Title:

Deep learning-based volumetric denoising enables efficient acquisition of volume electron microscopy

Authors:

Bohao Chen,[1,2] Fangfang Wang,[3] Haoyu Wang,[3] Yanchao Zhang,[2,4] Zhuangzhuang Zhao,[2] Haoran Chen,[2,5] Hua Han,[2,4,*] Xi Chen,[2,*] Yunfeng Hua[3,6,*]

Affiliations:

[1]School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Science, Beijing, China

[2]State Key Laboratory of Brain Cognition and Brain-inspired Intelligence Technology, Institute of Automation, Chinese Academy of Sciences, Beijing, China

[3]Shanghai Institute of Precision Medicine, Ninth People's Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai, China

[4]School of Future Technology, University of Chinese Academy of Sciences, Beijing, China

[5]School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

[6]Lead contact

E-mail: hua.han@ia.ac.cn; xi.chen@ia.ac.cn; yunfeng.hua@shsmu.edu.cn

SUMMARY:

Volume electron microscopy (VEM) enables three-dimensional (3D) visualization of specimens through serial sectioning and nanometer-resolution imaging. To alleviate daunting difficulties in subsequent data analysis, people tend to employ excessively slow imaging and high resolution, limiting the acquisition throughput. Here, we titrated the combinations of acquisition parameters for preserving key structural information. This revealed a prior role of sufficient spatial sampling over high image contrast. To save acquisition time, we then attempted to restore serial images that contained either low-contrast snapshots or skipped sections using machine learning. Owing to constraints of cytoarchitecture rationality, 3D context-based (volumetric) denoising not only preserved more structural features but also generated fewer artifacts than other methods. Moreover, volumetric denoising lowered the requirement of imaging dwell time and thereby allowed serial block-face removal down to 20 nm because of reduced radiation damage. This work demonstrated how machine learning-based image processing enabled the optimization of VEM acquisition parameters.

MAIN TEXT

Introduction:

Volume electron microscopy (VEM) is a series of rapidly evolving technologies that enable the automated acquisition of a large image stack either from consecutive ultra-thin sections or iteratively exposed block faces of a resin-embedded sample.[1-3] Although both serial section and block face-types of VEM can provide a lateral resolution of about 5 nm, the minimal cutting thickness that can be maintained over hundreds of cycles differs between them due to distinct mechanics.[2,4-6] Typically, a commercial ultra-microtome can produce high-quality serial sections by cutting with the diamond blade at a step size of 30-40 nm, below which increased occurrences of varying cutting quality and section damage are observed.[7] In contrast, serial-block-face (SBF) scanning electron microscope (SEM) can homogeneously remove the sample surface at every 25 nm with an in-chamber ultramicrotome, as the resin block is much more robust against mechanical deformation than ultra-thin sections.[8] Among all the VEM, focused ion beam (FIB) SEM provides the highest axial resolution owing to its capability of sample surface milling with Ga ions at a precision of up to 5 nm.[6,9] Considering limited setup availability and reasonable running costs, it is therefore unsurprising that most VEM volumes (except for very few acquired by FIB-SEM) have been registered with anisotropic voxel sizes to date. Although the importance of improved axial resolution is stressed for the sake of not only an error-free structure annotation (utmost important for neurites) but also the development of automated segmentation algorithms, the implication of anisotropic resolution on the preservation of structural details and segmentation difficulty has not yet been studied comprehensively.[4]

Moreover, with the ambition to map tissue volumes at an ever-increasing scale, both spatial and temporal sampling rates need to be taken into account. Collective efforts have been made to accelerate VEM image acquisition, including sample contrast methods with saturated heavy metal deposition,[10-12] fast scopes empowered by a camera array[13] or multiple electron beams,[14] automated and parallelized imaging pipelines,[15] etc. Nevertheless, a trade-off among pixel size, acquisition duration, and signal-to-noise ratio (SNR) of images is often made by users of commercial VEM for a feasible or affordable imaging plan. Moreover, for SBF and FIB-SEM, improved

image quality by prolonging pixel dwell time is frequently accompanied by compromised axial resolution, as heightened electron dose may lead to severe radiation damage to the sample block face that hinders the next round of ultra-thin surface removal.[4,5] This raises a key question for almost every VEM acquisition, namely, should a higher priority be given to either fine imaging or thin sectioning?

For the above dilemma, recent developments in artificial intelligence (AI)-based image restoration may offer a promising solution without compromise. By learning features from a variety of published VEM datasets, state-of-the-art methods have enabled image denoising by diffusion[16] or general[17] models, as well as generating datasets of isotropic voxel size through nonlinear interpolation.[18-21] However, the accuracy of model inference was predominantly assessed by comparison between the ground truth and synthetic data recovered from artificial degradations or among different model outcomes, leaving the general application of those specifically trained models poorly tested. Moreover, image denoising and pixel interpolation are independent tasks, for which AI models demonstrate differential performance, leading us to seek a machine-friendly VEM acquisition strategy with regard to maximum recovery of structural information using the AI model of choice.

Here, we utilized a public FIB-SEM dataset and simulated the influence of various acquisition conditions on dataset quality. Next, we compared specific AI performances on improving the dataset quality of test stacks acquired using a hybrid slow and fast imaging routine on an SBF-SEM. Last, we demonstrated a use-dependent optimization of SBF-SEM acquisition for a self-supervised AI denoising.

Results:

Structural information fidelity under different acquisition settings

To explore quantitatively the implication of section thickness and image background noise on the VEM dataset quality, we took a publicly accessible high-quality FIB-SEM dataset of the mouse brain tissue[22] to generate a series of synthetic datasets through noise masking and spatial down-sampling (Fig. 1A). First, the original image stack with an isotropic voxel size of 5 nm was down-sampled in both lateral (XY) and axial (Z) dimensions to yield new stacks at 10-nm and 15-nm voxel sizes. Next, to simulate low-SNR datasets acquired from the same sample due to shortened pixel dwell times, Poisson-Gaussian noise of anticipated intensity was added to the raw images (Fig. 1B). Meanwhile, Z-extraction of the stack was carried out by excluding every two, three, or four sections, respectively. To ease comparisons of dataset quality among combinations of thicker cutting and faster imaging, we equalized the scales of Z-extraction and noise level by normalizing to the estimated acquisition time based on reported reference values[23] (Fig. S1, also see Materials and Methods).

Quantitative assessments of these synthetic datasets were carried out by computing their effective spatial resolution (Fig. 1C and D). Notably, the obtained effective resolution of data with a high Z-sampling rate was robust to varying image SNRs (Fig. 1D), forming a “comfort zone” for relatively stable data quality upon different imaging speeds estimated based on the SNR. These observations were found to be consistent across all the simulated scenarios of further data down-sampling (Fig. 1D inset). Besides modeling with synthetic datasets, this spatial sampling-dependent influence of SNR on the VEM dataset quality was validated further by physically acquired FIB-SEM datasets from a different brain sample (Fig. 1E- G). These observations are instructive for planning VEM acquisition when a trade-off between spatial resolution and imaging speed needs to be made.

To quantify the putative impact of VEM data quality on the accuracy of organelle segmentation, we compared the performance of automated mitochondrial segmentation among the synthetic datasets (Fig. 2A). Similarly, accurate segmentations (low 1-IoU values) were achievable from the datasets of high Z-sampling rate despite a modest SNR (Fig. 2B). Further, we show that the segmentation error increases linearly with

the effective resolution (Fig. 2C), arguing for a primary role of high spatial resolution in minimizing errors during automated organelle segmentation.

Deep learning-based restoration of serial images

Considering that anisotropic VEM voxel sizes are primarily caused by mechanical challenges in reliably cutting thinner slices or creating uniform block faces at smaller step sizes, the above results substantiate a yet underestimated role of section thickness in counteracting noise and maintaining VEM dataset quality for automated image processing. We speculate that spatial oversampling contributes to robust digital preservation of biological structures, which are continuous across consecutive sections. To test whether such 3D structural constraints could be utilized for AI-based restoration of serial images, we designed a hybrid imaging routine that allowed the acquisition of arbitrary interlaced high-SNR reference and low-SNR snapshot slices by alternating imaging speeds. In practice, a mouse brain sample was cut consecutively at 50 nm and imaged at 2 µs as well as 0.5 µs pixel dwell time using a commercial SBF-SEM (Fig. 3A). The obtained images were arranged in a use-dependent fashion to generate different types of test datasets (Fig. 3B), including a reference stack (R-stack, all scanned at 2 µs/pixel) and a snapshot stack (S-stack, all scanned at 0.5 µs/pixel) as well as a subset of references by skipping every other slices (½R-stack) and a hybrid dataset with a fixed combination of ¾ snapshots and ¼ references (S/r-stack).

