

Motion-corrected Noise2Noise denoising for structural analysis of cells and tissues with (cryo-)FIB/SEM

Vicente González-Ruiz^{1,*} and José-Jesús Fernández^{2,*}

¹Informatics Dept., CeinsaUAL, CINN-CSIC Associated Unit, Almería University, 04120 Almería, Spain

²Spanish National Research Council (CINN-CSIC). Asturias Health Research Institute (ISPA), 33011 Oviedo, Spain

*Corresponding author. vruiz@ual.es, jj.fernandez@csic.es

Abstract

Motivation: (Cryo-)FIB/SEM enables three-dimensional (3D) ultrastructural analysis of cells and tissues at nanoscale resolution. However, its characteristically low signal-to-noise ratio severely hinders interpretation of the volume data. Noise2Noise has emerged as a state-of-the-art denoising paradigm with outstanding performance in 3D electron microscopy -particularly cryo-ET- but its direct application to (cryo-)FIB/SEM remains unfeasible due to the inherent acquisition workflows.

Results: A method to enable Noise2Noise denoising in (cryo-)FIB/SEM is introduced. The input volume is split into even and odd slices followed by the application of motion correction methods to estimate and compensate for the structural differences between the two subvolumes. The resulting motion-corrected subvolumes are then used to train the Noise2Noise model, which is ultimately applied to denoise the original volume. Representative (cryo-)FIB/SEM volumes have been used to validate the method and to demonstrate the critical importance of the motion correction stage for volumes with a slice thickness greater than 10 nm.

Availability and Implementation: Code is available at <https://github.com/microscopy-processing/N2N-REO> and <https://doi.org/10.5281/zenodo.22303478>.

Contact: vruiz@ual.es, jj.fernandez@csic.es

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1. Introduction

FIB/SEM (Focused Ion Beam - Scanning Electron Microscopy) is a volume imaging technique with outstanding potential in cell biology owing to its unique ability to derive 3D structural information of cells and tissues with unlimited thickness at a few-nanometer intervals (Collinson *et al.*, 2023). Under cryogenic conditions, referred to as cryo-FIB/SEM, the structural preservation of the sample in close-to-native conditions is ensured (Schertel *et al.*, 2013; Dumoux *et al.*, 2023). The technique entails cycles of (a) FIB-milling of a thin layer (5-50 nm, so-called slice thickness) of the sample, followed by (b) SEM-imaging of the exposed surface, so as to collect stacks of images from several μm -thick volumes of the sample. The acquired data thus contain exquisite details of the cellular architecture at nanoscale resolution. The low signal-to-noise Ratio (SNR) inherent to this imaging technique, especially in cryo-FIB/SEM, poses significant challenges for visualization and interpretation, thereby making denoising an essential stage.

The standard denoising method in (cryo-)FIB/SEM is Gaussian filtering (Uwizeye *et al.*, 2021; Hintze *et al.*, 2024; Ch *et al.*, 2025; Velazco *et al.*, 2025) despite the overall blur it introduces. The application of more sophisticated methods that better preserve the structural features (e.g. non-local means, anisotropic nonlinear diffusion, deep-learning-based) still

remains scarce (Spehner *et al.*, 2020; Kapteijn *et al.*, 2022) because of their user-difficulty or computational complexity.

Noise2Noise (N2N) is a self-supervised denoising framework based on training a U-Net neural network using two independent noisy versions of a signal as both the input and the target; the model ultimately learns the expectation value, which represents the underlying noise-free signal (Lehtinen *et al.*, 2018). N2N has become the state-of-the-art in other 3D nano-imaging modalities, specifically cryo-electron tomography (cryo-ET), thanks to its striking performance (remarkable noise reduction, exceptional background flattening, excellent preservation of structural features) under extremely low contrast and low SNR conditions (Buchholz *et al.*, 2019a,b). However, (cryo-)FIB/SEM is not directly amenable to N2N because its destructive serial imaging workflows prevent the fulfillment of this central requirement (i.e. independent noisy observations of the same volume). To circumvent this situation, subsampling strategies of a single signal have been proposed to synthesize training pairs for N2N (Huang *et al.*, 2021; Lequyer *et al.*, 2022).

