



28 **ABSTRACT**

29 **Cryogenic electron tomography (cryo-ET) has rapidly advanced as a high-resolution**  
30 **imaging tool for visualizing subcellular structures in 3D with molecular detail. Direct image**  
31 **inspection remains challenging due to inherent low signal-to-noise ratios (SNR). We**  
32 **introduce CryoSamba, a self-supervised deep learning-based model designed for**  
33 **denoising cryo-ET images. CryoSamba enhances single consecutive 2D planes in**  
34 **tomograms by averaging motion-compensated nearby planes through deep learning**  
35 **interpolation, effectively mimicking increased exposure. This approach amplifies coherent**  
36 **signals and reduces high-frequency noise, substantially improving tomogram contrast**  
37 **and SNR. CryoSamba operates on 3D volumes without needing pre-recorded images,**  
38 **synthetic data, labels or annotations, noise models, or paired volumes. CryoSamba**  
39 **suppresses high-frequency information less aggressively than do existing cryo-ET**  
40 **denoising methods, while retaining real information, as shown both by visual inspection**  
41 **and by Fourier shell correlation analysis of icosahedrally symmetric virus particles. Thus,**  
42 **CryoSamba enhances the analytical pipeline for direct 3D tomogram visual interpretation.**

## 43 INTRODUCTION

44 Cryogenic electron tomography (cryo-ET) has become an important tool in structural biology for  
45 imaging three-dimensional biological structures with molecular resolution in their native context  
46 (Baumeister et al., 1999; Medalia et al., 2002). Achieving this dual capability requires an extremely  
47 low electron dose per tilt image to avoid sample damage, resulting in data with very low signal-  
48 to-noise ratios (SNR) (Gan et al., 2012). Enhancement of signal through downstream processing  
49 can be achieved by aligning and merging multiple instances of an invariant biological structure  
50 (when one exists), an approach known as subtomogram averaging (STA) (Wan et al., 2016). STA  
51 has effectively produced high-resolution electron density maps of viruses (Schur et al., 2016),  
52 ribosomes (Erdmann et al., 2021), and nuclear pores (Mosalaganti et al., 2022).

53

54 Traditionally, contrast in cryoET volumes has been enhanced using low pass filtering and pixel  
55 binning (Lučić et al., 2005). Recently, deep learning approaches have emerged as superior  
56 alternatives (Buchholz et al., 2019; Bepler et al., 2020). These methods adapt to the intricacies  
57 of the data, but because cryo-ET data generally lack ground truth high SNR images for direct  
58 supervision, most deep-learning denoising algorithms rely on self-supervision (Lehtinen et al.,  
59 2018), using paired 3D volumes from an even/odd split of the cryo-ET tilt-series (Buchholz et al.,  
60 2019; Bepler et al., 2020), synthetic or annotated data (Zeng et al., 2024), or noise modeling (Li  
61 et al., 2022).

62

63 Despite their effectiveness in enhancing SNR, these approaches inevitably distort the original  
64 data, particularly by suppressing high-frequency details (Bepler et al., 2020). Consequently,  
65 denoised tomograms are typically reserved for interpretability tasks, such as identifying the  
66 regions of interest for STA or other downstream processing, which then uses the original raw  
67 data. These distortions can impede interpretability if high spatial frequency details are essential  
68 for distinguishing objects of interest. Therefore, it is desirable for cryo-ET pipelines to incorporate  
69 denoising methods that minimize such deformations.

70

71 We have developed a deep learning-based denoising method for cryo-ET that enhances contrast  
72 with minimal deformation when compared to other current techniques. CryoSamba, the software  
73 that applies this approach, operates in a fully self-supervised manner, training directly on the raw  
74 three-dimensional volume without requiring additional data such as paired volumes or  
75 simulations. CryoSamba is very efficient, with only three million parameters, making it feasible to  
76 run even on GPU-equipped current laptops.

77

78 We demonstrate CryoSamba's efficacy on five cryoET datasets with three distinct voxel  
79 resolutions (achieved by pixel binning). CryoSamba substantially increases SNR for all voxel  
80 resolutions, verified both visually and quantitatively. Analysis of spatial frequencies in Fourier  
81 space confirms that our approach suppresses higher frequencies less severely than do current  
82 methods. We benchmarked CryoSamba's performance by evaluating the Fourier Shell  
83 Correlation (FSC) (Harauz et al., 1986) for subtomogram-averaged virus-particle images, starting  
84 with tomograms before and after denoising. Comparison with the known virus structure showed  
85 that CryoSamba denoising preserved higher resolution information more faithfully than did Topaz-  
86 Denoise or CryoCARE, two widely used denoising methods.

## 87 **RESULTS**

### 88 **Cryo-ET data sets**

89 To assess the denoising capabilities of CryoSamba, we employed five distinct tomograms derived  
90 from cryo-ET of various biological samples. Two of these tomograms were obtained from the  
91 edges of plunge-frozen human BSC1 cells grown overnight on top of the electron microscopy  
92 grids, showcasing cross-sections of the plasma membrane, mitochondria, and an early  
93 endosome. These images also contained numerous free ribosomes and actin filaments within the  
94 cytosol, as well as rhesus rotavirus particles in the surrounding medium, from a study of the initial  
95 stages of virus entry (Herrmann et al., 2021; de Sautu et al., 2024).

96

97 The remaining three tomograms came from lamellae prepared by cryo-focused ion beam (cryo-  
98 FIB) milling of plunge-frozen yeast cells. These samples showed cross-sections of mitochondrial  
99 and endoplasmic reticulum (ER) membranes, ribosomes, either free in the cytosol or attached to  
100 the ER cytosolic face, and cross-sections of the double-membrane nuclear envelope. They also  
101 included some actin filaments in the cytosol.

