

MIT Open Access Articles

Learning Guided Electron Microscopy with Active Acquisition

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Mi, Lu et al. "Learning Guided Electron Microscopy with Active Acquisition." MICCAI: International Conference on Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science, 12265, Springer, 2020, 77-87. © 2020 Springer Nature

Published Version: http://dx.doi.org/10.1007/978-3-030-59722-1_8

Publisher: Springer International Publishing

Permanent Link: <https://hdl.handle.net/1721.1/130096>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



Learning Guided Electron Microscopy with Active Acquisition

Lu Mi¹, Hao Wang¹, Yaron Meirovitch^{1,2}, Richard Schalek², Srinivas C. Turaga³, Jeff W. Lichtman², Aravinthan D.T. Samuel², Nir Shavit¹

¹ Massachusetts Institute of Technology, MA, USA

² Harvard University, MA, USA

³ HHMI Janelia Research Campus, VA, USA

lumi@mit.edu

Abstract. Single-beam scanning electron microscopes (SEM) are widely used to acquire massive data sets for biomedical study, material analysis, and fabrication inspection. Datasets are typically acquired with uniform acquisition: applying the electron beam with the same power and duration to all image pixels, even if there is great variety in the pixels' importance for eventual use. Many SEMs are now able to move the beam to any pixel in the field of view without delay, enabling them, in principle, to invest their time budget more effectively with non-uniform imaging. In this paper, we show how to use deep learning to accelerate and optimize single-beam SEM acquisition of images. Our algorithm rapidly collects an information-lossy image (e.g. low resolution) and then applies a novel learning method to identify a small subset of pixels to be collected at higher resolution based on a trade-off between the saliency and spatial diversity. We demonstrate the efficacy of this novel technique for active acquisition by speeding up the task of collecting connectomic datasets for neurobiology by up to an order of magnitude. Code is available at <https://github.com/lumi9587/learning-guided-SEM>.

Keywords: Electron Microscope · Active Acquisition · Determinantal Point Process.

1 Introduction

Scanning electron microscopes are widely used for nanometer-scale imaging in diverse applications including structural biology [14,9], materials analysis [22], and semiconductor fabrication [21]. In most cases, an electron beam is applied with the same power and duration to all image pixels. This is fundamentally inefficient since the saliency of each pixel might be heterogeneous.

In this paper, we design an imaging strategy for a scanning electron microscope (SEM) that mimics the human visual system. The human visual system quickly decides saliency by first using a low-resolution (non-foveal) collection and then applies the high-resolution fovea to dwell on important parts [25]. Most scanning electron microscopes are ideally suited for such active and adaptive

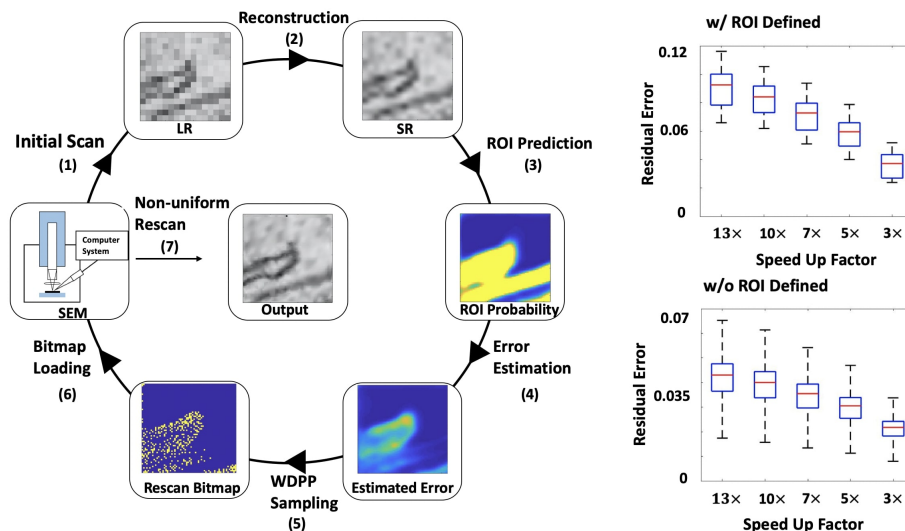


Fig. 1: **Left:** Overview of our learning guided electron microscope with ROI defined. The SEM firstly performs the initial scan to collect the low-resolution image I_{LR} (1). The reconstruction step is applied to generate a super-resolved image I_{SR} (2). Then the pixel of interest (3) and estimated error (4) is identified. The next step is to implement diversified sampling based on estimated error using weighted determinantal point process (WDPP) (5). Finally a sparse bitmap is sent to SEM (6) to perform rescan (7). See more details in Section 2.1. **Right:** The box plots represent different speedup factors with residual error (L_1 loss) using our active acquisition pipeline on *SNEMI3D*.