In this work, we introduce a N2N-based approach to single-volume denoising in (cryo-)FIB/SEM. Since the serially acquired cryo-FIB/SEM images contain independent noise, they are ideal for N2N. The proposed method thus splits the image stack into even and odd images to create the two required noisy versions of the volume. A key problem to address is that the




Figure 1 Motion-corrected Noise2Noise denoising for (cryo-)FIB/SEM. The volume is split into two subvolumes (even/odd), which are warped with motion correction to mitigate their differences due to the large slice thickness (even $\rightarrow$ odd/odd $\rightarrow$ even). Combinations of the subvolumes are then used to form training pairs ((even,odd $\rightarrow$ even) and (even $\rightarrow$ odd,odd), in turquoise and cyan arrows, respectively) to train a U-Net neural network. For the training, the subvolumes are actually split into 3D patches that are randomly selected to serve as the pairs input-target, and the network is updated according to a mean-squared-error-based loss function. Once trained, the network is applied to the original input volume (inference) to yield the denoised version.

structures within these two split subvolumes may differ non-negligibly, which is detrimental to N2N performance. This inter-volume difference is caused by the relatively large slice thickness in (cryo-)FIB/SEM (typically 10-50 nm). To overcome this issue, we apply motion correction methods to warp and align the structures, making the two subvolumes better-suited for effective N2N application.

2. Method

A self-supervised N2N approach to denoising single (cryo-)FIB/SEM volumes can utilize a subsampling procedure to produce training pairs (Huang et al., 2021; Lequyer et al., 2022). Therefore, given a noisy volume X , two subvolumes $X^{(\text{even})}$ and $X^{(\text{odd})}$ are generated from the even and odd slices, respectively. Because the images are acquired independently in the serial imaging workflow of (cryo-)FIB/SEM, these subvolumes provide two independent noisy observations of approximately the same clean signal. The U-Net model is trained with this pair, and the trained model is subsequently used during the inference process to denoise X and thus yield a denoised volume Y . We refer to this approach as N2N-EO.

Nevertheless, the previous procedure has several drawbacks compared to the ideal scenario involving two true noisy instances of the exact same clean signal, as occurs in cryo-ET (Buchholz et al., 2019a,b). First, as a consequence of the subsampling (which effectively doubles the slice thickness), there is a potential risk of reduced axial resolution and structural blurring during the inference. Second, because the clean underlying signal is not identical in both subvolumes, the model may learn from inaccurate spatial correspondences and also capture the aliasing effect produced by the subsampling, which in the end, is another source of blurring. In summary, the doubled slice thickness and the structural discrepancies between the two subvolumes may lead to blurred solutions in the N2N-EO approach, particularly for larger slice thickness values.

2.1. Motion-corrected N2N approach

To mitigate the spatial blurriness problem outlined above, we propose to estimate and compensate for the structural difference between the two subvolumes (even and odd) generated from the input noisy volume. To this end, we employ motion correction

methods to apply slice-wise, non-rigid registration of the subvolumes. Figure 1 sketches our proposed approach, termed N2N-REO (registered even/odd), for denoising (cryo-)FIB/SEM volumes, which is described in detail below.

To estimate the motion, we utilize dense optical flow (OF) estimation (Farnebäck, 2003; González-Ruiz et al., 2023). The OF characterizes the apparent motion of structures between the equivalent slices of the two subvolumes, $X_z^{(\text{even})}$ and $X_z^{(\text{odd})}$, where z denotes the slice index. The OF can be estimated assuming that the intensity of the pixels that represent those objects remains approximately constant:

$$X_z^{(\text{odd})}[y+v, x+u] \approx X_z^{(\text{even})}[y, x], \quad (1)$$

where each coordinate (y, x) maps an area of the image $X_z^{(\text{even})}$ to an area of the image $X_z^{(\text{odd})}$ through the corresponding displacement vector (v, u) . This allows us to build a projection to spatially bring the former to the latter image:

$$\mathbf{d}_z^{(\text{even} \rightarrow \text{odd})}(X_z^{(\text{even})}) = X_z^{(\text{even} \rightarrow \text{odd})} \approx X_z^{(\text{odd})}, \quad (2)$$

where $\mathbf{d}_z^{(\text{even} \rightarrow \text{odd})} = \{(v, u)\}$ is the displacement vector field. The same applies in the reverse direction:

$$X_z^{(\text{even})}[y+v, x+u] \approx X_z^{(\text{odd})}[y, x], \quad (3)$$

$$\mathbf{d}_z^{(\text{odd} \rightarrow \text{even})}(X_z^{(\text{odd})}) = X_z^{(\text{odd} \rightarrow \text{even})} \approx X_z^{(\text{even})}. \quad (4)$$