102

103 The 3D tomographic reconstructions were produced from tilted images recorded at a nominal  
104 (unbinned) pixel size of 2.62 Å. Before denoising, we applied 3D contrast transfer function (CTF)  
105 correction using NovaCTF (Turoňová et al., 2017) to mitigate potential confounding effects from  
106 defocus. The tomograms used in our denoising tests were derived from data at various binning  
107 levels: unbinned (2.62 Å/pixel), 3x binned (7.86 Å/pixel), or 6x binned (15.72 Å/pixel), and some  
108 tests used tomograms generated from the even or odd frames of tilt series.

109

### 110 **Deep learning CryoSamba training pipeline**

111 The general strategy is shown schematically in Fig. 1A. It repurposes the deep learning model  
112 Enhanced Bi-Directional Motion Estimation (EBME) (Jin et al., 2023), initially designed to enhance  
113 the frame rate of 2D videos through synthetic video frame interpolation. We treat our tomograms  
114 as “videos”, converting the z spatial direction in the tomogram into the time dimension of EBME.  
115 EBME then generates a series of “motion-compensated” xy planes. For each plane, the model  
116 generates a set of “best guesses” from pairs of planes equally spaced to either side of the xy  
117 plane in question. It then takes the average of these best guesses as the new estimate of that  
118 plane. This averaging process produces a denoising effect like that obtained by an increase in  
119 imaging exposure time (Mildenhall et al., 2018).

120

121 To train the EBME model, we selected three sequential xy planes from a tomogram, spaced  
122 equally along the z-axis at positions  $z-L$ ,  $z$ , and  $z+L$  (Fig. 1A). The outer planes were input into a  
123 neural network, depicted by the flow module in Fig. 1B, which computed two deformation fields to  
124 morph these planes towards the middle one. The transformed planes were merged using a U-Net  
125 (Ronneberger et al., 2015) represented by the fusion module in Fig. 1B, creating an interpolated  
126 copy of the central plane. We then minimized a loss function—reflecting the disparity between  
127 this interpolated plane and the original—by gradient descent and backpropagation (LeCun et al.,  
128 2015). Training proceeded by using plane triplets across all  $z$  values and spacings  $L$ , from  $L=1$   
129 up to  $L=L_{\max}$ , and concluded when the loss stabilized. This training process effectively reduced  
130 high-frequency noise uncorrelated across planes, such as Gaussian and shot noise; these  
131 sources of noise were further damped by averaging the interpolated images for all  $L$  spacings  
132 corresponding to the same  $z$ .

133

### 134 **CryoSamba denoising pipeline**

135 The general strategy of the denoising steps is represented schematically in Fig. 1C. The inference  
136 phase of CryoSamba initiates upon completion of the training step carried out with the tomogram  
137 being denoised. For any given  $z$  in the tomogram, we input adjacent xy planes at  $z-L$  and  $z+L$  into  
138 the EBME model, to generate a denoised version at  $z$ . This process is repeated for a given  $z$   
139 plane by varying  $L$  from 1 to  $L_{\max}$  and averaging them to create a final xy plane at the position  $z$ .  
140 In our tests, visual inspection was consistent with enhanced denoising. The interpolation quality  
141 decreased as  $L$  increased, leading to a tradeoff between denoising strength and blurring of fine  
142 details; an optimal image was achieved by visual ‘fine-tuning’ the value of  $L_{\max}$ . Fig. 1 shows an  
143 example of the effect of this process in the appearance of rotavirus-particle cross section. The  
144 series of denoised images were obtained using different  $L_{\max}$  values (Fig. S1) with best chosen  
145 images obtained with  $L_{\max}$  of 3 (6), 6 (12) and 10 (20) for training (inference) and voxel resolutions  
146 of 15.72Å, 7.86Å and 2.62Å, respectively.

147

148 We minimize loss of information introduced by the averaging process during the inference phase  
149 (Fig. S2) by using the "one step back, one step forward" strategy outlined in Fig. 1D. This strategy,  
150 like Samba dance steps (Guillermoprieto, 1991), involves generating synthetic xy planes at  $z-L/2$   
151 and  $z+L/2$ , between the experimentally determined  $z-L$  and  $z$ , and between  $z$  and  $z+L$ ,  
152 respectively, which are then used as additional input during inference. By restricting use of the  
153 "one step back, one step forward" approach to the inference phase, we prevent a potential model

154 collapse during training, that could have led for example to trivial solutions such as simply copying  
155 input frames (Reda et al., 2019).

156

157 The inference phase for a given  $z$  finishes by averaging the interpolated planes associated with it  
158 and using the average to replace the data in the original  $xy$  plane (Fig. 1E); this process is  
159 repeated for all  $z$  positions, ultimately yielding a uniformly denoised 3D volume, as schematically  
160 illustrated in Fig. 1C.

161

162 We note that in the CryoSamba denoising strategy, we train separately using each tomogram we  
163 wish to denoise. We found that the results using this approach were more reliable than those from  
164 one in which we performed a single training with many tomograms and then used the trained  
165 model to denoise naïve tomograms not included in the training pool. The fully self-supervised  
166 character of CryoSamba makes this approach possible, since it only requires for training the same  
167 3D volume that one then wishes to denoise. In practical terms, the underlying deep learning model  
168 is relatively light (approximately three million parameters), and the training times are quite short.