image acquisition; their scan generators can instantaneously deflect the beam to any arbitrary position and then stably dwell to acquire each pixel [1,20]. Because of fast beam re-positioning, an arbitrary distribution of pixel locations has negligible impact on scanning time. Therefore, for an image where only 5% of pixels need scanning, one can ideally speed up its acquisition by 20 folds [1,23].

Most images are characterized by two types of saliency: *interest* and *information density*. Conventional SEM ignores both types of saliency and scans the whole image using identical resolution with low efficiency [7]. In this work, we develop a new method to accelerate and optimize SEM acquisition. To do this, we have built an active acquisition algorithm that intelligently budgets the operation of a scanner to focus on regions with high saliency (either high interest or high information density) within an image.

We apply our technology to one critical area using SEM in connectomics, which aims for the reconstruction of synaptic connectivity maps for brain tissue [17,14]. So far, connectomics has only been applied to a small number of specimens because of the enormous burden in acquiring and analyzing datasets that can easily span terabytes and petabytes [13,29]. For example, scanning a cubic millimeter of brain tissue at the resolution needed for connectomics

($4nm \times 4nm \times 30nm$ per voxel) requires two thousand trillion voxels (2PBs of data). A typical SEM running at one million voxels per second would require 63 years to do this.

Here, we design to accelerate SEM image acquisition for connectomics by exploiting the sparsity of salient pixels in images. The essential goal in connectomics is to map neural circuitry [12,18]. The only structures that contain pixels with high interest are membrane of neurons, and objects associated with synapses. All other intracellular objects and extracellular space have low saliency. For typical mammalian tissues, neural membranes account for only 5-10% of the images and synapses account for even less [14]. This sparsity of salient pixels suggests significant potential for speeding up acquisition. In this work, without loss of generality, we define the membrane boundary of individual neurons as the region of interest (ROI), to demonstrate the effectiveness of our pipeline.

In contrast to previous work using image reconstruction [5,8,16,11,28,27,6,2], or multi-beam approaches that use highly parallelized but expensive microscope systems [4], our work aims to guide widely available single-beam SEMs to collect salient image pixels, thereby reconstructing essential regions at high resolution, as shown in Fig. 1. Our major contributions are:

- We are the first, to our knowledge, to cast the acquisition of electron microscopes as a learning-guided sampling problem and thereby capable to achieve significant speedup.
- We present an effective and principled sampling technique, weighted determinantal point process (WDPP), that optimizes pixel selection based on their saliency and spatial diversity.
- We present a new active-acquisition pipeline for SEM that executes non-uniform pixel-wise scanning, and demonstrate a potential speedup rate of up to an order of magnitude on real-world connectomic datasets.

2 Methodology

In this paper, we formulate a learning-guided sampling problem to speed up SEM acquisition. The goal is to intelligently sample a subset of pixels in a way that balances the following trade-off:

$$\mathcal{C} = \underbrace{|I_{HR} - R(I_{HR} \odot B)|}_{\text{reconstruction loss}} + \lambda \underbrace{\sum_{i,j} B^{(i,j)}}_{\text{acquisition cost}} .$$

Here I_{HR} denotes a high-resolution image, B is the bitmap with the superscript indexing positions to indicate locations of sampled pixels, $R(\cdot)$ is the reconstruction function, \odot is the Hadamard (element-wise) product, and λ is a hyperparameter balancing the trade-off between reconstruction loss and acquisition cost. Unlike other works [16,11,28,27] focusing on improving $R(\cdot)$ given a low-resolution image $I_{LR} = I_{HR} \odot B$, our work assumes a fixed $R(\cdot)$ and instead tries to find a reasonable bitmap B that can achieve low reconstruction loss with

low acquisition cost. We do this via the proposed WDPP sampling technique, which selects pixels based on spatial diversity as well as saliency (quantified as estimated error). In the following, we describe the overview of our active acquisition pipeline, introduce its key components (i.e., binarized error estimation in Section 2.2 and WDPP in Section 2.3) as well as our technical contributions.