The OF problem therefore consists of estimating the displacement vector field $\mathbf{d}$ given the intensity constancy assumption (Eqs. 1, 3) and under the constraint that $\mathbf{d}$ must be smooth. Here we use the Farnebäck algorithm (Farnebäck, 2003), which minimizes the local Mean Square Error between the projection $X_z^{(\text{even} \rightarrow \text{odd})}$ and its corresponding reference slice $X_z^{(\text{odd})}$ (Eq. 2), and the projection $X_z^{(\text{odd} \rightarrow \text{even})}$ and its reference $X_z^{(\text{even})}$ (Eq. 4), using a sliding window that models the neighborhood with quadratic polynomials. The size w of this window controls the trade-off between estimation accuracy and robustness against noise. Another key parameter is the number of levels l of a Gaussian pyramid, which accounts for large displacements.

After OF estimation, the motion is corrected for by warping (i.e. non-rigidly registering) every slice of each subvolume

Table 1. Datasets.

ID	reference	sample	condition	instrument	pixel size	slice thickness
1	Knott et al. (2011) ¹	Brain tissue	chem.fix.	Zeiss NVision 40	5 nm	5 nm
2	Hennies et al. (2020) ²	HeLa cell	HPF/FS	Zeiss Crossbeam 540	5 nm	8 nm
3	Martin-Solana et al. (2024)	Brain tissue	HPF/FS	Zeiss Crossbeam 550	5 nm	15 nm
4	Kelley et al. (2025) ³	Chlamy. cell	cryo	Thermo Sci. Hydra V	5 nm	20 nm
5	Dumoux et al. (2023) ⁴	Brain tissue	cryo	Thermo Sci. Hydra V	6.7 nm	50 nm
6	Dumoux et al. (2023) ⁵	HeLa cell	cryo	Thermo Sci. Hydra V	6.7 nm	50 nm

¹ <https://www.epfl.ch/labs/cvlab/data/data-em>, ^{2–5} <https://www.ebi.ac.uk/EMPIAR>, IDs: 10311², 11275³, 11415⁴, 11419⁵.

towards its corresponding slice in the complementary subvolume, thereby generating the registered subvolumes:

$$\mathbf{X}^{(\text{even} \rightarrow \text{odd})} = \{\mathbf{X}_z^{(\text{even} \rightarrow \text{odd})}\} = \{\mathbf{d}_z^{(\text{even} \rightarrow \text{odd})}(\mathbf{X}_z^{(\text{even})})\}, \quad (5)$$

$$\mathbf{X}^{(\text{odd} \rightarrow \text{even})} = \{\mathbf{X}_z^{(\text{odd} \rightarrow \text{even})}\} = \{\mathbf{d}_z^{(\text{odd} \rightarrow \text{even})}(\mathbf{X}_z^{(\text{odd})})\}. \quad (6)$$

This enables the generation of registered input-target pairs with enhanced structural consistency while preserving noise independence: $(\mathbf{X}^{(\text{even})}, \mathbf{X}^{(\text{odd} \rightarrow \text{even})})$ and $(\mathbf{X}^{(\text{even} \rightarrow \text{odd})}, \mathbf{X}^{(\text{odd})})$ (indicated by the turquoise and cyan arrows, respectively, in Fig. 1). Similar to practices in cryo-ET ([Buchholz et al., 2019a,b](#)), the subvolumes are divided into 3D patches that are randomly sampled to form the input-target pairs used to train the U-Net architecture underlying the N2N denoising model (Fig. 1). Once trained, the model is applied to the original noisy volume $\mathbf{X}$, operating on a 3D patch-by-patch basis, to yield the final denoised volume $\mathbf{Y}$ (see the inference branch in Fig. 1).