169

170 To illustrate the effectiveness of CryoSamba, we visually compared tomograms before and after  
171 denoising. These tomograms were generated with 3x binning (7.86 Å/pixel) from yeast (Fig. 2A-  
172 C) and mammalian BSC-1 cells (Fig. 2D). The denoised images show enhanced SNR, estimated  
173 as describe below, in a single  $xy$  plane, approximately at the tomogram midsection, orthogonal to  
174 the electron beam direction ( $z$ ). Inspection of sequential planes in the 3D volumes from yeast  
175 samples (Movies 1 and 2) confirmed improved signal-to-noise throughout.

176

177 Images denoised with CryoSamba distinctly showed the characteristic phospholipid bilayer  
178 profile, a "double track" with a spacing of ~5 nm, in membranes of the endoplasmic reticulum (ER)  
179 (Fig. 2A), mitochondria (Fig. 2B), and early endosomes (Fig. 2C, D). The enhanced clarity also  
180 resolved ribosomes in the yeast cytosol, either free (Fig. 2A, B) or interacting with the ER  
181 membrane (Fig. 2A), as well as the ~5 nm spacing of subunits along actin filaments in the  
182 mammalian cell (Fig. 2D). Additional examples of cross sections of the plasma membrane in a  
183 BSC-1 cell and membrane-less rotaviruses in the surrounding medium are shown in (Fig. 3,  
184 panels 1-4 and panels 5-7, respectively).

185

186 To quantify the enhancement in visual quality, we assessed the signal-to-noise ratio (SNR)  
187 determined for single viral particles in two tomograms from BSC-1 cells and throughout three

188 tomograms from three yeast lamellae samples (Table 1). We determined SNR by two distinct  
189 methods. In one (Table 1), we generated pairs of tomograms using consecutive even and odd  
190 frames from the tilt series and calculated SNR by comparing identical xy planes in the raw data  
191 tomograms and those processed with either CryoSamba or the widely used cryo-ET denoising  
192 algorithm Topaz-denoise (Bepler et al, 2020). This method could not be used for data processed  
193 with CryoCARE (Buchholz et al, 2019), another widely used denoising algorithm, as it required  
194 consecutive even/odd tilt images for denoising; splitting them further would have degraded the  
195 CryoCARE output. To include CryoCARE in our evaluation, we used two additional approaches  
196 (Table 2): in one case, we computed the SNR using two adjacent xy planes in tomograms from  
197 the raw data and from tomograms denoised by CryoSamba, Topaz-Denoise and CryoCARE. In  
198 the second case, we estimated the SNR using the same xy plane in tomograms from the raw data  
199 and from tomograms denoised by CryoSamba, Topaz-denoise and CryoCARE. These results,  
200 summarized in Tables 1 and 2, showed comparable, substantial enhancements in SNR from both  
201 CryoSamba and Topaz-Denoise. Although CryoCARE also showed enhanced SNR, the increase  
202 came at the expense of the lower resolution imposed by the denoising algorithm (see below).

203

#### 204 **Comparative Impact on Resolution by Denoising with CryoSamba, CryoCARE and Topaz-** 205 **Denoise**

206 We compared the impact on resolution of denoising with CryoSamba, Topaz-Denoise and  
207 CryoCARE, by direct visual analysis of selected 2D planes and by examining the corresponding  
208 2D Fourier transforms. As shown in Fig. 4A, while all denoising methods enhanced contrast,  
209 CryoSamba preserved high-frequency information more effectively, particularly evident by  
210 preservation of the double-track appearance of lipid bilayers in the membrane cross sections. Fig.  
211 2 shows additional examples of membranes surrounding the endoplasmic reticulum and  
212 endosomes imaged in the yeast and BSC-1 cell samples. Inspection of the 2D Fourier transform  
213 supported our conclusions from visual inspection of the images. CryoSamba retained higher  
214 spatial frequencies that Topaz-denoise and CryoCARE flattened.

215

216 We also assessed by visual inspection the denoising performance of CryoSamba, Topaz-denoise  
217 and CryoCARE for the same tomogram processed with different voxel resolutions. We selected  
218 a BSC-1 cell region containing a membrane bound organelle and rotavirus particles, and  
219 generated tomograms to final voxel resolutions of 2.62 Å, 7.86 Å, and 15.72 Å, from unbinned, 3x  
220 and 6x binning, respectively. (Fig. 5). CryoSamba yielded less blurriness, higher contrast and  
221 better-preserved double-layer appearance of the cross section of the membrane and the



222 substructure of the virions across all resolutions, specially at 2.62 Å. Topaz-Denoise performed  
223 well at 7.86 Å, which is close to the resolution of its training data (~10 Å), but poorly at 2.62 Å and  
224 15.72 Å. CryoCARE had the poorest performance overall. The ability of CryoSamba to effectively  
225 handle a broad range of voxel resolutions highlights a useful versatility.

226

### 227 **CryoSamba and sub-tomogram averaging**

228 To compare CryoSamba with CryoCARE and Topaz-denoise, we used a basic STA procedure to  
229 identify and analyze rotavirus particles. We used 3D template matching and manual classification  
230 to select 54 particles across two 3x binned tomograms of BSC-1 cells (Tang et al., 2007). We  
231 carried out a single iteration of template-based alignment and averaging with final icosahedral  
232 symmetrization to achieve a simple average for feature analysis (see Methods for details). Further  
233 iterations did not improve map resolution. Since the template did not include rotavirus VP4 spikes,  
234 the presence of those spikes in the final averages indicated minimal template bias.