2.1 Active Acquisition Pipeline

Below we describe individual steps of the active acquisition pipeline in Fig. 1.

Initial Scan, Reconstruction, and ROI Prediction: As the first step, SEM performs the *initial scan* of a low-resolution image I_{LR} with negligible cost. The next step is to apply a *reconstruction* model R to I_{LR} to generate super-resolved image $I_{SR} = R(I_{LR})$. The model can use either learning-based reconstruction methods or simple interpolation rules such as bicubic. The third step (*ROI Prediction*) is using ROI detector F_{ROI} to predict the saliency score for each pixel. Note that we are interested in two types of tasks in this paper: tasks with and tasks without ROI defined. *ROI Prediction* is not applied for the task without ROI defined. This task only considers regions with high information density as saliency. These steps correspond to (1)-(3) in Fig. 1.

Binarized Error Estimation: The fourth step, shown as (4) of Fig. 1, is to *estimate the prediction error*. For the tasks without ROI defined, the ground truth error is defined as L_1 loss $|I_{HR} - I_{SR}|$; for the tasks with ROI defined, the ground truth error is defined as L_1 loss $|F_{ROI}(I_{HR}) - F_{ROI}(I_{SR})|$. These errors will be estimated through an efficient and simple learning based method we propose in Section 2.2.

Diversified Sampling: This step, as shown in (5) of Fig. 1, is to perform WDPP sampling based on the estimated error map from the previous step. The goal is to select K samples contributing the largest estimated error while balancing the spatial diversity at the same time. Details are in Section 2.3.

Bitmap Loading and Rescan: Once the locations of sampled pixels (produced by WDPP) are encoded into bitmaps and loaded into the SEM, the SEM will perform rescan based on the sparse bitmap. The final output I_{OUT} is then reconstructed with recollected pixels during rescan as well as pixels in I_{LR} collected in the initial scan.

2.2 Binarized Error Estimation

One key component of the pipeline is our proposed *binarized error estimation*. It is a simple and efficient supervised learning method to estimate pixel-wise error.

As shown in Fig. 2, we first binarize the continuous pixel-wise error using the mean of the error distribution ϵ as the threshold and then train a UNET [24] to perform classification. For the task with ROI defined, our goal is to train a classification network to output the pixel-wise probability (we refer to this probability as *estimated error* in the following text):

$$P(|F_{ROI}(I_{HR}^{(i,j)}) - F_{ROI}(I_{SR}^{(i,j)})| > \epsilon),$$




Fig. 2: The overview of our end-to-end learning framework to train the error estimation networks F_{ES} . For the task without ROI defined (**left**), the inputs of F_{ES} are concatenation of I_{SR} and I_{LR} . For the task with ROI defined (**right**), the inputs of F_{ES} are concatenation of $F_{ROI}(I_{SR})$ and I_{SR} . The regression task is reformulated as a binary classification problem after applying a threshold to the ground truth error. ROI detector F_{ROI} and reconstruction model R are fixed during the training of F_{ES} .

where the superscript (i, j) indexes positions. For the task without ROI defined, our goal is to predict the probability:

$$P(|I_{HR}^{(i,j)} - I_{SR}^{(i,j)}| > \epsilon).$$

Our preliminary experiments show significant improvement in error estimation compared to directly regressing the error. This is because most pixels in an image have very low error, significantly biasing the model to output low values.

2.3 Diversified Sampling

After the estimated error is acquired, one naive strategy is to rescan K pixels with the highest estimated error. However, due to strong correlation between neighboring pixels, a more cost-effective way is to sample pixels according to both saliency (quantified by estimated error) and spatial diversity. To this end, we formulate the problem as a determinantal point process (DPP) [15].

