Segmentation results: visualization of 2D segmentation probability map of membranes in the X-Z direction.

along the X-Z and the Y-Z directions than the X-Y direction between X-ray and EM modalities, leading to a more challenging reconstruction task along these two directions.

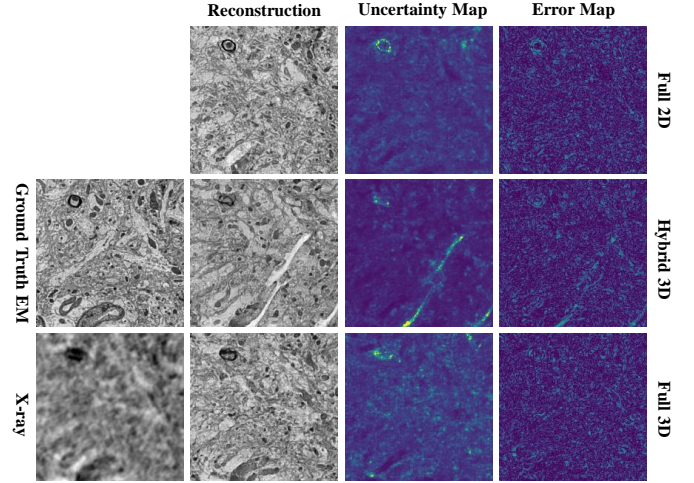
### 3.1.2. Public EM segmentation benchmark dataset

To pre-train the segmentation network, we use a related EM dataset of mouse cortex with corresponding membrane labels from Kasthuri. et al [1].

## 3.2. Performance evaluation

We evaluate our proposed framework in comparison with two baseline methods on both the reconstruction task and the downstream segmentation task: (1) The original Pix2Pix model [4], which uses a 2D generator and a 2D discriminator. In this paper, we denote it as Full 2D and train with NLL instead of  $L_1$  for a fair comparison. (2) A Hybrid 3D model consisting of a 3D generator and a 2D discriminator, similar to the method proposed in [7].

Fig. 2 shows the reconstruction results in three directions. Enhancement of spatial consistency can be observed from images reconstructed by Full 3D methods compared to the Full 2D and Hybrid 3D methods, particularly along the X-Z and Y-Z directions. Similarly in Fig. 3, when the reconstructed EM images are evaluated by the downstream task, i.e., 2D membrane segmentation, the visual quality of the results generated by Full 3D methods is better than two baseline methods. Quantitatively, as shown in Table 1 and 2, Full 3D methods outperform the Full 2D and the Hybrid 3D methods with considerable margins in most cases, in terms of PSNR and SSIM for reconstruction, as well as JS and Dice for segmentation. Meanwhile, Table 1 shows that using NLL (Full 3D) instead of  $L_1$  (Full 3D w/  $L_1$ ) not only provides uncertainty measurement but also enhances the reconstruction quality. And Table 2 shows that using segmentation-consistency loss (Full 3D + Seg) can improve the segmentation quality.



**Fig. 4.** Visualization of uncertainty estimation. Reconstruction & Uncertainty: Gaussian output distribution (mean & variance) generated by the models. Error:  $L_1$  norm between the reconstruction and the ground truth EM.

### 3.3. Uncertainty estimation

Intuitively, error maps are unable to capture visual ambiguity, and are not always accessible in real application scenarios. Uncertainty estimation, on the other hand, can highlight the uncertain areas that are consistent with the artifacts in the reconstructed images. For example, in the second row corresponding to Hybrid 3D in Fig. 4, the model clearly hallucinates a long and slender object in the reconstruction which is not shown in the ground truth EM. The uncertainty map correctly highlights this artifact, providing a richer interpretation than the error map.

## 4. CONCLUSIONS

In this work, we propose a cross-modality reconstruction approach to reconstruct high-quality EM-like images from low-quality X-ray images for connectomics. We find that the Full 3D model is the key to improving the reconstruction and

membrane segmentation results. Meanwhile, we also show the utilization of negative log-likelihood could improve the model performance as well as provide interpretable uncertainty maps. We believe our proposed approach will potentially allow the reconstruction of large-scale brain structures achieved in a fast and scalable way.