235

236 This analysis revealed notable differences in final map resolution. STA from tomograms derived  
237 from the raw data directly and from denoising by CryoSamba and Topaz-Denoise clearly  
238 delineated outlines of rotavirus proteins VP6 and VP7, in a T=13 icosahedral arrangement, and  
239 the projecting rotavirus VP4 spikes (Fig. 6B). The subtomogram average of the image denoised  
240 with CryoSamba also revealed VP1, the RNA polymerase surrounded by the dsRNA of the viral  
241 genome and thus in a particularly noisy environment (Fig. 6B, far right). Topaz-Denoise barely  
242 detected VP1 while CryoCARE could not.

243

244 We made a quantitative test of the impact of denoising on resolution by calculating a Fourier Shell  
245 Correlation (FSC) (Harauz et al., 1986) between half maps obtained by splitting the STA particles  
246 into even and odd sets and independently reconstructing separate subtomogram averages (the  
247 “gold standard” approach). CryoSamba maintained the closest correspondence to the raw data;  
248 Topaz-Denoise deviated at spatial frequencies beyond 28 Å<sup>-1</sup> while CryoCARE deviated  
249 significantly at 37 Å<sup>-1</sup>. In other words, Topaz-Denoise and CryoCARE tend to smooth high-spatial-  
250 frequency data more severely than does CryoSamba (Fig. 6C). By preferentially enhancing  
251 contrast at the expense of losing high-resolution data, denoising with Topaz-Denoise or  
252 CryoCARE may promote self-correlation or overfitting. This potential overfitting in the Topaz-  
253 Denoise and CryoCARE data is suggested by the initial rise of their FSC curves before it reaches  
254 zero, a pattern consistent with non-independence in the masked data (Scheres et al., 2012).

255

## 256 **DISCUSSION**

257 Our comparisons suggest that CryoSamba enhances both contrast and SNR without suppressing  
258 or distorting signal in the 3D tomogram, particularly at higher spatial frequencies. Its capabilities  
259 derive from its fully self-supervised algorithm, which denoises by directly engaging the 3D image  
260 itself. The optical flow interpolation method (Fig. 1B) reproduces 2D planes by reutilizing adjacent  
261 planes, not by generating new ones from scratch. This approach decreases artifacts and aids in  
262 the accurate reconstruction of high-frequency signals, which are often challenging for  
263 convolutional neural networks to capture (Rahaman et al., 2019). Moreover, the "one step back,  
264 one step forward" feature (Fig. 1C) further increases fidelity by amending residual deformations  
265 along the Z axis, effectively recycling information from the target plane to its copies.

266  
267 CryoSamba circumvents the inherent limitations common to denoising techniques that rely on  
268 synthetic data (Zeng et al., 2024), which may not generalize effectively to experimental images of  
269 varying characteristics or resolutions. It also avoids the pitfalls of methods that depend on noise  
270 modeling (Li et al., 2022), which can fail to fully grasp the complexities of noise in experimental  
271 data. Moreover, CryoSamba is not contingent on the use of paired volumes (Buchholz et al., 2019;  
272 Bepler et al., 2020), which are prone to misalignments or subtle signal discrepancies.

273  
274 A noteworthy advantage of CryoSamba is its ability to effectively denoise images with extremely  
275 low SNR in the raw data. For example, in Fig. 3, CryoSamba successfully denoised a tomogram  
276 processed at 2.62 Å/voxel, making apparent image details that in the raw image were almost  
277 indistinguishable from background. This capability of CryoSamba significantly broadens the direct  
278 visual inspection and potential analytical possibilities for postprocessing of data that might  
279 otherwise be considered unusable, particularly in cases involving thicker samples or data  
280 collected with reduced electron beam exposure times. The SNR and the improved contrast of the  
281 high-resolution raw image processed with CryoSamba matched those of the same volume  
282 typically associated with higher SNRs when down sampled by 3x or 6x pixel binning.

283  
284 In contrast to denoising with Topaz-Denoise and CryoCARE, which in part function as low-pass  
285 filters by eliminating higher frequencies, CryoSamba better preserves most high-frequency  
286 information. This ability of CryoSamba to maintain high resolution in denoised tomograms opens  
287 new avenues for using the denoised images in post-processing tasks that would traditionally  
288 require downsampling the raw image, potentially compromising the visibility of finer details and  
289 smaller structures.

290

291 A visual comparison of averaged subtomograms of rotavirus particles added to BSC-1 cells,  
292 together with a quantitative analysis of the corresponding 3D Fourier domain spectrum from the  
293 raw and denoised data using CryoSamba, Topaz-Denoise, and CryoCARE (Fig. 6), shows that  
294 CryoSamba can reduce noise while preserving essential features of the structure. It therefore has  
295 the potential to enhance the reliability of template-based particle picking and associated  
296 automated segmentation in cellular tomography.

297

298 Finally, we have engineered CryoSamba to be relatively computationally lightweight and  
299 compatible with most commercially available workstations. It is offered in two equivalent forms:  
300 one as a command-line tool and the other as a graphical user interface (GUI) executable,  
301 enhancing ease of installation and use. Upon approval of the peer-reviewed version of this  
302 manuscript, both versions of CryoSamba, along with tomograms suitable for testing, will be freely  
303 accessible via GitHub.