Moreover, another key contribution in our work is to propose a weighted DPP and construct a proper DPP kernel L balancing saliency and spatial diversity. Specifically, given an image with size of $M \times M$, we construct an $N \times N$ kernel $L = U^\gamma S U^\gamma$, where $N = M^2$ is the total number of pixels in the image, S is a $N \times N$ symmetric matrix indicating location similarity. U is an $N \times N$ diagonal

matrix; each diagonal entry u_{ii} indicates pixel i 's saliency, which is quantified by the estimated error described in Section 2.2. The exponent, γ , controls the trade-off between saliency and spatial diversity. For pixel i and j , we have

$$L_{ij} = u_{ii}^\gamma S_{ij} u_{jj}^\gamma, \quad S_{ij} = e^{-[(x_i - x_j)^2 + (y_i - y_j)^2] / \sigma_s^2}, \quad (1)$$

where σ_s is a hyperparameter. With this new diversified sampling algorithm, our pipeline can select K pixels simultaneously for rescan, while guaranteeing efficiency. The algorithm is shown in Algorithm 1. In brief, conventional DPP sampling promises the diversity of sampled points for each iteration; in the current iteration, DPP finds a point which is diverse from all previous points. In contrast, our proposed WDPP finds a point which strikes a balance between diversity and saliency in each iteration.

Algorithm 1: WDPP Sampling

Input: Location similarity matrix S and diagonalized quality matrix U .

Construct kernel matrix $L = U^\gamma S U^\gamma$.

Compute the eigen-decomposition $(v_n, \lambda_n)_{n=1}^N$ of L .

$J \leftarrow \emptyset$.

for $n = 1, 2, \dots, N$ **do**

$J \leftarrow J \cup \{n\}$ with prob. $\frac{\lambda_n}{\lambda_n + 1}$.

end

$V \leftarrow \{v_n\}_{n \in J}, Y \leftarrow \emptyset$.

while $|V| > 0$ **do**

 Select i from y with $Pr(i) = \frac{1}{|V|} \sum_{v \in V} (v^\top e_i)^2$.

$Y \leftarrow Y \cup i$.

$V \leftarrow V_\perp$, an orthonormal basis for the subspace of V orthogonal to e_i .

end

Output: Y .

3 Experiments

In this section, we provide an in-depth analysis of all components in our pipeline. We use two real-world connectomics datasets, *SNEMI3D* from a mouse cortex (with a resolution of $3 \times 3 \times 30$ nm/pixel), and *Human* from a human cerebrum (with a resolution of $4 \times 4 \times 30$ nm/pixel), to evaluate our algorithm. We generate low-resolution images from original images using nearest neighbor with down-sampling rates of $\times 4$, $\times 8$, and $\times 16$.

3.1 Reconstruction and ROI Detection

We reconstruct I_{SR} from I_{LR} using a UNET. Specifically, we explore two variants of UNET, one trained with an adversarial loss plus an L1 loss (SRGAN) [11]

Model	DS rate	Task	Ours	Entropy	Interest	Task	Ours	Trans	Gradient
Bicubic	× 4	w/ ROI	0.549	0.420	0.256	w/o ROI	0.488	–	0.323
	× 8		0.460	0.306	0.329		0.437	–	0.267
	× 16		0.307	0.249	0.213		0.349	–	0.200
SRGAN	× 4	w/ ROI	0.367	0.402	0.252	w/o ROI	0.599	0.451	0.347
	× 8		0.495	0.262	0.294		0.514	0.351	0.215
	× 16		0.389	0.111	0.317		0.461	0.293	0.136
SRUNET	× 4	w/ ROI	0.448	0.416	0.238	w/o ROI	0.451	0.399	0.324
	× 8		0.382	0.374	0.228		0.405	0.311	0.250
	× 16		0.412	0.163	0.046		0.315	0.203	0.187

Table 1: Correlation between ground truth error and error estimated from our method on *SNEMI3D*. Baselines are interest, entropy, infer-transformation, and gradient.

and one trained with only an L1 loss (SRUNET) [24]. Interestingly, we find that whether the adversarial loss improves reconstruction quality depends on I_{HR} 's noise level (see results in Supplementary). Note that reconstruction alone does not guarantee high-quality output, which is why we need adaptive rescan after binarized error estimation and diversified sampling.