These datasets provided a unique opportunity to compare AI performances on tasks including interpolation, 2D denoising, and newly established volumetric denoising (see Fig. S2, Materials and Methods for the network architecture and more details) based on a shared ground truth (Fig. 3C). First, as one would anticipate, significantly fewer artifacts and blurring regions were observed in the denoised images of S/r- and S-stacks using additional information from 3D structural context. Second, it was unexpected that the outcomes of volumetric denoising models were comparable between using input datasets with and without high-SNR images (S/r- and S-stack). This implies that doping of reference frames in the input stack has a minimal contribution to the inference accuracy of denoising models. Third, we noticed better structural preservation in terms of less aliasing and blurring in the volumetric denoised images (S/r- and S-stack) than the interpolated ½R-stack, again stressing the pivotal

role of axial sampling beyond the image SNR. These observations were further confirmed by quantifying the effective resolution of the AI output stacks (Fig. 3D). In fact, volumetric denoising achieved even a slightly worse score of data quality using S/r- than S-stack as the input (Fig. 3E), because the denoised outcome of S/r-stack contains those reference frames with certain residual noise (Fig. S3). More importantly, the volumetric denoised image series preserved maximally inter-slice continuity compared to interpolated or 3D denoised ones (Fig. 3F).

Optimization of VEM acquisition for volumetric denoising

As for SBF-SEM, the axial resolution is mainly limited by the mechanical cutting of the resin block. Prolonged pixel dwell time heightened the radiation damage, which alters the mechanical properties of the block face and makes it more difficult to cut. Effective volumetric denoising of a dataset containing exclusively fast-scanned images creates a margin for cutting at smaller step sizes by SBF-SEM, and in turn, the AI performance may benefit from an increased axial sampling rate. To test this idea, we acquired three image stacks from a sample using 0.5 μs dwell time but at different cutting thicknesses (50 nm, 25 nm, and 20 nm). Due to the absence of high-SNR reference frames, we adopted a completely self-supervised strategy for the model training. Based on visual inspection, the AI delivered reasonable denoising with input stacks acquired at all tested cutting thicknesses (Fig. 4A-C, Supplementary Video 1 and 2), but further quantification revealed an improvement of the effective resolution in the XZ plane from 850.3 ± 361.8 nm to 209.5 ± 6.6 nm (50 nm cutting), 306.5 ± 98.9 nm to 104.1 ± 3.3 nm (25 nm cutting), and 263.6 ± 97.0 nm to 83.6 ± 2.0 nm (20 nm cutting) after blind volumetric denoising by AI (Fig. 4D). Note that the achieved effective XZ resolution by denoising were found to decrease linearly with the physical reduction of cutting thickness (Fig. 4E), serving as a good indication for minimal residual noise in the model outcomes. By comparing the cutting-thickness-independent XY component of effective resolution, we further excluded that such improvements in data quality were due to inter-sample variability (Fig. S4). As importantly, in comparison with denoising without using 3D context information, our volumetric approach gained additional benefit from reduced cutting thickness in suppressing the occurrences of unrealistic large structural changes between adjacent slices (Fig. 4F).

Quantitatively, significant improvements in inter-slice continuity were observed in datasets of brain tissues that were sliced thinner than 75 nm (Fig. 4G).

Discussion

In principle, VEM imaging of biological specimens is a 3D-digitalization process, in which the sampling rate in space and time corresponds to voxel size and its SNR. For most subcellular structures, the lateral resolution of modern electron microscopy is more than sufficient; therefore, structural details in XY are often oversampled. In the third dimension, however, the increase in section number is substantially limited by technical challenges or time constraints, hindering isotropic voxel sizes of VEM. In the present work, we have comprehensively studied the trends of VEM dataset quality changes induced by different trade-offs between Z-sampling and image SNR (Fig. 1D, G). By doing so, we characterized the so-called "comfort zone" for VEM acquisition at adequate Z-sampling rates, within which the improvement of effective resolution and segmentation accuracy is saturated despite enhanced SNR. Although tissue type-specific calibration may be required, this approach is particularly instructive for imaging parameter selection to maximize the acquisition efficiency of VEM. Especially, emerging techniques, such as gas-cluster ion beam[24] and parallel ion beam source,[25] are capable of pushing the thickness of large-scale surface etching down to 10-nm precision, paving the way for industrial-scale tissue mapping at desired resolutions.

Moreover, our quantification provides a fact-based augment for the pivotal role of sufficient spatial sampling in maintaining structural information fidelity, superior to high image contrast (Fig. 1D, G). In line with this notion, our quantitative comparison among AI performance on various image restoration tasks (Fig. 3D) suggests that volumetric denoising outperforms interpolation in real experiment settings enabled by the hybrid acquisition routine. This finding is significant in two aspects. First, it provides proof-of-principle evidence for learning structural features with spatial continuity by AI to accurately classify stable signals and random noise in the VEM stacks. Second, as computer devices are much more powerful in memorizing and comparing large image sequences with a wide contrast range than human beings, this stimulates new thinking of suitable VEM sampling strategies primarily for AI-based image processing rather than visual assessment. Indeed, those high-SNR reference images that are favorable to visual inspection turn out to be not essential for accurate

volumetric denoising by AI (Fig. 3E). Nevertheless, our attempt with such a hybrid acquisition may serve as a human-in-the-loop paradigm for future method development of AI-based image processing beyond denoising, for instance, super-resolution or defocus correction.

Notably, self-supervised volumetric denoising no longer requires any reference frame (Fig. 4D and G), enabling calibration-free applications on VEM volumes of different tissue types and setups. More importantly, it provides a new perspective for serial block-face imaging with improved Z-resolution or more robust surface removal, which are negatively influenced by radiation damage in the practice of both SBF and FIB-SEM. To the best of our knowledge, the actual cutting thickness of SBF-SEM is still limited by 25 nm,[26-28] while our AI-denoising enables an acquisition with shortened dwell time and improved axial resolution at least to 20 nm using a commercial setup (Fig. 4E). As inter-section differences become smaller due to a reduced cutting thickness, we expect an improved Z-alignment of serial images, which may further contribute to the performance of volumetric denoising.

Lastly, this study provides an example of how AI can be systematically integrated into the data acquisition of VEM, rather than being applied for a *post hoc* image enhancement. When the model outcomes are used to inform optimal acquisition parameter selection in real time, VEM imaging can be accelerated while maintaining biologically meaningful resolution. Besides, the hybrid acquisition routine proposed in this study is fundamentally distinct from conventional fixed-parameter EM acquisition schemes, allowing imaging conditions to be adaptively tailored to experimental goals during acquisition. Importantly, this flexibility is achieved without compromising the reliable identification of key ultrastructural features, thereby enabling application-specific trade-offs between throughput and image quality. In line with this idea, specific AI has been recently developed for predicting regions containing synapses and thin neurites to be rescanned at high resolution and prolonged dwell time.[29]

In summary, the present work argues for a pivotal role of axial resolution in maintaining VEM data quality by numerical scoring. In line with this notion, volumetric denoising appears to be a more efficient and effective task for image stack

restoration than generative interpolation using machine learning. This result triggers new thinking of adapting our current VEM sampling strategies in favor of AI-based processing, which can substantially differ from our current gold standard – human assessment, which is highly limited by our visual perception.

Materials and Methods

Animal experiment

Adult C57BL/6 mice (male) were purchased from Shanghai Jihui Experimental Animal Feeding Co., Ltd. (Shanghai, China). All procedures were approved by the ethical committee of Shanghai Ninth People's Hospital (No. SH9H-2019-A567-1) and conducted at Shanghai Institute of Precision Medicine.