304

## 305 **CONCLUSION**

306 One important advantage of CryoSamba is that it operates directly on 3D reconstructed  
307 tomograms, therefore bypassing the need for accessing the original tilt series data or employing  
308 specific tomography reconstruction algorithms. Consequently, CryoSamba denoising can be  
309 incorporated at earlier stages of post-processing. While this direct approach to denoising 3D  
310 images significantly broadens its applicability, it is important to note that CryoSamba does not  
311 correct for imaging defects inherent in current tomographic data acquisition protocols, such as  
312 the missing wedge.

## 313 **MATERIALS AND METHODS**

### 314 **Sample Preparation**

315 *S. cerevisiae* cells were grown overnight in YPD medium, before being diluted to 0.2 OD. One  
316 yeast strain was arrested at the metaphase-anaphase transition by degrading CDC20 using an  
317 auxin-inducible degron tag; the other strain was synchronized by  $\alpha$ -factor for two hours, then  
318 released from G1 phase and plunge-frozen approximately 80 minutes later, such that most cells  
319 should have been in metaphase. Cells were deposited on Quantifoil R2/2 copper grids (Electron  
320 Microscopy Science) and plunge-frozen using a Leica EM-GP2 plunge (Leica).

321  
322 For rotavirus samples, BSC-1 cells grown on gold grids and plunge frozen after incubation with  
323 rotavirus particles was performed as described in (Abdelhakim et al., 2014). Quantifoil gold grids  
324 R 2/2 200 mesh coated with a holey SiO<sub>2</sub> film (Electron Microscopy Science) were glow  
325 discharged at 15 mA for 30 sec and then sterilized with 70% ethanol for 15 min. After washing 4  
326 times with sterile water the grids were incubated overnight with 0.1% poly-L-lysine hydrobromide  
327 (Gibco). BSC-1 cells (at a concentration of  $5 \times 10^5$  cells/ml) were plated on the grids in DMEM  
328 (Thermo Fisher Scientific) supplemented with 10% hi-FBS (Thermo Fisher Scientific) and 1%  
329 Glutamax (Thermo Fisher Scientific) and incubated for 5 h at 37 °C and 5 % CO<sub>2</sub>. Cells on grids  
330 were washed 3 times with DMEM (without FBS) and then the virus previously activated with 5  
331  $\mu$ g/ml trypsin was added at a MOI of 10 and incubated at 37 °C and 5 % CO<sub>2</sub> for 10 min. The  
332 grids were then blotted from the back with filter paper using the sensor blotting of the Leica Em  
333 GP2 Plunge Freezer, frozen in liquid ethane, and stored in liquid nitrogen.

334

### 335 **Tomogram Collection and Reconstruction**

336 We collected five tomograms for this study. All five tomograms were collected on a Thermo Fisher  
337 Krios G3i, with a BioQuantum energy filter and a K3 direct electron detector camera. All were  
338 collected in dose-fractionation mode. Yeast tomograms were collected in counting mode at 0.15  
339 s per frame for 1.348 s and positioning of collection was controlled by PACEtomo (Eisenstein et  
340 al., 2023). Virus tomograms were collected in counting mode at 0.15 s per frame for 1.198 s and  
341 positioning of the collection was controlled by SerialEM (Mastronarde, 2005). All collection was  
342 done at 2.620 Å/pix, or a magnification of 33k. The total dose was  $\sim 140$  e/Å<sup>2</sup>. Motion correction  
343 and initial alignment were done using an in-house on-the-fly pipeline, taking advantage of  
344 MotionCor2 (Zheng et al., 2017) and alignframes from IMOD (Kremer et al., 1996), respectively.  
345 Defocus files were generated by CTFFIND4 (Rohou et al., 2015).

346

347 After initial alignment, tilt series were then 3D-CTF-corrected and reconstructed using the  
348 NovaCTF pipeline (Turoňová et al., 2017), such that the whole tomogram was CTF-corrected  
349 prior to denoising and subtomogram averaging. Briefly, defocus files with a user-specified defocus  
350 step of 15 nm were applied to the CTFFIND4 files. Then, NovaCTF generated a CTF-corrected  
351 tilt series for each specified defocus step. Those CTF-corrected tilt series were individually aligned  
352 using Aretomo (Zheng et al., 2022). The resultant images were then flipped (preserving  
353 handedness of tomograms) using IMOD before filtering using the NovaCTF's implementation in  
354 the IMOD radial filtering package. Then, tomograms were reconstructed with the NovaCTF's 3D-  
355 CTF feature, using the relevant voxels from each CTF-corrected tilt series to reconstruct the final  
356 volume.

357

### 358 **Deep learning model architecture**

359 CryoSamba uses the motion-based video frame interpolation model EBME (Jin et al., 2023),  
360 which takes as input two images (implicitly belonging to a temporal sequence) and returns a new  
361 one at an arbitrary time-point between them. As shown in Fig. 1B, this model has two main steps:  
362 bi-directional motion estimation and frame fusion.

363

364 In the motion estimation step, each input frame is initially down sampled by increasing factors of  
365 two, to form an “inverted pyramid” of images with a predetermined number of levels. At the lowest  
366 resolution level, the two corresponding downsampled images are each processed by a  
367 convolutional neural network (CNN); the correlation of the CNN output features is used to estimate  
368 two deformation fields. These fields, which can be used to “warp” one image to the other and vice-  
369 versa, are known as Optical Flow (Horn et al., 1981) and represent the (directional) pixel motion  
370 between the two frames. EBME uses a warping process known as Softmax Splatting (Niklaus et  
371 al., 2020), which directly deforms the frames to each other and combines them by a small U-Net  
372 (Ronneberger et al., 2015). At every subsequent pyramid level, the images are warped towards  
373 each other by the (upscaled) corresponding estimated flow, processed by the CNN, correlated  
374 with each other, and combined with the previous level's outputs (as in a recurrent network (Jin et  
375 al., 2023)) to produce a refined version of the bi-directional flow. After the final level, the output  
376 flows are upsampled to the frame's native resolution and taken to the next step.