## 5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using animal subject data made available in open access by [2]. Ethical approval was not required as confirmed by the license attached with the open access data.

## 6. ACKNOWLEDGMENTS

This research is funded by the NSF (IIS-1607189, CCF-1563880, IOS-1452593, 1806818) and the NIH (EB032217, MH117808, MH128949). Aaron T. Kuan is supported by NIH (EB032217), Wei-Chung Allen Lee is supported by NIH (MH128949), Lu Mi is supported by a Mathworks Fellowship, and Yaron Meirovitch is supported by NIH (5U24NS109102) & NIH (U01 NS108637). This work is also supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n°852455). We acknowledge the ESRF for granting beamtime for proposal LS2892.

## 7. REFERENCES

- [1] Narayanan Kasthuri, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, Amelio Vázquez-Reina, Verena Kaynig, Thouis Raymond Jones, et al., “Saturated reconstruction of a volume of neocortex,” *Cell*, vol. 162, no. 3, pp. 648–661, 2015.
- [2] Aaron T Kuan, Jasper S Phelps, Logan A Thomas, Tri M Nguyen, Julie Han, Chiao-Lin Chen, Anthony W Azevedo, John C Tuthill, Jan Funke, Peter Cloetens, et al., “Dense neuronal reconstruction through x-ray holographic nano-tomography,” *Nature neuroscience*, vol. 23, no. 12, pp. 1637–1643, 2020.
- [3] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [4] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros, “Image-to-image translation with conditional adversarial networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [5] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [6] Mehdi Mirza and Simon Osindero, “Conditional generative adversarial nets,” *arXiv preprint arXiv:1411.1784*, 2014.
- [7] Guodong Zeng and Guoyan Zheng, “Hybrid generative adversarial networks for deep mr to ct synthesis using unpaired data,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 759–767.
- [8] Yongsheng Pan, Mingxia Liu, Chunfeng Lian, Tao Zhou, Yong Xia, and Dinggang Shen, “Synthesizing missing pet from mri with cycle-consistent generative adversarial networks for alzheimer’s disease diagnosis,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 455–463.
- [9] Zizhao Zhang, Lin Yang, and Yefeng Zheng, “Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network,” in *Proceedings of the IEEE conference on computer vision and pattern Recognition*, 2018, pp. 9242–9251.
- [10] Linjing Fang, Fred Monroe, Sammy Weiser Novak, Lyndsey Kirk, Cara R Schiavon, Seungyeon B Yu, Tong Zhang, Melissa Wu, Kyle Kastner, Alaa Abdel Latif, et al., “Deep learning-based point-scanning super-resolution imaging,” *Nature methods*, vol. 18, no. 4, pp. 406–416, 2021.
- [11] Hongda Wang, Yair Rivenson, Yiyin Jin, Zhensong Wei, Ronald Gao, Harun Günaydin, Laurent A Bentolila, Comert Kural, and Aydogan Ozcan, “Deep learning enables cross-modality super-resolution in fluorescence microscopy,” *Nature methods*, vol. 16, no. 1, pp. 103–110, 2019.
- [12] Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Culley, et al., “Content-aware image restoration: pushing the limits of fluorescence microscopy,” *Nature methods*, vol. 15, no. 12, pp. 1090–1097, 2018.
- [13] Alex Kendall and Yarin Gal, “What uncertainties do we need in bayesian deep learning for computer vision?,” *Advances in neural information processing systems*, vol. 30, 2017.




- [14] Stefan Klein, Marius Staring, Keelin Murphy, Max A Viergever, and Josien PW Pluim, “Elastix: a toolbox for intensity-based medical image registration,” *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 196–205, 2009.
- [15] Denis P Shamonin, Esther E Bron, Boudewijn PF Lelieveldt, Marion Smits, Stefan Klein, and Marius Staring, “Fast parallel image registration on cpu and gpu for diagnostic classification of alzheimer’s disease,” *Frontiers in neuroinformatics*, vol. 7, pp. 50, 2014.
- [16] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [17] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.