The mouse brain samples were prepared according to the previously published procedure[10] with minor modifications. In brief, animals were anesthetized with isoflurane inhalation and transcardially perfused with 30 mL 1.5 M cacodylate buffer (CB, pH 7.4), followed by an ice-cold fixative mixture containing 2% paraformaldehyde and 2.5% glutaraldehyde in 0.08 M CB. The dissected brain tissues were post-fixed at 4°C in the same fixative for 48 hours and then sliced 500-μm thick using a vibratome (Leica VT 1200S) in 0.15 M CB. The *en bloc* staining was done by multiple rounds of incubation with heavy metals, including 2% OsO_4 (in 0.15 M CB), 2.5% ferrocyanide (in 0.15 M CB), 2% OsO_4 (in 0.15 M CB) at room temperature (RT) for 1.5, 1.5, and 1.0 hours, respectively. After being washed with 0.15 M CB and water for 30 min each, the samples were incubated at RT in thiocarbonhydrazide (saturated aqueous solution) for 1 hour at RT, 2% OsO_4 aqueous solution for 1.5 hours at RT, and 1% uranyl acetate (aqueous solution) for 2 hours at 50°C with intermediate washing step (twice in water for 30 min each). The sample dehydration was carried out through a graded acetone series (50%, 75%, 90%, 30 min each, all cooled at 4°C) and in pure acetone (100%, three times, 30 min each). Finally, the samples were infiltrated through a 1:1 mixture of acetone and Spurr's resin monomer (4.1g ERL 4221, 0.95g DER 736, 5.9g NSA, and 1% DMAE) overnight on a rotator, then embedded in pure resin for 8 hours before being baked in a 70°C oven for 72 hours.

Data acquisition of SBF-SEM

The sample was mounted on an aluminum metal rivet using epoxy adhesive for 12 hours and trimmed to a rectangular shape (0.8×0.8 mm^2) using a sample trimmer (Leica, TRIM2). Next, the sample surface was prepared by a diamond knife using an ultramicrotome (Leica, UC7). Sputter coating of a thin layer of gold (Leica, EM ACE600) was applied to the sample block, which was then imaged using a field-emission SEM (Zeiss, Gemini300) equipped with an in-chamber ultramicrotome (Gatan, 3View2XP) and back-scattered electron detector (Gatan, OnPoint).

The SEM imaging was carried out with an incident beam energy of 2 keV and a pixel size of 15 nm. For Fig. 2, serial images (4000×4000 pixels, 365 slices) were acquired at a dwell time of 0.5 μs with a cutting thickness of 50 nm. For each imaging and cutting cycle, repeated scans with a dwell time of 2.0 μs were manually triggered to capture redundant images of improved quality. These serial images were assigned in use-dependent orders to compose the reference stack (all frames acquired with 2.0-μs dwell time), the snapshot stack (all frames acquired with 0.5-μs dwell time), and the fast-slow-scan dataset (a combination of one 2.0-μs-dwell-time frame and three 0.5-μs-dwell-time frames). For Fig. 3, the snapshot stack was reused for the 50-nm-cutting experiment. And two independent datasets were acquired using the same imaging parameters, except that the cutting thickness was reduced to 25 nm (902 slices) and 20 nm (626 slices). All SBF-SEM datasets contained at least 301 slices for model training

and testing after Z-alignment using a self-written MATLAB script based on cross-correlation maximum between consecutive slices before analysis.[30] Data simulation on a public FIB-SEM volume

To simulate dataset quality under different image settings (Fig. 1), a publicly available dataset was used, which was acquired from the mouse CA1 hippocampus region using FIB-SEM at an isotropic voxel size of 5 nm (https://www.epfl.ch/labs/cvlab/data/data-em/).[22] For the XY down-sampling, binning of image pixels at a factor of 2 and 3 was used, yielding a final XY pixel size of 10 and 15 nm. To perform down-sampling along the Z-direction at a factor of k, the new image stack kept every k slice from the original stack. For a given volume, this process resulted in a k-fold increase in the slice thickness, a k-fold reduction in the number of imaged slices, and an equivalent 1/k fraction of the original acquisition time.

To simulate the effect of a reduced pixel dwell time, we manipulated the image signal-to-noise ratio by introducing Poisson-Gaussian noise of different levels to the original images. The outcomes were scored using the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE)[23] for a quantitative calibration. In detail, the variance of the noise used for masking was adjusted so that the obtained BRISQUE score of the masked image matched closely the documented value for images acquired with a specific dwell time.[16] By doing so, we were able to generate a series of datasets with comparable noise levels to images scanned with different dwell times of 0.5 µs, 1.0 µs, 1.5 µs, 2.0 µs, 4.0 µs, and 16.0 µs, respectively (Fig. S1).

Generating Authentic Datasets under Different VEM Conditions

From the mouse cortex sample, a 3D volume was imaged using the Plasma-FIB setup (Thermo Scientific Hydra Bio). Except for pixel dwell time, the same imaging parameters were used. This yielded a series of authentic FIB-SEM datasets, each consisting of 332 serial sections of 2048 × 2048 pixels at a 5 × 5 × 5 nm^3 isotropic voxel size, scanned at different dwell times of 2 µs, 1.5 µs, 1 µs, 0.75 µs, 0.5 µs, 0.25 µs, respectively. The image stacks of different SNRs were then down-sampled along the Z-axis in a use-dependent manner for further analysis.

Computing effective spatial resolution of EM volumes

Sectional Fourier Shell Correlation (SFSC) analysis was used to estimate the effective spatial resolution of the cubed VEM stacks.[31] This method determined directional resolution by calculating the normalized cross-correlation within wedge-shaped sectors of individual Fourier shells. For analysis of an EM volume, the input data was first split into two pairs of sub-volumes based on pixel coordinate on the XY plane:

Pair 1: One sub-volume containing all pixels at (even, even) coordinates and the other containing pixels at (odd, odd) coordinates.

Pair 2: A second pair formed from pixels at (even, odd) and (odd, even) coordinates.

The SFSC for each pair was then calculated using the normalized cross-correlation between their Fourier transforms:

$$FSC_{pair_i}(r_i, \alpha_i) = \frac{\left(\sum_{r\in r_i, \alpha\in\alpha_i} F_1(r, \alpha, \phi) \cdot F_2(r, \alpha, \phi)\right)}{\sqrt{\sum_{r\in r_i, \alpha\in\alpha_i} F_1^2(r, \alpha, \phi) \cdot \sum_{r\in r_i, \alpha\in\alpha_i} F_2^2(r, \alpha, \phi)}}, \quad (1)$$

where $F_1(r, \alpha, \phi)$ and $F_2(r, \alpha, \phi)$ denote the voxels in two Fourier-transformed sub-volumes within pair_i that are located at a given distance r_i from the origin and within an orientation sector, defined by α and ϕ. Finally, the SFSC curves on each pair were averaged to yield the final SFSC measurement for the input volume:

$$SFSC = \frac{SFSC_{\text{pair}_1} + SFSC_{\text{pair}_2}}{2} \quad (2)$$

To ensure computational efficiency while obtaining a robust measurement for the entire volume, our analysis was performed on 50 sub-volumes of a uniform shape (128 × 128 × 128) randomly sampled from the input volume. The final effective spatial resolution reported for a given volume is the average of these individual SFSC measurements.

Python package for the SFSC analysis was downloaded from an open-source Github repository (sakoho81/miplib: A Python software library with a variety of functions for (optical) microscopy image restoration, reconstruction and analysis.) and implemented for this study. The resolution cut-off was determined by a default criterion (SNR-0.5) recommended from official examples, and the obtained directional resolution profiles were visualized as a polar plot. For comparison among EM volumes, we used the resolution cut-offs on the 90° axis (corresponding to the XZ plane), on which this value was expected to be equally sensitive to image quality and slice thickness.

Quantification of automated mitochondrial segmentation

The data quality was assessed based on the mitochondrial segmentation accuracy of a published 3D AI segmentation model combining a U-Net with Swin-Transformer blocks (STU-Net, uni-medical/STU-Net: The largest pre-trained medical image segmentation model (1.4B parameters) based on the largest public dataset (>100k annotations), up until April 2023., arXiv:2304.06716). Models were trained independently for down-sampled image stacks of isotropic voxel sizes (5, 10, and 15 nm). The training data and its corresponding ground truth were provided by the aforementioned publicly available FIB-SEM dataset. The training was conducted for 200 epochs using a Stochastic Gradient Descent (SGD) optimizer with momentum of 0.99 and a fixed learning rate of 1×10^{-3}. The model from the final epoch was used for inference. Each model was applied to its corresponding set of testing data with the same pixel size, and the segmentation performance was quantified on a slice-by-slice basis by calculating the Intersection-over-Union (IoU) against the human ground truth, which was provided by the same public dataset and processed with the corresponding down-sampling factors. We plotted the segmentation error as 1-IoU (Fig. 1E) to ease visual comparison with the effect resolution of the corresponding dataset (Fig. 1D).

AI-based Z-interpolation of serial images

The input image stack was acquired at a 2-µs dwell time and spaced at 100 nm in Z (Fig. 2B). For the interpolation, we fed the pretrained Real-Time Intermediate Flow Estimation (RIFE) model,[32] which was previously shown to have outstanding

performance on the FIB-SEM data,[33] with the input stack. By a 2-fold interpolation, the RIFE artificially improved the Z-spacing of the original dataset from 100 nm to 50 nm.