377

378 In the fusion step, the two images are first warped towards an arbitrary time-point between them  
379 by the estimated bi-directional flow. This process assumes a linear motion model (Jin et al., 2023;  
380 Xu et al., 2019), where each flow is simply scaled by the temporal distance between the

381 destination and its origin. In parallel, each image is also processed by a “context” CNN. The flows,  
382 warped frames, and contextual features are all combined as input to a “refinement” U-Net, which  
383 synthesizes the desired intermediate frame.

384

385 We used EBME with three pyramid levels for both training and inference, as lower values  
386 decreased performance and higher values, which increase computational costs, did not bring  
387 noticeable improvements. We also used the “high synthesis” mode (Jin et al., 2023) of EBME,  
388 which upscales flows and images by a factor of two before the fusion step, subsequently  
389 downscaling the final output to the native resolution at the end. This mode allows high-resolution  
390 processing and greatly increases the interpolation quality of small objects. The arbitrary time-point  
391 in the fusion step was always chosen to be exactly halfway between the two input images, as  
392 “asymmetric” interpolation did not work well when training the model with our chosen datasets.  
393 Finally, we used the “reflect” padding mode in all convolutional layers, to reduce the typical CNN  
394 artifacts (Alsallakh et al., 2021) at the borders of the synthesized frames, which were especially  
395 prominent after the “Samba” processes. The final model had 2.9 million weights.

396

### 397 **Data preprocessing**

398 CryoSamba directly accepts as input a list of 3D image volumes in three different formats (which  
399 are automatically recognized), with no need for preprocessing and/or conversion by the user: as  
400 a single “.mrc” (or “.rec”) file, as a single “.tif” file, or as a folder containing an alphanumerically  
401 ordered sequence of 2D “.tif” images. The volumes are then initialized as memory-mapped arrays,  
402 which means that only the subregions that are currently being processed are loaded in memory,  
403 being freed afterwards. This allows us to fit very large volumes (which are typical in Cryo-ET) into  
404 systems with limited RAM.

405

406 For preprocessing, the data were converted (if necessary) to floating point values, and the  
407 intensity of each whole volume was normalized between -1 and 1. The volumes were then divided  
408 into overlapping planes of shape 256x256x1 (xyz), each of which corresponded to a square crop  
409 of a single xy plane. The crops overlapped with each other by 16 pixels along both x and y. If the  
410 number of voxels along x and/or y was not divisible by the corresponding plane size (considering  
411 the overlaps), the smaller cropped arrays at the edges were padded until they had the same  
412 shape (256x256 pixels) as the others. The padded pixels were filled not with a constant value,  
413 but with the original crop’s pixel values reflected from its border, to preserve the continuity and  
414 smoothness of the full, padded crop.

415

416 The plane crops were then combined into blocks, each consisting of three planes at the same  
417 position in x and y and at the positions z-L, z, and z+L along Z. The values of L ranged from L=1  
418 to a fixed value  $L=L_{\max}$ , while z covered the whole z-axis except for the border values  $z < L_{\max}$  and  
419  $z > z_{\max} - L_{\max}$  that were not included. The final dataset, which consisted of the full set of blocks,  
420 was loaded into eight Nvidia A100 GPUs via Pytorch's Distributed Data Parallel protocol (DDP)  
421 (Paszke et al., 2019; Zhao et al., 2020), in batches of 32 blocks each via 4 CPU workers.

422

### 423 **Training pipeline**

424 For training, the  $L_{\max}$  of the datasets was chosen as 3, 6 and 10 for pixel bins of 6x, 3x and 1x,  
425 respectively (Fig. S1). The resulting datasets were separated into training (95%) and validation  
426 (5%) splits, which were used to train the EBME model by backpropagation with the Adam  
427 optimizer (Kingma et al., 2015), with learning rate of  $2e-4$ , a warmup scheduler for 300 iterations  
428 and a learning rate decay scheduler with a multiplicative factor of 0.99995. For each training  
429 iteration (i.e., when a single data point was passed through the model), the first and last planes  
430 of a (three-plane) block were fed into the EBME model, whose output was then compared with  
431 the block's middle plane by photometric loss, i.e., a measure of their pixel intensity differences.  
432 The latter was a combination of a Charbonnier loss (Charbonnier et al., 1994) (a regularized  
433 version of mean squared error) with  $\alpha=0.5$  and  $\epsilon=1e-3$  and a ternary Census loss (a photometric  
434 loss invariant to illumination changes (Meister et al., 2018)) weighted by  $\lambda=0.1$ .

435

436 After every epoch, i.e., after every block of the full dataset had been passed through the model  
437 once, the block order of the training split was randomly shuffled. Furthermore, each block was  
438 transformed by randomly flipping along any of the three spatial dimensions. This  
439 transformation/deformation process, known as data augmentation (Shorten et al., 2019),  
440 increases data diversity without fundamentally altering its nature and reduces overfitting in neural  
441 networks. Training ended when both training and validation losses stabilized, which on average  
442 occurred after  $L_{\max} * 30k$  iterations.

443

### 444 **Inference pipeline**

445 For inference,  $L_{\max}$  was chosen as 6, 12 and 20 for pixel bins of 6x, 3x and 1x, respectively (Fig.  
446 S1). Each dataset block was passed into a Samba module (Fig. 1D) with the underlying EBME  
447 model trained on the same data. The corresponding outputs for blocks with different inter-plane  
448 spacings L but same middle plane z were then averaged and used to compose the final volume.