For ROI detection, the ROI detector F_{ROI} also uses a UNET. It is trained with I_{HR} as input and human-annotated membrane as ground-truth labels.

3.2 Error Estimation Analysis

To guide the rescan process using estimated error (proposed in Section 2.2) is an effective and efficient way for active acquisition. For the first task without ROI defined, the residual error after reconstruction is $|I_{HR} - I_{SR}|$, and we use gradient [3] and infer-transformation for uncertainty estimation [19,26] as baselines. For the second task with ROI defined, the residual error after reconstruction is $|F_{ROI}(I_{HR}) - F_{ROI}(I_{SR})|$. We use $F_{ROI}(I_{SR})$ (interest) and its corresponding entropy as baseline.

The first metric is the pixel-wise correlation between the estimated error and ground-truth error. As shown in Table 1, our method can estimate error much more accurately than the baselines (visualizations in Supplementary). This method is also robust to different reconstruction methods and down-sampling rates. The second metric is sparsification error curve [10], which shows how residual error decays as SEM increases the rescan rate. Specifically, we select the top K pixels according to different measurements, e.g., estimated error (our method) and entropy, and set the corresponding ground truth errors of these top K pixels to zero. We then study the decay of error when K increases. Fig. 3 shows the results for different methods including random sampling [5,8]. Our method achieves the fastest decay.




Fig. 3: Sparsification error curves for estimated error (from our method) and baselines on *SNEMI3D*.




Fig. 4: The effect of γ to balance the trade-off between spatial diversity and saliency for WDPP sampling. All rescan bitmaps contain the same number of samples.

3.3 Sampling with Saliency & Diversity

With the error estimation U , we can then construct the kernel matrix L according to Equation (1) and run the WDPP sampling algorithm to select pixels to rescan. Fig. 4 shows the rescan bitmap produced by WDPP when $\gamma \in \{1, 2, 5\}$. As expected, (1) WDPP can naturally trade off saliency and spatial diversity during sampling; (2) compared to using only error estimation as rescan bitmaps, WDPP can cover larger areas within a sampling budget. Column 4 to 6 in Fig. 5 show the final outputs from SEM following different rescan schemes, i.e., random, estimated error, and estimated error with WDPP, demonstrating WDPP can significantly improve output image quality given a fixed rescan budget.

4 Performance Evaluation

In this section, we evaluate the overall speedup using our active acquisition pipeline, as shown in Fig. 1. The quality of the final output I_{OUT} is compared with I_{HR} . For the task with ROI defined, the residual error is quantified as $|F_{ROI}(I_{HR}) - F_{ROI}(I_{OUT})|$. For the task without ROI defined, the residual error is quantified as $|I_{HR} - I_{OUT}|$ (evaluations using PSNR and SSIM are shown in Supplementary). We evaluate the initial scan with a down-sampling ratio of $\times 4$; we use different total scan rates (initial scan plus rescan) inversely proportional to the speedup factor, since the total run time of our computational pipeline on a single GPU is a small fraction ($< 3\%$) of the SEM imaging time. The results demonstrate that our pipeline can achieve a speedup factor of up to an order of magnitude with relatively small error.

5 Conclusion

We propose a novel and efficient learning-guided sampling algorithm based on learned saliency and spatial diversity. Our active acquisition pipeline demonstrates a potential speedup rate of up to an order of magnitude for SEM in connectomic data collection. In a broader sense, our work addresses research issues across many fields where high-throughput SEM is an essential tool for discovery. Techniques we propose in this work may also be widely used to speed up other imaging systems where sparse scanning can be applied.