3. Results

3.1. Datasets and parameters

A set of datasets representative of different samples (cells and tissues), preparation and imaging conditions (chemical fixation, high-pressure freezing and freeze substitution -HPF/FS-, cryo), instruments and slice thicknesses (5-50 nm) was selected to evaluate the method (Table 1).

The actual implementation of N2N-EO and N2N-REO relies on cryoCARE N2N ([Buchholz et al., 2019a](#)) and non-rigid registration using the Farnebäck optical flow algorithm available in OpenCV (pseudocode in Supplementary Material, Figs. S1 and S2). For the N2N training, the default parameters were employed (a U-Net of depth 3, kernel size 3, 16 initial filters, batch size 16 and 50 epochs) with cubic 3D patches of 16³ voxels. For the OF, the key parameters window size and number of levels were set to their optimal values of 5 and 3, respectively, as determined in a previous study ([González-Ruiz et al., 2023](#)).

3.2. Metrics for objective evaluation

Complementary to the visual analysis, we used a selected set of metrics for an objective evaluation to demonstrate the superior performance of N2N-REO. The first metric is based on the single-image Fourier Ring Correlation (FRC), also known as Self-FRC ([Koho et al., 2019](#); [Fernandez et al., 2020](#)). Here, every slice of the volume is decomposed into subimages of half

the original size by taking every other pixel along both spatial dimensions. The 2x-subsampled subimages are then subjected to FRC to yield the correlation of the slice content as a function of the spatial frequency. This metric accounts for both SNR and blurring (lower SNR or greater blur results in a lower FRC at the corresponding spatial frequency), thereby enabling the assessment of the effective resolution of each slice. In this work, the average Self-FRC computed over all slices is utilized as the final metric, referred to as ASFRC.

Furthermore, we also incorporated well-known similarity metrics from the computer vision field: the Structural Similarity Index Measure (SSIM), the Pearson's Correlation Coefficient (PCC) and the Peak SNR (PSNR) ([Szeliski, 2022](#); [Wang et al., 2004](#)). These metrics were computed between each denoised volume and its corresponding original noisy volume to assess the preservation of the original signal upon denoising (with larger values indicating better preservation). The combination of these metrics with ASFRC aims to evaluate the robustness to blurring and the structure-preserving capability of our method.

3.3. Influence of the slice thickness

To study the effect of the slice thickness on the performance of the denoising methods, a series of datasets with slice thicknesses of 5, 10, 20, 30, 40 nm was synthesized from Dataset 1 (slice thickness 5 nm) by omitting the corresponding number of intermediate slices. These generated datasets were then subjected to denoising using both N2N-EO and N2N-REO. Figure 2 shows a selected slice present in all denoised volumes (the 5 nm case not shown for conciseness). N2N-REO (Fig. 2B) exhibits noise reduction with good structure preservation regardless of the slice thickness. In contrast, N2N-EO shows a clear susceptibility to blurring, particularly severe at 30 and 40 nm (Fig. 2C and Suppl. Fig. S3). This behaviour is especially noticeable in mitochondria cristae, vesicles and membranous areas in general. Therefore, as indicated by the divergence between N2N-REO and N2N-EO, motion correction becomes paramount beyond 10 nm, growing increasingly critical as the thickness increases.

To further substantiate the performance of the methods with the slice thickness, we performed a quantitative analysis by computing objective metrics from the volumes. First, ASFRC curves were obtained for all volumes to estimate the overall cleanliness and structural sharpness. Fig. 3 shows that N2N-REO consistently outperforms N2N-EO at all spatial frequencies, with the performance gap becoming particularly pronounced for slice thicknesses greater than 10 nm. Furthermore, N2N-EO




Figure 2 Performance of motion-corrected N2N as a function of slice thickness. (A) Original volume. The dashed box is zoomed-in in the bottom panel. Some areas of interest (membranes, vesicles, mitochondria cristae) are marked with arrows. (B) Results of N2N-REO (motion correction) for volumes with slice thickness 10 to 40 nm (left to right). (C) Equivalent results of N2N-EO (no motion correction). Bar: 200 nm.