449 During inference, the latter was memory-mapped as a binary file, to fit into memory and to allow  
450 resuming the whole procedure from the last step in case it had to be interrupted. After running  
451 over the whole dataset, the final volume was converted back to its original intensity range, data  
452 type and file format.

453  
454 The dataset blocks were also flipped with respect to x, y, and both x and y, and separately  
455 processed by the deep learning model. The three outputs were flipped again (undoing their  
456 transformation) and averaged with the original denoised output. This procedure (which can be  
457 optionally turned on or off in CryoSamba's final version), known as Test-Time Augmentation (TTA)  
458 (Ayhan et al., 2018), led to a small increase in inference quality at the cost of a nearly four-fold  
459 increase in inference time.

460  
461 **Hyperparameter tuning**  
462 The data block's xy dimensions were chosen as 256x256 to include sufficient contextual  
463 information in each block, to ensure that border artifacts would not occupy a significant portion of  
464 the blocks, and to take a reasonable training/inference time without overloading the GPUs. Note  
465 that the x and y values can be chosen with a certain freedom due to the convolutional nature of  
466 the EBME model, which accepts inputs of different shapes without having to be modified or re-  
467 trained.

468  
469 The loss and optimizer hyperparameters were chosen to minimize the resulting loss functions.  
470 The number of EBME pyramid levels, which is flexible enough to have different values for training  
471 and inference (Jin et al., 2023), was chosen as 3 for both routines through a grid search aimed at  
472 maximizing the final denoising performance as assessed by visual inspection.

473  
474 The most important hyperparameter to tune was  $L_{\max}$ . For training, larger values improve the final  
475 denoising results up to a certain point, after which the improvements are negligible and increasing  
476  $L_{\max}$  only increases the computational burden (longer processing times and/or more memory  
477 usage). For inference, increasing this parameter increases the contrast enhancement effect but  
478 also the overall blurriness, so an optimized value must be found to maximize the final denoising  
479 quality. To do this, we made a simple script that, after a training run with a sufficiently large  $L_{\max}$ ,  
480 performed inference on a small crop of the whole dataset with varying values of  $L_{\max}$  and displayed  
481 the results side by side. The optimal inference value of  $L_{\max}$  was then chosen by visual inspection  
482 of these results (Fig. S1), and the optimal training value was retroactively chosen (to be useful for



483 the subsequent training runs) as half of this value. The reason for this choice is that, while training  
484 does interpolation with the z-L and z+L plane pairs of each data block, inference only does so for  
485 the corresponding (z-L, z) and (z, z+L) pairs (due to the Samba procedure), which are separated  
486 at half the distance than the former. We chose inference (training)  $L_{\max}$  values of 6 (3), 12 (6) and  
487 20 (10) for pixel bins of 6x, 3x and 1x, respectively.

488

### 489 **Computational requirements**

490 The experiments reported in this work were each run on eight Nvidia A-100 Tensor Core GPUs  
491 (Table 3). Model training used about 31 GB of VRAM in each GPU and took, on average, 0.4  
492 seconds for each iteration, with total training time (number of iterations) depending on the  
493 convergence speeds. Inference (using the trained model to denoise the data) used about 8GB of  
494 VRAM in each GPU, took about 0.9 seconds per iteration, with a completion time that was linearly  
495 proportional to the size of the volume data and the chosen  $L_{\max}$ . Even though training times in  
496 deep learning are usually longer than inference times, due to calculation of losses and  
497 propagation of gradients, this situation was inverted in our case due to the “Samba” and TTA  
498 procedures, which lead to 12 (3 times 4) forward passes in each iteration instead of just one.

499

500 In CryoSamba, the main parameters that affect VRAM usage without significantly altering  
501 performance are batch size and number of GPUs, which defaulted to 32 and 8, respectively. Batch  
502 size defines the quantity of data passed at once to the GPUs, and running the model on N devices  
503 allows us to increase effective batch size N fold without increasing computational time. For multi-  
504 GPU runs we used DDP: the deep learning model was run independently on each device by  
505 copying all model weights for each of them and dividing the data batches among them, while  
506 using advanced procedures to synchronize their gradients during training (Zhao et al., 2020). DDP  
507 drastically improves iteration times at the cost of massively increasing VRAM usage (divided  
508 across all GPUs). We also ran the models with mixed precision (Micikevicius et al., 2017) and the  
509 optimization routines of Pytorch’s Compiler (Ansel et al., 2024) when possible (depending on the  
510 GPU’s capabilities and operating system), both of which further improved the overall  
511 computational times.

512

513 By tuning all these parameters, we could also run CryoSamba with as little as 3GB of VRAM  
514 (using batch size of 2), which fits into most personal laptop GPUs. In Table 3, we show practical  
515 VRAM usage and iteration times for a low budget laptop setup with 6 GB of VRAM and one GPU.

516 CryoSamba is a democratic algorithm: it can run on workstations currently available to most  
517 research groups.

518

### 519 **Denosing using Benchmarking Algorithms**

520 Implementation of Topaz-Denoise was done using Topaz version 0.2.5\_cu11.2 (SBGrid build).  
521 The prebuilt U-Net model was used in all cases as we found it to provide superior results over  
522 self-trained models. We used the 8 GPUs available in our workstations with a patch size of 96,  
523 patch padding of 48 and no gaussian filter.