Quantification of Inter-Slice Structural Continuity

To quantitatively evaluate the 3D structural continuity, we calculated the mean Structural Similarity Index (SSIM) between consecutive slices.[34] The metric assesses the perceptual similarity between two images (x and y) by comparing luminance, contrast, and structure. This metric is defined as:

$$SSIM(x, y) = l(x, y)^{\alpha} \cdot c(x, y)^{\beta} \cdot s(x, y)^{\gamma}, \tag{3}$$

where luminance comparison $l(x, y)=\frac{2\mu_x\mu_y+C_1}{\mu_x^2+\mu_y^2+C_1}$, contrast comparison $c(x, y)=\frac{2\sigma_x\sigma_y+C_2}{\sigma_x^2+\sigma_y^2+C_2}$, and structure comparison $s(x, y)=\frac{\sigma_{xy}+C_3}{\sigma_x\sigma_y+C_3}$. The terms μ_x, μ_y, σ_x, σ_y and σ_{xy} are the local means, standard deviations, and cross-covariance for the corresponding images. The constants $C_1=(K_1L)^2$, $C_2=(K_2L)^2$, and $C_3=C_2/2$ stabilize the calculation, where L is the dynamic range of the pixel values. In our experiments, this metric was implemented using parameters $K_1=0.01$ and $K_2=0.03$. This calculation was applied in two distinct ways:

1. Independent analysis: To compare the overall continuity of different restoration methods, 50 sub-volumes (256 × 256 × 24) were independently and randomly sampled from each full volume. The mean SSIM was calculated between every two images within each cropped sub-volume to generate a distribution as a continuity measure.

2. Paired analysis: To compare 2D-denoising and V-denoising as a function of cutting thickness (25 nm to 200 nm), a paired-sampling analysis was conducted. For each thickness, 50 sub-volumes (256 × 256 × 24) were randomly sampled from identical coordinates in both the 2D-denoised and V-denoised image stacks. The mean SSIM was calculated for each sub-volume, resulting in 50 paired scores for statistical testing. The calibration baseline (dotted line) was established by sampling 500 random, non-adjacent slice pairs from volumes and computing the mean SSIM.

Volumetric denoising network architecture and model training

For volume denoising, we adopted a cascaded U-Net-based deep denoising model named GShiftNet[35] as the core architecture (Fig. S4). The proposed model contained an initial feature extraction block enhanced by channel attention, followed by an encoder for fusing multi-frame features by channel-wise circular shifts and 8-direction spatial shifts. The processed features were then passed through three cascaded U-Net sub-networks for deep representation learning. Throughout the network, Channel Attention Blocks (CABs) refined features by adaptively reweighting channel responses via global average pooling, making the model focus on informative signals while maintaining stability through residual connections. The final reconstruction module aggregated multi-stage features and combined them with the input via a global residual path, followed by a convolutional layer to produce the denoised output.

All volumetric denoising training was performed using 3D patches of a fixed size (128 × 128 × 5). While the initial patches were all randomly sampled from the data volumes, the precise method of generating the training pairs differed for each workflow, as detailed in the following subsections. Standard data augmentation techniques, including random flips and rotations, were applied to each patch prior to the network feeding. Each epoch consisted of 540 training pairs and 32 validation pairs. All models were trained for a fixed 100 epochs, and the model checkpoint from the final epoch was used for all subsequent inference. During training, the model was optimized using a weighted sum of the L1 and structural similarity loss (defined as 1–SSIM). The loss function L_{total} used in training is defined as:

$$L_{total}=L_1+(1\text{-}SSIM) \tag{4}$$

Training was performed with Adam (initial LR 2 × 10-4, $\beta_1 = 0.9$, $\beta_2 = 0.99$) and a plateau-based learning rate scheduler. All experiments were run on a Linux server equipped with a single Nvidia RTX 4090 GPU.

Supervised training of volumetric denoising with serial images

For the supervised denoising experiments (Fig. 2), two distinct data preparation pipelines were implemented to generate source and target pairs to train our denoising model:

1. For the fast and slow imaging of S/r-stack, whose input sequence consisted of five consecutive frames (the 1st and 5th frames were captured by slow-scan; the 2nd to 4th frames were from fast scan), an approach was employed to prevent the model from learning a trivial identity map for the slow-scanned frames. To construct paired training samples, a larger patch of size (256 × 256 × 5) was initially extracted from both the hybrid stack and the corresponding high-SNR ground truth. From these, the input and supervision patches (128 × 128 × 5) were derived using a deterministic checkerboard subsampling scheme: the input patch was composed of pixels located at even rows and even columns (top-left), while the supervision patch was constructed from the diagonally opposite subset (bottom-right). This design enabled robust end-to-end denoising of hybrid patches.

2. For the fast-scan mode, training pairs were generated by directly extracting a 3D patch (128 × 128 × 5) from the snapshot stack and pairing it with the corresponding patch from the ground truth at the same spatial location. This pairing strategy enabled the model to learn a direct mapping from noisy inputs to their clean counterparts. For the comparable 2D supervised denoising model (Fig. 2), the training process was identical, except that the input patch depth was reduced to one (128 × 128 × 1). This minimal modification ensured a fair comparison between the 2D and 3D approaches by keeping the model architecture, number of training samples, and optimization strategy consistent.

Self-supervised training of volumetric denoising with serial images

For the self-supervised denoising experiments lacking paired low-noise references (Fig. 4), training pairs were generated directly from the raw images. The pairing strategy followed the same principle as the supervised model training scheme

described above. Specifically, a randomly cropped patch (256 × 256 × 5) was divided into two sub-patches using a deterministic checkerboard pattern, which were then assigned as the source and target, respectively. This blind-spot design encouraged the network to learn underlying structural representations rather than simply fitting noise patterns.

Volumetric denoising inference of serial images

For inference on large volumes, a patch-based approach was implemented using an overlapping sliding window strategy. To prevent block artifacts in the denoised stacks, adjacent patches (128 × 128 × 5) were processed with a 12.5% overlap in the XY plane (16 × 16), and 20% (one layer) overlap in the Z-axis. The final denoised pixels in these overlapping regions were the mean of the denoised outcomes from the adjacent patches. While this inference pipeline was applied to all denoising strategies, the hybrid fast-slow mode included an additional post-processing step to ensure a fair comparison with the interpolation method. After the initial seamless volume was generated, the slices corresponding to the positions of the original slow-scan frames (e.g., the 1st and 5th for every patch) were replaced with the original, unprocessed high-SNR slow-scan images. This step guaranteed that both volumetric denoising methods and the RIFE interpolation method share identical data at the ground truth layers.

Statistical analysis

All statistical tests were performed using GraphPad Prism 10.1.2 built-in functions. One-sample t tests for Fig. 3D (½SO, ****$p < 0.0001$, $n = 50$, $t = 49.78$, $df = 49$; S/r-stack, ****$p < 0.0001$, $n = 50$, $t = 17.07$, $df = 49$, two-tailed; S-stack, ****$p < 0.0001$, $n = 50$, $t = 13.29$, $df = 49$, all two-tailed), Fig. 4D (effective XZ resolution: 50-nm, ****$p < 0.0001$, $n = 50$, $t = 12.58$, $df = 49$; 25-nm, ****$p < 0.0001$, $n = 50$, $t = 14.72$, $df = 49$; 20-nm, ****$p < 0.0001$, $n = 50$, $t = 13.28$, $df = 49$, all two-tailed), and Fig. 4G (inter-slice continuity: 25-nm, ****$p < 0.0001$, $n = 50$, $t = 23.16$, $df = 88.19$; 50-nm, ****$p < 0.0001$, $n = 50$, $t = 7.85$, $df = 78.16$; 75-nm, *$p = 0.0195$, $n = 50$, $t = 2.39$, $df = 79.01$, all two-tailed). Two-sample t tests for Fig. 3E (interpolated *vs.* S/r-denoised, ****$p < 0.0001$, $n = 50$, $t = 12.15$, $df = 98$; interpolated vs. S-denoised, ****$p < 0.0001$, $n = 50$, $t = 15.44$, $df = 98$; S/r-denoised *vs.* S-denoised, ****$p < 0.0001$, $n = 50$, $t = 18.38$, $df = 98$), Fig. 3F (R-stack *vs.* ½R-stack, ****$p < 0.0001$, $n = 50$, $t = 16.14$, $df = 62.18$; R-stack *vs.* Interpolated, ****$p < 0.0001$, $n = 50$, $t = 9.039$, $df = 97.10$; Interpolated *vs.* ½R-stack, ****$p < 0.0001$, $n = 50$, $t = 30.00$, $df = 64.86$; S-stack *vs.* 2D-denoised, ****$p < 0.0001$, $n = 50$, $t = 15.63$, $df = 57.26$; S-stack *vs.* V-denoised, ****$p < 0.0001$, $n = 50$, $t = 17.87$, $df = 56.43$; 2D-denoised *vs.* V-denoised, *$p = 0.0262$, $n = 50$, $t = 2.258$, $df = 97.72$), Fig. 4D inset (effective XZ resolution: 50-nm *vs.* 25-nm, ****$p < 0.0001$, $n = 50$, $t = 101.2$, $df = 98$; 50-nm *vs.* 20-nm, ****$p < 0.0001$, $n = 50$, $t = 129.6$, $df = 98$; 25-nm *vs.* 20-nm, ****$p < 0.0001$, $n = 50$, $t = 37.23$, $df = 98$).