Figure 3 Average Self Fourier Ring Correlation (ASFRC) as a function of slice thickness. From left to right, ASFRC curves for slice thickness of 10, 20, 30 and 40 nm. Colors: original noisy volume (blue), N2N-EO (red), N2N-REO (green). At X axis, the maximum normalized spatial frequency (0.5) corresponds to that from the 2x-subsampled subimages used for the ASFRC calculation.

degrades significantly at large values (30, 40 nm). These results indicate that N2N-REO provides denoised volumes that retain information at higher resolution. Second, SSIM, PCC and PSNR metrics were calculated to assess the preservation of the original signal (Suppl. Table S1). The ratio N2N-REO/N2N-EO for all metrics (Fig. 4) confirms the superiority of N2N-REO, exhibiting a consistent upward trend for slice thicknesses exceeding 10 nm.

To provide an additional independent evaluation, we conducted an experiment using synthetic phantoms with a clean reference and quantitative metrics. The results align with the findings in this section and confirm the superiority of N2N-REO (see Supplementary Material, Section S3).

N2N relies on the strict assumption that the noise in the training pairs is independent. We conducted an experiment confirming that the warping applied in N2N-REO does not violate this requirement (Supplementary Material, Section S4).

3.4. Evaluation with experimental data

We then applied the methods to the experimental datasets listed in Table 1. The results for Dataset 1 (slice thickness 5 nm) were




Figure 4 Metrics PSNR, SSIM and PCC as a function of slice thickness. The metrics are presented as a ratio N2N-REO/N2N-EO.

visually indistinguishable from those presented in Fig. 2 (panels labelled as 10 nm), showing a sharp signal on a clean volume, with no discernible difference between N2N-REO and N2N-EO.

Fig. 5 presents the results for the two datasets prepared with HPF/FS. In both cases, the structural signal is preserved, noise




Figure 5 Denoising FIB/SEM datasets prepared with HPF/FS. From left to right, a slice of the original volume, denoised with N2N-REO (motion correction) and with N2N-EO (no motion correction). (A) HeLa cell (Dataset 2). (B) Brain tissue (Dataset 3). Some areas with more evident effects are marked (arrows). Bar: 200 nm.

is effectively suppressed, and, particularly in the second case, the background is substantially flattened. In accordance with the findings in the previous section, N2N-REO and N2N-EO display negligible differences for Dataset 2 (slice thickness 8 nm) whereas for Dataset 3 (15 nm) N2N-REO outperforms N2N-EO, as the latter exhibits blurring that is particularly noticeable in membranous regions.

Finally, we evaluated the methods with the three datasets under cryo-conditions. Cryo-FIB/SEM of vitrified samples is in general characterized by extremely low contrast, low SNR conditions and slice thicknesses typically exceeding 20 nm, thus representing challenging case studies. As presented in Fig. 6 and Suppl. Figs. S4 and S5, N2N-REO (motion correction) succeeds in significantly reducing the noise and flattening the background while demonstrating outstanding abilities to preserve sharpness in the structures. In contrast, N2N-EO (no motion correction) undergoes substantial overall blurring, leaving certain membranous constituents (e.g. Golgi stacked cisternae, mitochondria cristae) barely perceptible in some areas. These results are consistent with the large slice thickness inherent to cryo-FIB/SEM and corroborate the paramount importance of motion correction in N2N denoising.

To quantitatively evaluate the quality of the denoised volumes, the objective metrics were calculated. The ASFRC curves (Fig. 7) show that N2N-REO is superior to N2N-EO at all spatial frequencies in all cases. ASFRC thus confirms that N2N-REO yields denoised volumes with details of higher resolution and avoids blurring compared to N2N-EO. Consistent with the previous section, the performance gap in ASFRC between the two approaches becomes apparent beyond 10 nm and widens for larger slice thicknesses. Additionally, the SSIM, PCC and PSNR metrics (Suppl. Table S2) for N2N-REO are consistently higher

than those for N2N-EO, further demonstrating its superior ability to preserve original structural details.

The mismatch in slice spacing between the training and inference stages (even/odd subvolumes with double spacing versus the original volume, respectively) might, in principle, introduce structural blurring and reduced axial resolution in N2N-REO. Nevertheless, our results suggest that these effects are not apparent, thereby validating the N2N-REO approach for structure-preserving denoising in cryo-FIB/SEM.