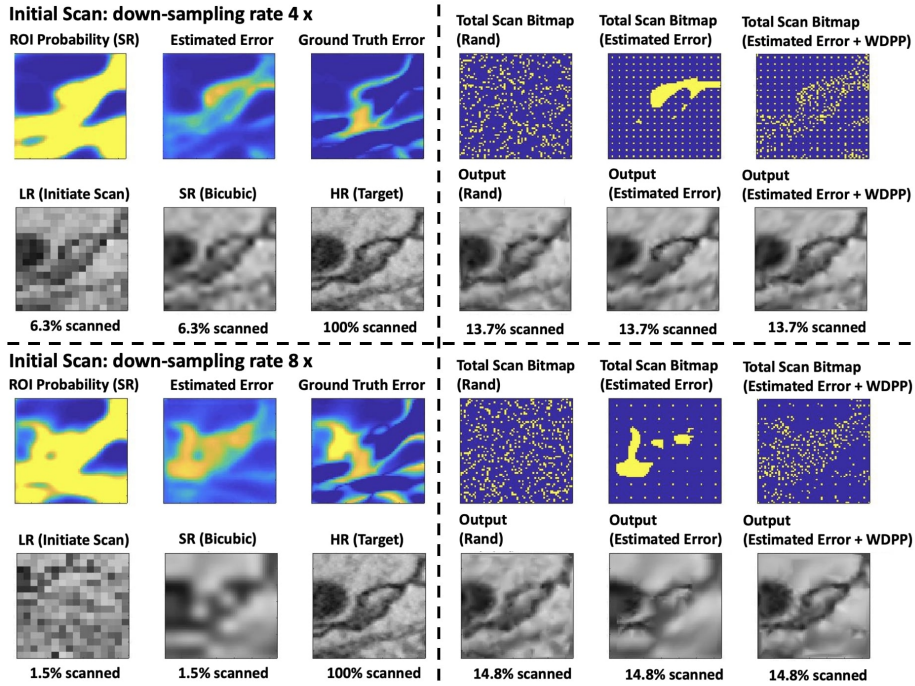


Fig. 5: The final outputs after rescan on *SNEMI3D*. We compare WDPP sampling based on estimated error maps with other baselines, using down-sampling rates of $\times 4$ and $\times 8$ for initial scan. The right part shows different total scan bitmaps and their corresponding outputs using different methods. The total scan bitmaps contain the pixel locations scanned in both initial scan and rescan.

Acknowledgement We would thank Remco Schoenmakers, Pengfei Guo for insightful comments and suggestions. This is an area of active investigation in our laboratory and others (Uri Manor, personal communication). This research was supported by the National Science Foundation (NSF) under grants IIS-1607189, CCF-1563880, IOS-1452593 and NSF 1806818.

References

1. Anderson, H.S., Ilic-Helms, J., Rohrer, B., Wheeler, J., Larson, K.: Sparse imaging for fast electron microscopy. In: Computational Imaging XI. vol. 8657, p. 86570C. International Society for Optics and Photonics (2013)
2. Buchholz, T.O., Krull, A., Shahidi, R., Pignino, G., Jékely, G., Jug, F.: Content-aware image restoration for electron microscopy. In: Methods in cell biology, vol. 152, pp. 277–289. Elsevier (2019)
3. Dahmen, T., Engstler, M., Pauly, C., Trampert, P., De Jonge, N., Mücklich, F., Slusallek, P.: Feature adaptive sampling for scanning electron microscopy. Scientific reports **6**, 25350 (2016)