References

1. Collinson, L.M., Bosch, C., Bullen, A., Burden, J.J., Carzaniga, R., Cheng, C., Darrow, M.C., Fletcher, G., Johnson, E., Narayan, K., et al. (2023). Volume EM: a quiet revolution takes shape. Nat. Methods *20*, 777-782. 10.1038/s41592-023-01861-8.

2. Peddie, C.J., Genoud, C., Kreshuk, A., Meechan, K., Micheva, K.D., Narayan, K., Pape, C., Parton, R.G., Schieber, N.L., Schwab, Y., et al. (2022). Volume electron microscopy. Nat. Rev. Methods Primers *2*, 51. 10.1038/s43586-022-00131-9.
3. Helmstaedter, M. (2025). Synaptic-resolution connectomics: towards large brains and connectomic screening. Nat. Rev. Neurosci. Online ahead of print. 10.1038/s41583-025-00998-z.
4. Briggman, K.L., and Bock, D.D. (2012). Volume electron microscopy for neuronal circuit reconstruction. Curr. Opin. Neurobiol. *22*, 154-161. 10.1016/j.conb.2011.10.022.
5. Wanner, A.A., Kirschmann, M.A., and Genoud, C. (2015). Challenges of microtome-based serial block-face scanning electron microscopy in neuroscience. J. Microsc. *259*, 137-142. 10.1111/jmi.12244.
6. Xu, C.S., Hayworth, K.J., Lu, Z., Grob, P., Hassan, A.M., Garcia-Cerdan, J.G., Niyogi, K.K., Nogales, E., Weinberg, R.J., and Hess, H.F. (2017). Enhanced FIB-SEM systems for large-volume 3D imaging. Elife *6*, e25916. 10.7554/eLife.25916.
7. Schalek, R., Kasthuri, N., Hayworth, K., Berger, D., Tapia, J., Morgan, J., Turaga, S., Fagerholm, E., Seung, H., and Lichtman, J. (2011). Development of high-throughput, high-resolution 3D reconstruction of large-volume biological tissue using automated tape collection ultramicrotomy and scanning electron microscopy. Microsc. Microanal. *17*, 966-967. 10.1017/s1431927611005708.
8. Denk, W., and Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. PLoS Biol. *2*, e329. 10.1371/journal.pbio.0020329.
9. Knott, G., Marchman, H., Wall, D., and Lich, B. (2008). Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. J. Neurosci. *28*, 2959-2964. 10.1523/JNEUROSCI.3189-07.2008.
10. Hua, Y., Laserstein, P., and Helmstaedter, M. (2015). Large-volume en-bloc staining for electron microscopy-based connectomics. Nat. Commun. *6*, 7923. 10.1038/ncomms8923.
11. Mikula, S., and Denk, W. (2015). High-resolution whole-brain staining for electron microscopic circuit reconstruction. Nat. Methods *12*, 541-546. 10.1038/nmeth.3361.
12. Song, K., Feng, Z., and Helmstaedter, M. (2023). High-contrast en bloc staining of mouse whole-brain and human brain samples for EM-based connectomics. Nat. Methods *20*, 836-840. 10.1038/s41592-023-01866-3.
13. Bock, D.D., Lee, W.C., Kerlin, A.M., Andermann, M.L., Hood, G., Wetzel, A.W., Yurgenson, S., Soucy, E.R., Kim, H.S., and Reid, R.C. (2011). Network anatomy and in vivo physiology of visual cortical neurons. Nature *471*, 177-182. 10.1038/nature09802.
14. Eberle, A.L., Mikula, S., Schalek, R., Lichtman, J., Tate, M.L.K., and Zeidler, D. (2015). High-resolution, high-throughput imaging with a multibeam scanning electron microscope. J. Microsc. *259*, 114-120. 10.1111/jmi.12224.
15. Hayworth, K.J., Xu, C.S., Lu, Z., Knott, G.W., Fetter, R.D., Tapia, J.C., Lichtman, J.W., and Hess, H.F. (2015). Ultrastructurally smooth thick partitioning and volume stitching for large-scale connectomics. Nat. Methods *12*, 319-322. 10.1038/nmeth.3292.
16. Lu, C., Chen, K., Qiu, H., Chen, X., Chen, G., Qi, X., and Jiang, H. (2024). Diffusion-based deep learning method for augmenting ultrastructural imaging and volume electron microscopy. Nat. Commun. *15*, 4677. 10.1038/s41467-024-49125-z.
17. Sun, H., Ye, E., Lu, C., Jiang, Z., Bai, W., Ren, S., Wang, C., Zhang, J., Shi, R., Ma, L., et al. (2025). Preprint at Research Square, 10.21203/rs.3.rs-6602548/v1.
18. Ferede, F.A., Khalighifar, A., John, J., Venkataraman, K., and Khairy, K. (2025). Z-

UPSCALING: optical flow guided frame interpolation for isotropic reconstruction of 3D EM volumes. 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI).

19. Joshi, S., Forjaz, A., Han, K.S., Shen, Y., Queiroga, V., Selaru, F.A., Gérard, M., Xenes, D., Matelsky, J., Wester, B., et al. (2025). InterpolAI: deep learning-based optical flow interpolation and restoration of biomedical images for improved 3D tissue mapping. Nat. Methods *22*, 1556-1567. 10.1038/s41592-025-02712-4.
20. Wang, Z., Sun, G., Li, G., Shen, L., Zhang, L., and Han, H. (2022). STDIN: Spatio-temporal distilled interpolation for electron microscope images. Neurocomputing *505*, 188-202. 10.1016/j.neucom.2022.07.037.
21. Heinrich, L., Bogovic, J.A., and Saalfeld, S. (2017). Deep learning for isotropic super-resolution from non-isotropic 3D electron microscopy. In Medical Image Computing and Computer-Assisted Intervention − MICCAI 2017, pp. 135-143. 10.1007/978-3-319-66185-8_16.
22. Lucchi, A., Li, Y., and Fua, P. (2013). Learning for structured prediction using approximate subgradient descent with working sets. 2013 IEEE Conference on Computer Vision and Pattern Recognition.
23. Mittal, A., Moorthy, A.K., and Bovik, A.C. (2011). Blind/referenceless image spatial quality evaluator. 2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR).
24. Hayworth, K.J., Peale, D., Januszewski, M., Knott, G.W., Lu, Z., Xu, C.S., and Hess, H.F. (2020). Gas cluster ion beam SEM for imaging of large tissue samples with 10 nm isotropic resolution. Nat. Methods *17*, 68-71. 10.1038/s41592-019-0641-2.
25. Ma, H., Lv, Y., Wang, M., Zhao, Z., Zhang, L., Chang, S., Liu, J., Chen, B., Dong, X., Zhang, Y., et al. (2025). Preprint at bioRxiv, 10.1101/2025.04.25.650569.
26. Helmstaedter, M., Briggman, K.L., Turaga, S.C., Jain, V., Seung, H.S., and Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. Nature *500*, 168-174. 10.1038/nature12346.
27. Svara, F., Forster, D., Kubo, F., Januszewski, M., Dal Maschio, M., Schubert, P.J., Kornfeld, J., Wanner, A.A., Laurell, E., Denk, W., and Baier, H. (2022). Automated synapse-level reconstruction of neural circuits in the larval zebrafish brain. Nat. Methods *19*, 1357-1366. 10.1038/s41592-022-01621-0.
28. Briggman, K.L., Helmstaedter, M., and Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. Nature *471*, 183-188. 10.1038/nature09818.
29. Meirovitch, Y., Chandok, I.S., Park, C.F., Potocek, P., Mi, L., Sawmya, S., Li, Y., Athey, T.L., Susoy, V., Karlupia, N., et al. (2026). SmartEM: machine learning-guided electron microscopy. Nat. Methods *23*, 193-204. 10.1038/s41592-025-02929-3.
30. Jiang, Y., Wang, H., Boergens, K.M., Rzepka, N., Wang, F., and Hua, Y. (2025). Efficient cell-wide mapping of mitochondria in electron microscopic volumes using webKnossos. Cell Rep. Methods *5*, 100989. 10.1016/j.crmeth.2025.100989.
31. Koho, S., Tortarolo, G., Castello, M., Deguchi, T., Diaspro, A., and Vicidomini, G. (2019). Fourier ring correlation simplifies image restoration in fluorescence microscopy. Nat. Commun. *10*, 3103. 10.1038/s41467-019-11024-z.
32. B, P., S, N., and Sabaresan, V. (2024). Video frame interpolation using real-time intermediate flow estimation. 2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS).
33. Gambini, L., Gabbett, C., Doolan, L., Jones, L., Coleman, J.N., Gilligan, P., and Sanvito, S. (2024). Video frame interpolation neural network for 3D tomography across different