3.5. Comparison with state-of-the-art methods

Next, we aimed to compare our motion-corrected N2N approach with respect to standard and state-of-the-art denoising methods in (cryo-)FIB/SEM. To this end, the challenging Cryo-dataset 4 was selected and subjected to standard Gaussian filtering (GF), Structure-Preserving Gaussian denoising (SPGD) (González-Ruiz *et al.*, 2023; Zhu *et al.*, 2024) and Anisotropic Nonlinear Diffusion (AND) (Moreno *et al.*, 2018; Kapteijn *et al.*, 2022). Fig. 8 presents the results for areas containing Golgi complexes, showing that all methods succeed in reducing the noise to a similar level. Nevertheless, the drawbacks of GF, SPGD and AND are evident (see areas denoted by arrows). GF blurs the structures while SPGD and AND -which are known to reinforce membranous structures- are prone to creating artefactual patterns and still introduce some blur in specific areas. Overall, N2N-REO manages to reduce noise while naturally preserving structural features without generating artefacts.




Figure 6 Denoising datasets from vitrified samples acquired with cryo-FIB/SEM. From left to right, a slice of the original volume, denoised with N2N-REO (motion correction) and with N2N-EO (no motion correction). (A) *Chlamydomonas* cell (Dataset 4). (B) Brain tissue (Dataset 5). (C) HeLa cell (Dataset 6). Some membranous areas exhibiting more evident effects are marked (arrows). Bar: 200 nm.




Figure 7 Average Self Fourier Ring Correlation (ASFRC) for the six datasets in Table 1. From left to right, the slice thickness of the datasets is in increasing order: 5, 8, 15, 20, 50 and 50 nm. Colors: original noisy volume (blue), N2N-EO (red), N2N-REO (green). The maximum value (0.5) on X axis represents the maximum spatial frequency corresponding to the 2x-subsampled subimages used for the ASFRC calculation.

3.6. Application to Volume Electron Microscopy (vEM)

(Cryo-)FIB/SEM represents a prominent modality within a suite of electron microscopy techniques, collectively known as vEM, which enable the in-depth investigation of the 3D structure of cells and tissues (Collinson *et al.*, 2023). Most vEM techniques rely on serial sectioning and imaging of the sample. Consequently, they could also benefit from our motion-corrected N2N denoising approach. To demonstrate this broader

applicability, we applied N2N-REO to selected datasets from two vEM modalities. Fig. 9 and Suppl. Fig. S9 illustrate the outstanding noise reduction, background flattening and signal preservation achieved in both cases. Therefore, our approach can be used to facilitate interpretation of the volumes in vEM.

4. Conclusion

We have introduced a self-supervised Noise2Noise approach to denoising single volumes in (cryo-)FIB/SEM. It splits the




Figure 8 Comparison with other methods used in (cryo-)FIB/SEM. Selected areas of cryo-dataset 4. From left to right, slice from the original volume and denoised with N2N-REO, Gaussian filtering, Structure-Preserving Gaussian denoising and Anisotropic Nonlinear Diffusion. Areas susceptible to artefacts or blurring during denoising are marked with arrows. Bar: 200 nm.




Figure 9 Motion-corrected N2N denoising applied to Array Tomography of chemically fixed cells (Nakamura *et al.*, 2025) (EMPIAR, ID: 12608) with slice thickness of 50 nm. Bar: 500 nm.

input volume into two subvolumes containing even/odd slices to conform the training pairs. The method includes a key motion-correction stage just prior to training, aiming at estimating and compensating for the structural discrepancies between the two subvolumes. Experimental results on (cryo-)FIB/SEM datasets demonstrate that the motion correction significantly improves the noise attenuation, structural preservation and overall quality of the volumes, particularly for slice thicknesses larger than 10 nm. These findings suggest that data splitting combined with non-rigid registration provides an effective solution for self-supervised denoising in volume electron microscopy scenarios where repeated acquisitions of the same volume are unfeasible.

Funding

This work was supported by MICIU/AEI/ 10.13039/501100011033 and ERDF/EU [grants PID2022-139071NB-I00, PID2024-155521OB-I00].