4. Eberle, A., Mikula, S., Schalek, R., Lichtman, J., Tate, M.K., Zeidler, D.: High-resolution, high-throughput imaging with a multibeam scanning electron microscope. *Journal of microscopy* **259**(2), 114–120 (2015)
5. Eldar, Y.C., Kutyniok, G.: *Compressed sensing: theory and applications*. Cambridge university press (2012)
6. Fang, L., Monroe, F., Novak, S.W., Kirk, L., Schiavon, C.R., Seungyoon, B.Y., Zhang, T., Wu, M., Kastner, K., Kubota, Y., et al.: Deep learning-based point-scanning super-resolution imaging. *bioRxiv* p. 740548 (2019)
7. Flegler, S.L., Flegler, S.L.: *Scanning & Transmission Electron Microscopy*. Oxford University Press (1997)
8. Gan, L.: Block compressed sensing of natural images. In: 2007 15th International conference on digital signal processing. pp. 403–406. IEEE (2007)
9. Helmstaedter, M., Briggman, K.L., Denk, W.: High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nature neuroscience* **14**(8), 1081–1088 (2011)
10. Ilg, E., Cicek, O., Galesso, S., Klein, A., Makansi, O., Hutter, F., Brox, T.: Uncertainty estimates and multi-hypotheses networks for optical flow. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 652–667 (2018)
11. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
12. Januszewski, M., Kornfeld, J., Li, P.H., Pope, A., Blakely, T., Lindsey, L., Maitin-Shepard, J.B., Tyka, M., Denk, W., Jain, V.: High-precision automated reconstruction of neurons with flood-filling networks. *Nature Methods* **15**, 605–610 (2018), <https://www.nature.com/articles/s41592-018-0049-4>
13. Jarrell, T.A., Wang, Y., Bloniarz, A.E., Brittin, C.A., Xu, M., Thomson, J.N., Albertson, D.G., Hall, D.H., Emmons, S.W.: The connectome of a decision-making neural network. *Science* **337**(6093), 437–444 (2012)
14. Kasthuri, N., Hayworth, K.J., Berger, D.R., Schalek, R.L., Conchello, J.A., Knowles-Barley, S., Lee, D., Vázquez-Reina, A., Kaynig, V., Jones, T.R., Roberts, M., Morgan, J.L., Tapia, J.C., Seung, S., Roncal, W.G., Vogelstein, J.T., Burns, R., Sussman, D.L., Priebe, C.E., Pfister, H., Lichtman, J.W.: Saturated reconstruction of a volume of neocortex. *Cell* **162**(3), 648–661 (2015)
15. Kulesza, A., Taskar, B., et al.: Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning* **5**(2–3), 123–286 (2012)
16. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4681–4690 (2017)
17. Lichtman, J.W., Pfister, H., Shavit, N.: The big data challenges of connectomics. *Nature neuroscience* **17**(11), 1448–1454 (2014)
18. Meirovitch, Y., Mi, L., Saribekyan, H., Matveev, A., Rolnick, D., Shavit, N.: Cross-classification clustering: An efficient multi-object tracking technique for 3-d instance segmentation in connectomics. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8425–8435 (2019)
19. Mi, L., Wang, H., Tian, Y., Shavit, N.: Training-free uncertainty estimation for neural networks. *arXiv preprint arXiv:1910.04858* (2019)
20. Mohammed, A.: *Scanning electron microscopy (sem): A review* (12 2018)
21. Newell, T., Tillotson, B., Pearl, H., Miller, A.: Detection of electrical defects with semvision in semiconductor production mode manufacturing. In: 2016 27th Annual

- SEMI Advanced Semiconductor Manufacturing Conference (ASMC). pp. 151–156. IEEE (2016)
22. Pandey, K., Setua, D., Mathur, G.: Material behaviour: Fracture topography of rubber surfaces: An sem study. *Polymer testing* **22**(3), 353–359 (2003)
 23. Potocek, P., Trampert, P., Peemen, M., Schoenmakers, R., Dahmen, T.: Sparse scanning electron microscopy data acquisition and deep neural networks for automated segmentation in connectomics. *Microscopy and Microanalysis* pp. 1–10 (2020)
 24. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
 25. Thorpe, S., Fize, D., Marlot, C.: Speed of processing in the human visual system. *nature* **381**(6582), 520 (1996)
 26. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Test-time augmentation with uncertainty estimation for deep learning-based medical image segmentation. *arXiv preprint arXiv:1807.07356* (2018)
 27. Wang, H., Rivenson, Y., Jin, Y., Wei, Z., Gao, R., Günaydin, H., Bentolila, L.A., Kural, C., Ozcan, A.: Deep learning enables cross-modality super-resolution in fluorescence microscopy. *Nat. Methods* **16**, 103–110 (2019)
 28. 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 (2018)
 29. Yan, G., Vértés, P.E., Towilson, E.K., Chew, Y.L., Walker, D.S., Schafer, W.R., Barabási, A.L.: Network control principles predict neuron function in the *caenorhabditis elegans* connectome. *Nature* **550**(7677), 519 (2017)

Appendix

Dataset	Method	Task	$\times 4$	$\times 8$	$\times 16$	Task	$\times 4$	$\times 8$	$\times 16$
SNEMI3D	Bicubic		0.089	0.195	0.288		0.047	0.076	0.103
	SRGAN	w/ ROI	0.088	0.212	0.344	w/o ROI	0.058	0.091	0.124
	SRUNET		0.074	0.165	0.335		0.043	0.069	0.092
Human	Bicubic		0.510	0.700	0.726		0.146	0.167	0.199
	SRGAN	w/ ROI	0.064	0.111	0.220	w/o ROI	0.166	0.191	0.225
	SRUNET		0.664	0.712	0.725		0.134	0.154	0.184

Table 2: The residual error after reconstruction. For the task with ROI defined, the residual error is quantified as $|F_{ROI}(I_{HR}) - F_{ROI}(I_{SR})|$. For the task without ROI defined, residual error is quantified as $|I_{HR} - I_{SR}|$. Results are evaluated on different reconstruction methods, Bicubic, SRUNET, and SRGAN on *SNEMI3D* and *Human*.