length scales. Nat. Commun. *15*, 7962. 10.1038/s41467-024-52260-2.
34. Zhou, W., Bovik, A.C., Sheikh, H.R., and Simoncelli, E.P. (2004). Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. *13*, 600-612. 10.1109/tip.2003.819861.
35. Li, D., Shi, X., Zhang, Y., Cheung, K.C., See, S., Wang, X., Qin, H., and Li, H. (2023). A Simple baseline for video restoration with grouped spatial-temporal shift. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

Acknowledgements

Funding:
the Scientific Research Instrument and Equipment Development Project of Chinese Academy of Sciences (PTYQ2025TD0002 to H.H.)
the National Natural Science Foundation of China (32171461 to H.H., 82171133 to Y.H.)
the Innovative Research Team of High-level Local Universities in Shanghai (SHSMU-ZLCX20211700).

Author contributions:
Conceptualization: YH, XC
Supervision: YH, XC, HH
Methodology: FW
Software: BC, YZ, HC
Data Curation: BC, YZ, HC
Formal Analysis: BC, HW, YH, XC
Project Administration: YH, XC
Funding Acquisition: YH, HH
Writing - Original Draft: YH, BC
Writing - Review & Editing: All authors

Competing interests:
All other authors declare they have no competing interests.

Data and materials availability:
The FIB-SEM mouse CA1 hippocampus dataset, including the raw image stacks and the corresponding mitochondrial segmentation annotations used in this study, is available from the website (https://www.epfl.ch/labs/cvlab/data/data-em/). Both the fast-scan and slow-scan SBF-SEM image stacks are available on the fileserver (https://zenodo.org/records/17082440). The source codes used in this paper are available at https://github.com/hmzawz2/OptimizedAcquisition4vEM.

Figures and Tables

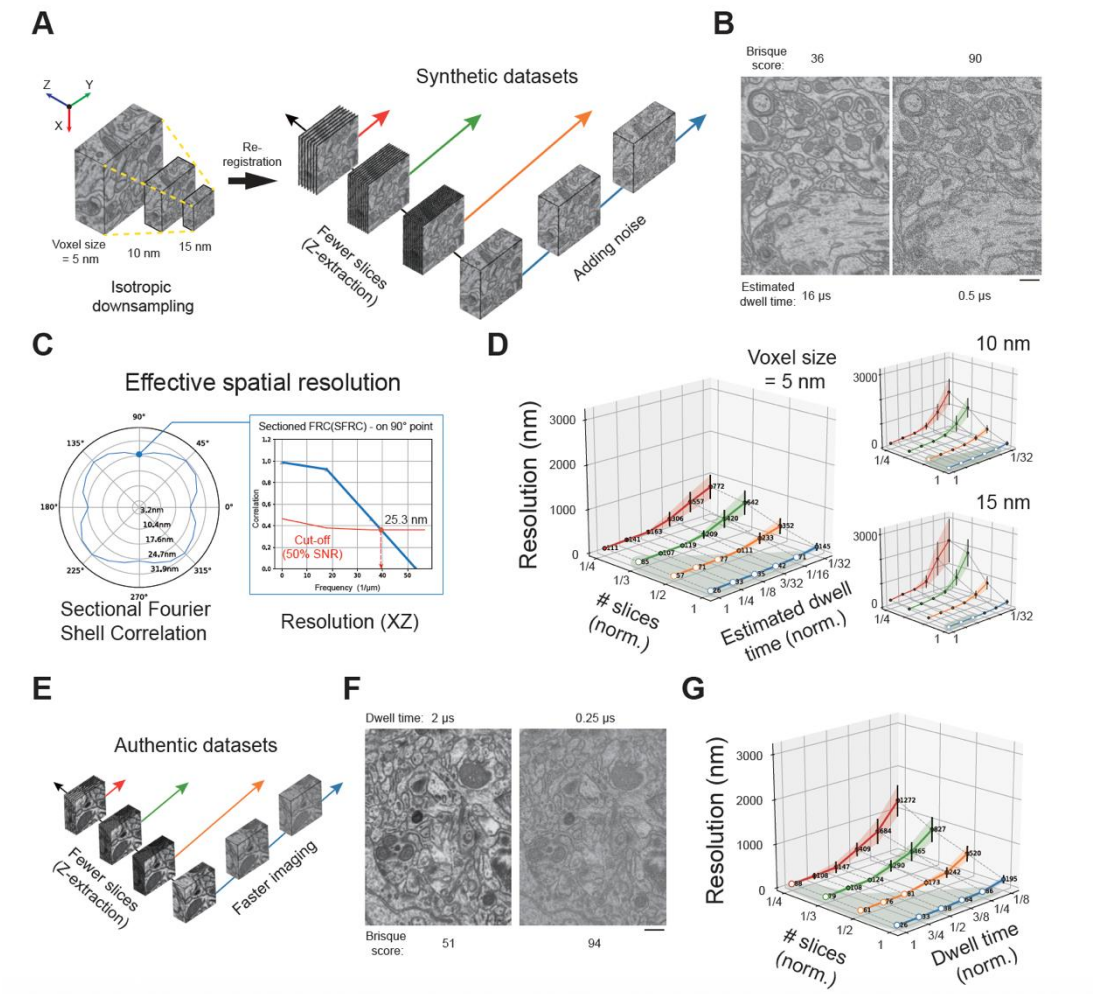

Fig. 1. Fidelity of structural information in different VEM samplings: (**A**) Data simulation procedure of a public FIB-SEM volume. Down-sampling yields three synthetic datasets with isotropic voxel sizes. The obtained datasets are progressively masked with noise of varying levels and extracted in the axial direction (Z). **(B)** To the original FIB-SEM images, Poisson-Gaussian noise was added with matched intensity in terms of the BRISQUE (Blind/Referenceless Image Spatial Quality Evaluator) scores reported for specific pixel dwell times (15). Example images masked by different noise levels are shown with corresponding BRISQUE scores, based on which the used dwell times are estimated. Scale bar: 500 nm. (**C**) Quantitative assessment of data quality using the effective spatial resolution, for which a sectional Fourier Shell Correlation (SFSC) is computed for each volume across different angles (as a polar plot). The resolution cut-off in the XZ plane (90° direction) is used to score changes in response to Z-extraction and noise masking. (**D**) Effective XZ resolutions of datasets under different Z-extraction and noise levels. Z-sampling rate (# slices) and noise mask (estimated dwell time) are normalized to the original imaging consumption and scaled approximately to the same fractions. Distinct rising phases of colored lines suggest that resolution worsening in datasets with low slice numbers is significantly accelerated upon heightened noise levels associated with fast scanning. Boundary conditions for the so-called “comfort zone” of acquisition (grey area), which is defined by an effective resolution below 100 nm (open dots).

Inset, a similar tendency is observed in modeling with datasets of down-sampled voxel sizes. (**E**) FIB-SEM validation data acquisition procedure of a mouse brain sample with different dwell times and extracted in the axial direction. (**F**) Representative images from the same region acquired at 2 µs/pixel (left) and 0.25 µs/pixel (right). The BRISQUE score of the slow-scanned image is 51 and for the fast-scanned one is 94, corresponding to different signal-to-noise ratios. Scale bar: 500 µm. (**G**) Effective XZ resolutions of validation datasets under different Z-extraction and imaging dwell times.