References

- Buchholz, T. *et al.* (2019a). Content-aware image restoration for electron microscopy. *Methods Cell Biol*, **152**, 277–289.
- Buchholz, T. *et al.* (2019b). Cryo-care: Content-aware image restoration for cryo-transmission electron microscopy data. In *2019 IEEE Intl Symp Biomedical Imaging*, pages 502–506.
- Ch, S. *et al.* (2025). A comprehensive overview of FIB-SEM applications for the evaluation of outer retina. *Front Cell Dev Biol*, **13**, 1586029.
- Collinson, L. *et al.* (2023). Volume EM: a quiet revolution takes shape. *Nat Methods*, **20**, 777–782.
- Dumoux, M. *et al.* (2023). Cryo-plasma FIB/SEM volume imaging of biological specimens. *eLife*, **12**, e83623.
- Farnebäck, G. (2003). Two-frame motion estimation based on polynomial expansion. *LNCS*, **2749**, 363–370.
- Fernandez, J. J. *et al.* (2020). PolishEM: image enhancement in FIB-SEM. *Bioinformatics*, **36**, 3947–3948.
- González-Ruiz, V. *et al.* (2023). Structure-preserving gaussian denoising of FIB-SEM volumes. *Ultramicroscopy*, **246**, 113674.
- Hennies, J. *et al.* (2020). AMST: alignment to median smoothed template for focused ion beam scanning electron microscopy image stacks. *Sci Rep*, **10**, 2004.
- Hintze, A. *et al.* (2024). Developmental changes of the mitochondria in the murine anteroventral cochlear nucleus. *iScience*, **27**, 108700.
- Huang, T. *et al.* (2021). Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proc IEEE/CVF conf computer vision and pattern recognition*, pages 14781–14790.
- Kapteijn, R. *et al.* (2022). Endocytosis-like DNA uptake by cell wall-deficient bacteria. *Nat Commun*, **13**, 5524.
- Kelley, R. *et al.* (2025). Toward community-driven visual proteomics with large-scale cryo-electron tomography of *Chlamydomonas reinhardtii*. *Mol Cell*, **86**, 213–230.
- Knott, G. *et al.* (2011). Focussed ion beam milling and scanning electron microscopy of brain tissue. *JoVE*, **53**, e2588.

- Koho, S. *et al.* (2019). Fourier ring correlation simplifies image restoration in fluorescence microscopy. *Nat Commun*, **10**, 3103.
- Lehtinen, J. *et al.* (2018). Noise2noise: Learning image restoration without clean data. In *Proc. 35th Intl Conf Machine Learning*, volume 80, pages 2965–2974.
- Lequyer, J. *et al.* (2022). A fast blind zero-shot denoiser. *Nat Mach Intell*, **4**, 953–963.
- Martin-Solana, E. *et al.* (2024). Disruption of the mitochondrial network in a mouse model of Huntington's disease visualized by in-tissue multiscale 3D electron microscopy. *Acta Neuropathol Commun*, **12**, 88.
- Moreno, J. J. *et al.* (2018). TomoEED: fast edge-enhancing denoising of tomographic volumes. *Bioinformatics*, **34**, 3776–3778.
- Nakamura, K. *et al.* (2025). Mitochondrial complexity is regulated at er-mitochondria contact sites via pdzd8-fkbp8 tethering. *Nat Commun*, **16**, 3401.
- Schertel, A. *et al.* (2013). Cryo FIB-SEM: Volume imaging of cellular ultrastructure in native frozen specimens. *J Struct Biol*, **184**, 355–360.
- Spehner, D. *et al.* (2020). Cryo-FIB-SEM as a promising tool for localizing proteins in 3D. *J Struct Biol*, **211**, 107528.
- Szeliski, R. (2022). *Computer Vision*. Springer.
- Uwizeye, C. *et al.* (2021). Morphological bases of phytoplankton energy management and physiological responses unveiled by 3D subcellular imaging. *Nat Commun*, **12**, 1049.
- Velazco, A. *et al.* (2025). Reduction of SEM charging artefacts in native cryogenic biological samples. *Nat Commun*, **16**, 5204.
- Wang, Z. *et al.* (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Trans Image Process*, **13**, 600–612.
- Zhu, H. *et al.* (2024). Salmonella exploits membrane reservoirs for invasion of host cells. *Nat Commun*, **15**, 3120.