## 8. SUPPLEMENTARY MATERIAL

### 8.1. Visualization of domain differences

In this paper, we use a segmentation network pre-trained on a public benchmark EM membrane segmentation dataset to provide segmentation consistency constraints, in the absence of segmentation labels in X-Ray2EM dataset. Both datasets are EM images of mouse cortex and we show in Fig. 5 that the domain difference between these two datasets has a negligible negative effect on the segmentation results.



**Fig. 5.** Comparison of domain differences between benchmark EM dataset and X-Ray2EM dataset.

### 8.2. Additional 2D membrane segmentation results

In Fig. 6, the reconstructed EM images are evaluated by downstream tasks, i.e., 2D membrane segmentation. It is consistent with the results (X-Z direction) in the main paper that along Y-Z and X-Y directions, the segmentation quality of EM reconstructed by Full 3D method is also better than the two baseline methods.

### 8.3. Ablation studies

#### 8.3.1. $L_1$ loss v.s. NLL

In our work, we use NLL as one of the training objectives instead of  $L_1$  loss. Results in the main paper show that the reconstruction results of Full 3D model trained with NLL are better than or comparable with that of Full 3D w/  $L_1$  model in most cases, indicating that using NLL for training not only provides the estimation of uncertainty but can also improve the reconstruction performance.

#### 8.3.2. W/ seg v.s. w/o seg

According to the results shown in the main paper, Full 3D + Seg method achieve comparable results to Full 3D w/o Seg on the reconstruction task and prevails on the 2D membrane segmentation task, indicating that features like neuronal boundaries are better preserved by applying segmentation-consistency constraints.

### 8.4. Software and hardware requirements


All experiments in the paper are implemented in Python 3.6.12 using PyTorch 1.8.0 framework on a computation server with a 2.30GHz Intel(R) Xeon(R) Gold 5218 CPU and a GeForce RTX 3090 GPU.

### 8.5. Segmentation network pre-training

The segmentation network has a U-Net architecture originated from [16]. We use a cross-entropy loss function and train the network for 3000 iterations with a batch size of 4. Adam [17] optimizer is used with an initial learning rate of 0.0001 and a weight decay of 0.00005.

### 8.6. Reconstruction framework training

The generator has a U-Net like architecture and the discriminator is a PatchGAN, both are 3D versions of the 2D Pix2Pix model proposed in [4]. Input to the generator is randomly cropped from the whole volume to a size of  $128 \times 128 \times 64$  and the output reconstruction also has the same size. As for the pre-trained 2D segmentation network following the generator, we randomly select one section from each of the three directions of the reconstructed EM volume to perform the inference process, leading to three segmentation probability



**Fig. 6.** Segmentation results: visualization of 2D segmentation probability map of membrane along the Y-Z and the X-Y directions.

maps. Only the generator and the discriminator participate in the training process, and the parameters of the segmentation network are fixed. The whole framework is also trained using Adam optimizer in an end-to-end fashion for 100 epochs, with a constant learning rate of 0.0002 for the initial 50 epochs and linearly decayed to 0 for the remaining 50 epochs. The batch size is 1. Loss function weights  $w_{\text{NLL}}$ ,  $w_{\text{GAN}}$ , and  $w_{\text{Seg}}$  are 0.00002, 1.0, and 1.0, respectively.

### 8.7. X-Ray2EM dataset

For the EM modality, we have a serial section TEM dataset. For the X-ray modality, we have an X-ray holographic nanotomography dataset [2]. These datasets can be found at:

- <https://www.lee.hms.harvard.edu/kuan-phelps-et-al-2020>

### 8.8. Public EM segmentation benchmark dataset

To pre-train the segmentation network, we use a serial section SEM dataset of mouse cortex with corresponding membrane labels from Kasthuri. et al [1]. This dataset can be found at:

- <https://lichtman.rc.fas.harvard.edu/vast/>