Fig. 6: The visualizations of reconstruction image quality for I_{SR} . Different reconstruction methods, Bicubic, SRGAN, and SRUNET are evaluated on *SNEMI3D* and *Human*. Adding adversarial loss (SRGAN) improves the reconstruction quality on *Human* (see details in Table 2), where I_{HR} is relatively noisy. The red and yellow boxes mark the artifacts after reconstruction.

Task	Model	Method	$\times 4$	$\times 8$	$\times 16$
w/ ROI	SRGAN	Ours	0.596	0.449	0.351
		Entropy	0.552	0.469	0.255
		Interest	0.255	0.234	0.190
Task	Model	Method	$\times 4$	$\times 8$	$\times 16$
w/o ROI	SRUNET	Ours	0.184	0.179	0.151
		trans	0.201	0.158	0.126
		Gradient	0.098	0.126	0.107

Table 3: Pixel-wise correlation between ground-truth error and error estimated from our method on *Human*. Our results are compared with other baselines, interest, entropy, variance from infer-transformation, and gradient.

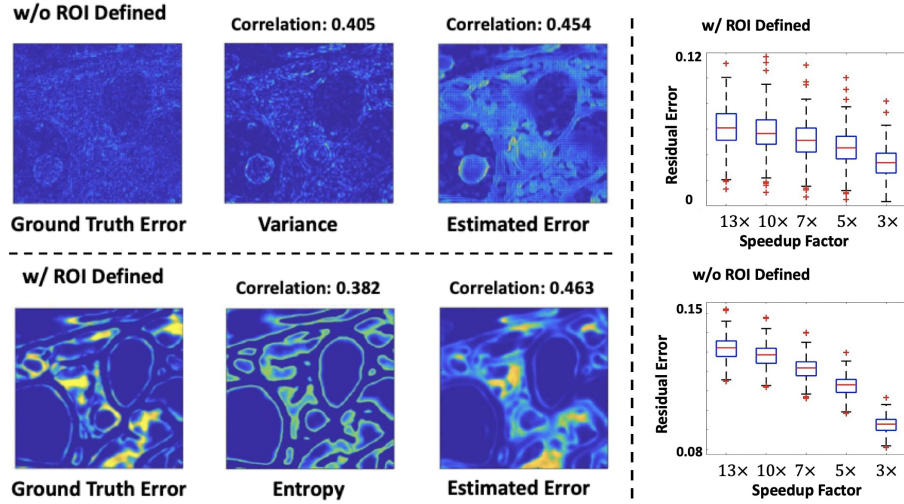


Fig. 7: **Left:** The visualizations of estimated error from our method, compared with other baselines: interest, entropy, and variance from infer-transformation on *SNEMI3D*. The value of pixel-wise correlation is also identified. **Right:** The box plots represent different speedup factors v.s. residual error (L_1 loss) using our active acquisition pipelines on *Human*. The residual error of final output I_{OUT} is quantified as $|F_{ROI}(I_{HR}) - F_{ROI}(I_{OUT})|$ for the task with ROI defined, and as $|I_{HR} - I_{OUT}|$ for the task without ROI defined.

Speedup	3 ×	5 ×	7 ×	10 ×	13 ×
L_1 Loss	0.021	0.030	0.035	0.039	0.042
SSIM	0.815	0.751	0.716	0.686	0.668
PSNR	29.48	27.26	26.16	25.21	24.57

Table 4: Speedup v.s. quality for the task without ROI defined on *SNEMI3D*. The quality of final output I_{OUT} is compared with I_{HR} and evaluated with L_1 Loss, SSIM and PSNR.