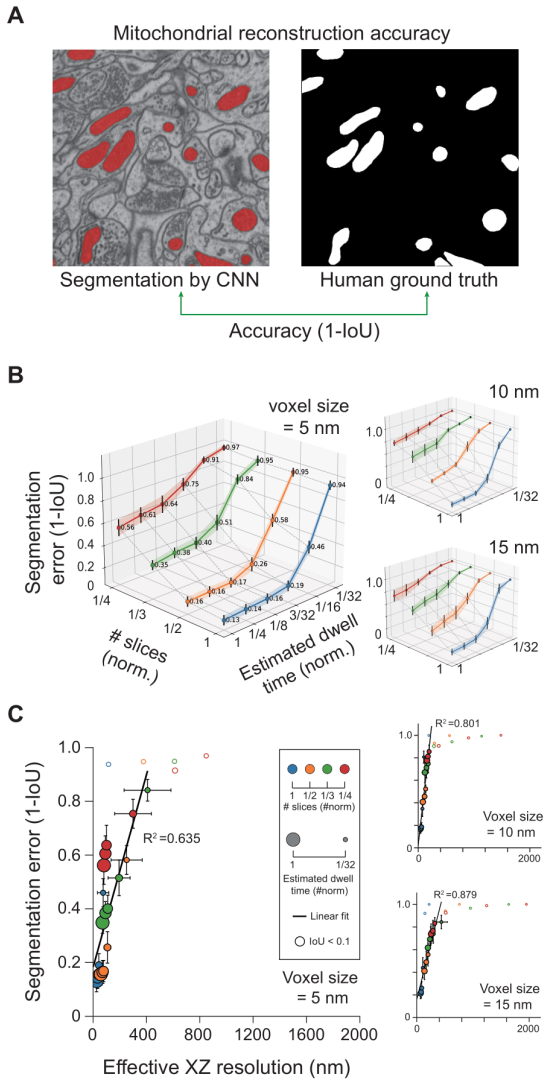

Fig. 2. VEM data quality influences the automated organelle segmentation: (**A**) Data quality assessment by segmentation accuracy of mitochondria using the convolutional neural network (CNN)-based model for mitochondria. The segmentation accuracy is scored using the Intersection-over-Union (IoU) between model inference and human annotation. **(B)** CNN segmentation error (as 1 - IoU) on synthetic datasets reveals declined model performance upon reduced Z-sampling and shortened dwell time. Inset: Similar scenarios for datasets with down-sampled voxel sizes. **(C)** Plot of CNN segmentation error (1-IoU) of mitochondria against the data effective resolution (data from Fig. 1). Linear fit, $R^2 > 0.635$ (outliers: IoU < 0.1). Inset: same plot with the datasets of 10 nm (top) and 15 nm (below) voxel size. These results demonstrate a strong positive correlation between the data's effective resolution and subsequent 3D reconstruction quality, validating the usage of the former as a general metric for data quality assessment.

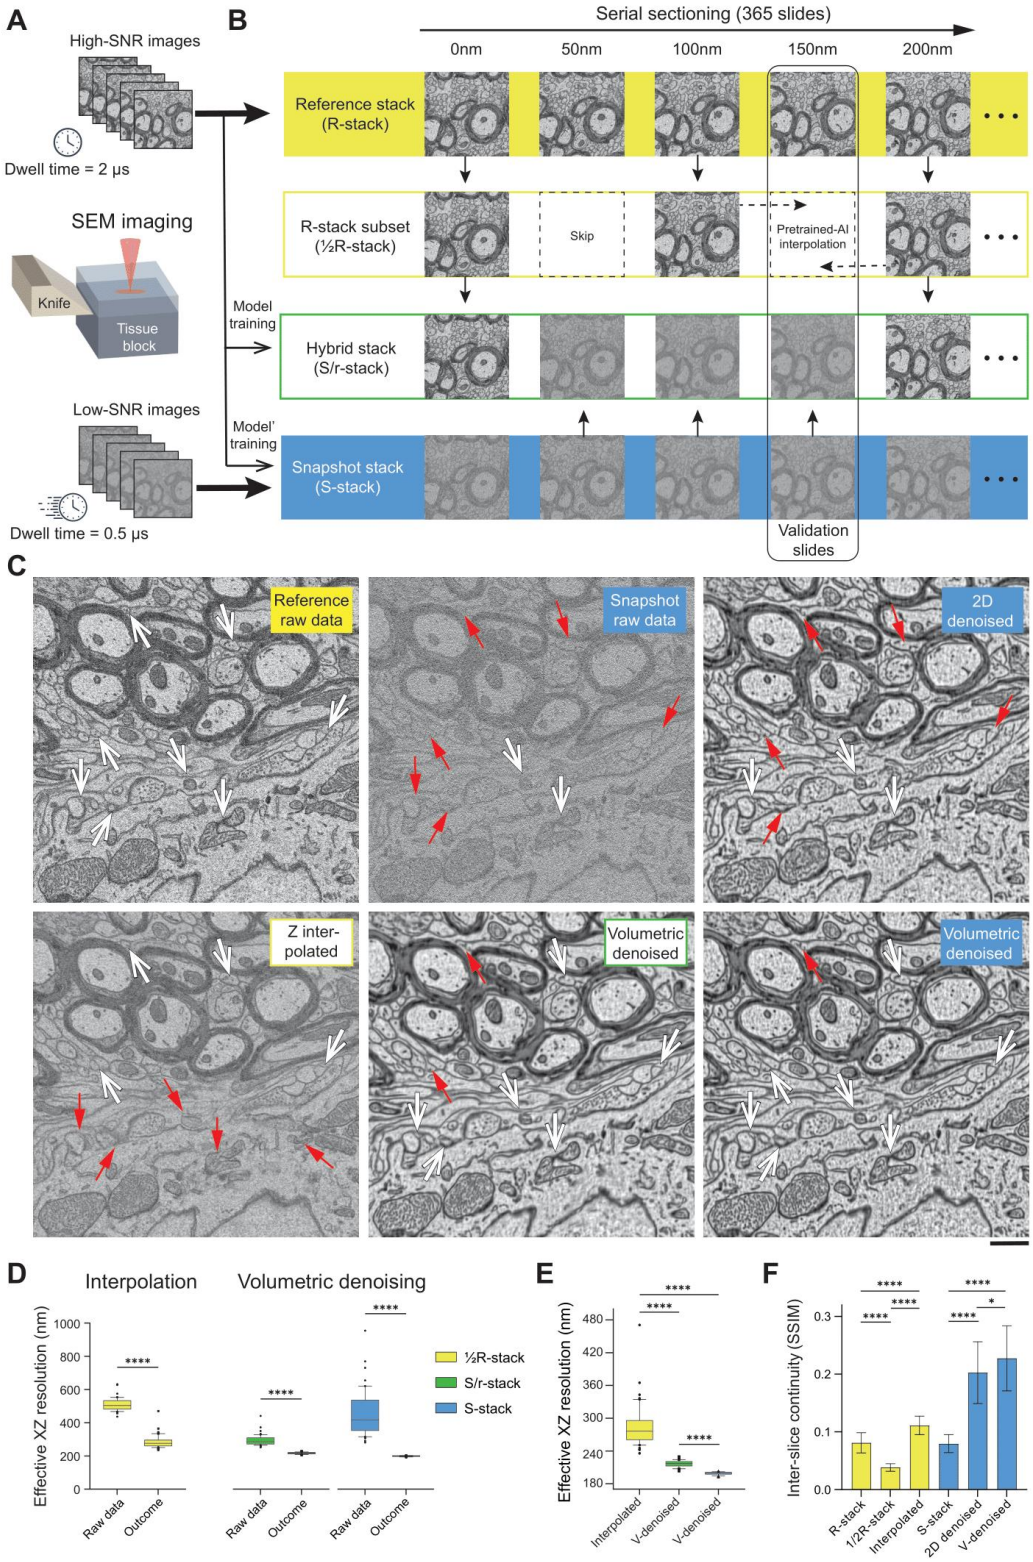
A
High-SNR images
Dwell time = 2 µs
SEM imaging
Knife
Tissue block
Low-SNR images
Dwell time = 0.5 µs
B
Serial sectioning (365 slides)
0nm
50nm
100nm
150nm
200nm
Reference stack (R-stack)
R-stack subset (½R-stack)
Skip
Pretrained-AI interpolation
Model training
Hybrid stack (S/r-stack)
Model' training
Snapshot stack (S-stack)
Validation slides
C
Reference raw data
Snapshot raw data
2D denoised
Z inter-polated
Volumetric denoised
Volumetric denoised
D
Interpolation
Volumetric denoising
Effective XZ resolution (nm)
½R-stack
S/r-stack
S-stack
Raw data
Outcome
E
Effective XZ resolution (nm)
Interpolated
V-denoised
V-denoised
F
Inter-slice continuity (SSIM)
R-stack
1/2R-stack
Interpolated
S-stack
2D denoised
V-denoised

Fig. 3. Denoising of image stack using 3D structural context: (**A**) Schematic illustration of a hybrid image acquisition using SBF-SEM. Repeat block-face scanning produces image pairs of low and high signal-to-noise-ratio (SNR). (**B**) Full stack of high-SNR images (R-stack, yellow solid) serves as a reference, whereas the snapshot dataset containing exclusively low-SNR images (S-stack, blue solid) is for the method development. Two synthetic image stacks: For the ½R-stack (yellow box), every other slow-scan-only slice is replaced with an interpolated virtual slice. S/r-stack (green box) is made up of fixed combinations of one reference and three snapshot slices. (**C**) Representative image of the same slice in the reference stack (up left) and snapshot dataset (up middle), as well as AI-denoised S-stack only using 2D information (up right). Similarly, the same slice was generated by 3D methods, including AI-based interpolation (bottom left) as well as supervised volumetric (V-) denoising on the S/r-stack (bottom middle) and the S-stack (bottom right). Arrows indicate faithfully preserved structures (white) and artifacts (red). Scale bar, 1000 nm. (**D**) Improvement of effective resolution by AI-based image restoration. (Left) By interpolation, the effective resolution for the ½R-stack drops from 510.0 ± 38.1 nm to 284.6 ± 39.2 nm. The V-denoising improves the effective resolution for the S/r-stack (middle) from 293.9 ± 32.6 nm to 216.5 ± 6.2 nm, as well as for the S-stack (right) from 458.0 ± 138.3 nm to 198.9 ± 2.7 nm. (**E**) Comparison of the achieved effective resolutions by different 3D approaches from (D). **(F)** Quantification of inter-slice structural continuity in datasets containing raw, interpolated, and denoised images. AI-interpolated ½R-stack achieved a higher score in Structural Similarity Index Measure (SSIM) than both the ½R-stack and the ground truth (R-stack). Compared to raw data (S-stack), both 2D and V-denoised datasets demonstrated significantly improved structural continuity, with the highest SSIM observed in the V-denoised dataset.

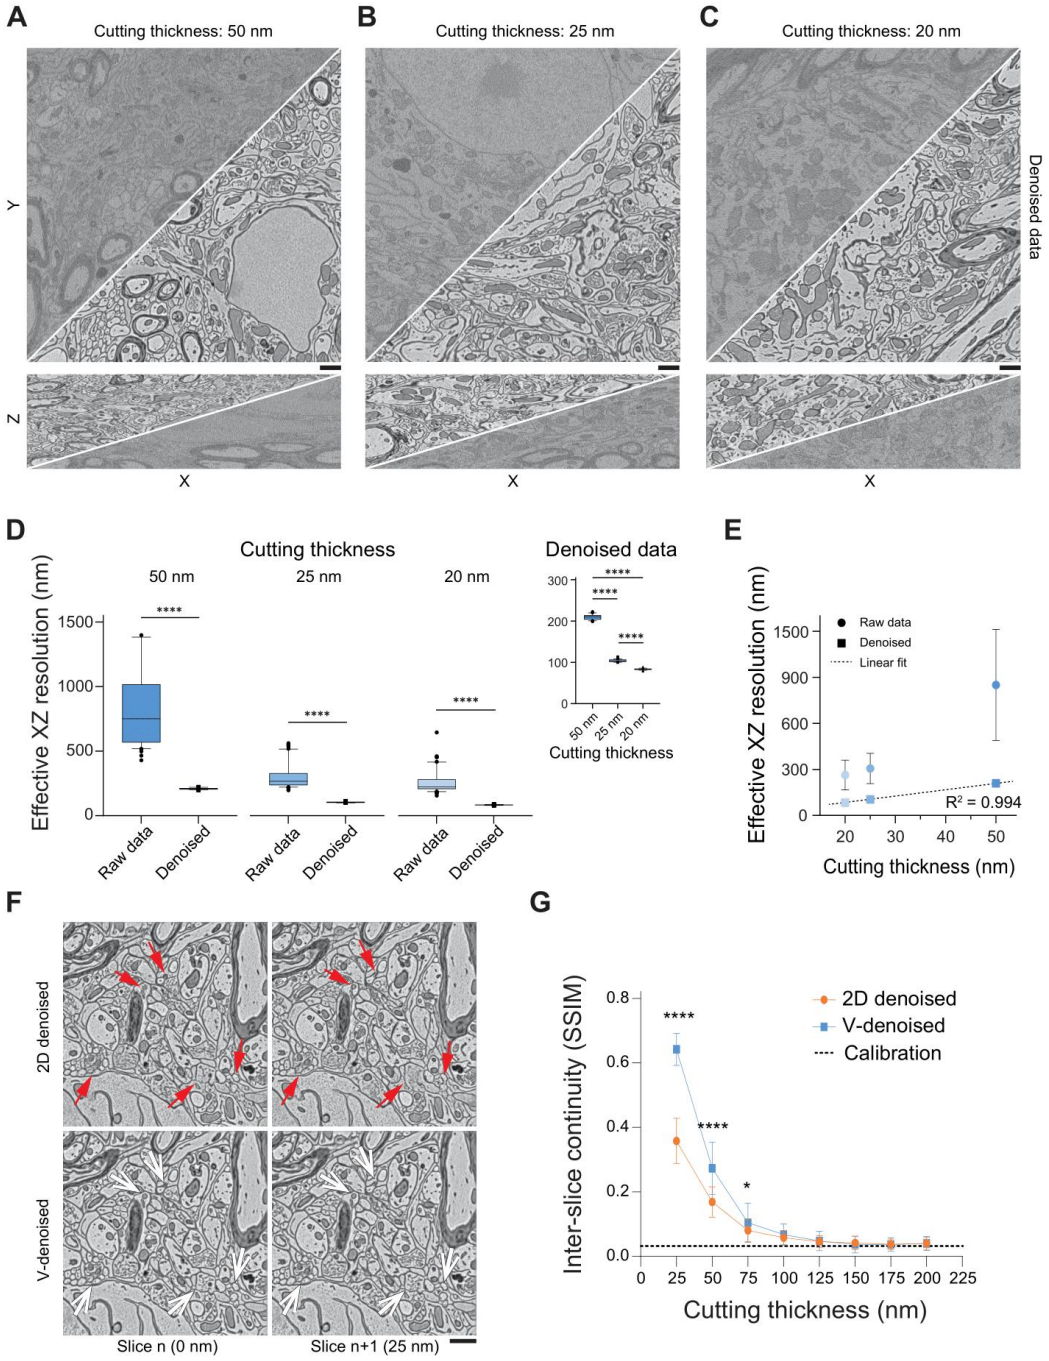

Fig. 4. Improved Z-sampling allowed by volumetric denoising: (**A-C**) Three SBF-SEM datasets acquired in the snapshot mode (0.5 µs/pixel) but at different cutting thicknesses (A: 50 nm, B: 25 nm, C: 20 nm) for the same sample. Representative raw image and volumetric denoised image by a self-supervised model are shown in both XY (top) and XZ (bottom) views. Scale bar, 1 µm. (**D**) Differential improvements in the effective XZ resolution after denoising are observed for datasets of different cutting thicknesses (50-nm: from 850.3 ± 361.8 nm to 209.5 ± 6.6 nm; 25nm: from 306.5 ± 98.9 nm to 104.1 ± 3.3 nm; 20-nm: from 263.6 ± 97.0 nm to 83.6 ± 2.0 nm). Inset: Plot of the achieved XZ

resolution from d. against physical cutting thickness. (**F**) 2D denoising (top row) and V-denoising (bottom row) were contact on the same dataset, and representative results of two consecutive slices (Slice n and Slice n+1) were shown. Upon a 25 nm cut, independent 2D denoising leads to more frequent sudden changes (red arrows), which are considered irrational and not observed in the V-denoised dataset (white arrows). Scale bar, 1 μm. (**G**) Measures of inter-slice structural continuity using SSIM as a function of cutting thickness. Datasets of increasing thicknesses (25 nm to 200 nm) were generated from the 25-nm dataset. Results of 2D denoising (blue circles) and V-denoising (orange squares) were compared. The advantage of V-denoising over 2D denoising in terms of preserving continuous structures is only observed in datasets with a slice thickness of less than 75 nm. Calibration of random structure continuity was carried out between two completely different patches within the volume (dotted line).