524

525 For CryoCARE, We used the CryoCARE memory efficient version  
526 ([https://github.com/juglab/cryoCARE\\_pip](https://github.com/juglab/cryoCARE_pip)) with default parameters, except for training batch size  
527 and inference number of tiles, which were chosen in order to fit the data into GPU memory.  
528 Training and inference were performed with one GPU.

529

### 530 **Subtomogram Averaging**

531 Particles were selected from two rotavirus tomograms with the command e2spt\_tempmatch.py in  
532 EMAN2 (Tang et al., 2007), using an atomic model of a rotavirus particle with the VP4 spikes and  
533 the interior masked out. 64 targets were selected automatically and extracted using  
534 e2spt\_boxer.py. After manual inspection and removal of incorrectly selected particles, particles  
535 were aligned and averaged using e2spt\_classaverage.py. Initial alignment was in c1, and  
536 icosahedral symmetry was then applied to the final map. The coordinates of particles were taken  
537 from the autogenerated “info” file by e2spt\_tempmatch.py and used to extract sub volumes in  
538 each of the denoised volumes. The same subtomogram averaging protocol was applied to each  
539 denoised particle set. Subtomogram averages were visualized in ChimeraX (Pettersen et al.,  
540 2021) and sharpened to equivalent relevant levels.

541

### 542 **Signal-to-Noise Calculations**

543 We estimated SNR as in Bepler et al (Bepler et al., 2020), as originally described by Frank and  
544 Al-Ali (Frank et al., 1975). The cross-correlation coefficient (CCC) was calculated between  
545 tomograms reconstructed from even tilts and tomograms denoised with CryoSamba or Topaz-  
546 Denoise reconstructed from odd tilts, and the SNR obtained from:

547

$$548 \quad SNR = 10 * \log_{10} \left( \frac{CCC}{1 - CCC} \right)$$

549

550 SNR was calculated for three different complete yeast tomograms and for two rotavirus particles,  
551 each selected from a different BSC-1 cell tomogram (Fig. S1, Table 1). This approach is not  
552 applicable to data denoised with CryoCARE because it requires information from all adjacent tilted  
553 frames, and hence using alternate tilts weakens its denoising capability. As a potential  
554 complementary approach, we obtained SNR values by comparing adjacent xy planes along the z  
555 axis within the raw and denoised tomograms (Table 2). We also obtained SNR values by  
556 comparing the same xy plane between the raw and denoised tomograms (Table 2). These values,  
557 however, were less informative because they correlated with the contrast enhancement  
558 regardless of preservation of the high frequency information.

## 559 **FIGURE LEGENDS**

### 560 **Figure 1. CryoSamba Pipeline for Deep-Learning Neural Network Denoising**

561 **(A)** Training Pipeline: This component uses sets of three xy planes at z-axis positions z-L, z, and  
562 z+L. The planes at z-L and z+L are input into the EBME deep learning model. The model's loss  
563 function evaluates the mismatch between its output and the middle plane at position z, supervising  
564 and refining the training. This iterative process is repeated for every z position (ranging from 1 to  
565 500 in this example), with updates to the learned model weights at each step. A complete cycle  
566 ("epoch") involves iterating through the entire dataset of images. Reaching a predetermined  
567 convergence criterion generally requires several epochs.

568 **(B)** EBME Model: This model comprises two main modules. The Flow module assesses the bi-  
569 directional optical flow between the two input images; the Fusion module integrates the input  
570 images along with the calculated flow to predict an interpolated image. This interpolated image is  
571 positioned as if halfway between the two inputs when viewed as part of a temporal sequence.

572 **(C)** Denoising Pipeline: Each xy plane from the tomogram is grouped with its adjacent z-1 and  
573 z+1 planes and processed through the trained CryoSamba module. This results in a denoised xy  
574 plane that supplants the original in the completed volume.

575 **(D)** Inference Module: For each triplet, the extremal planes z-1 and z+1, along with the middle z  
576 plane, are fed separately into the EBME model. The outputs from these inputs are then  
577 reintroduced to the EBME model to refine and produce the final version of the middle z plane.

578 **(E)** CryoSamba Module: This phase involves iterative processing, in which a specific plane z and  
579 a series of up to  $L_{\max}$  surrounding planes are input into the inference module. Each triplet, z-L, z,  
580 z+L (where  $1 \leq L \leq L_{\max}$ ), is processed through the Samba module as described in (D). The results  
581 are then averaged to produce a final, denoised version of plane z.

582

### 583 **Figure 2. Visual Comparison of Raw and CryoSamba Denoised 3D Images of Cryo-Electron** 584 **Tomograms.**

585 Visual comparison of the same xy planes from representative raw and CryoSamba denoised 3D  
586 images of cryo-tomograms, binned at 3x with a voxel resolution of 7.86 Å.

587 **(A-C)** Different tomograms derived from three distinct yeast cells. **(A)** Cross-sections of the  
588 endoplasmic reticulum and ribosomes. **(B)** Cross-section of mitochondria. **(C)** Endosome

589 containing intraluminal vesicles. The denoised images highlight the preservation of the double-  
590 layer appearance of the membranes, separated by approximately 4-5 nm.

591 **(D)** Tomogram from a BSC1 cell illustrating actin filaments and the cross-section of a membrane-  
592 bound organelle of unknown identity. Distance between the arrowheads is consistent with the  
593 expected 5.5 - 6 nm periodicity of actin monomers along the helical actin filament.

594 Scale bars: 50 nm (A-C) and 100 nm (D).

595

596 **Figure 3. Denoising using CryoSamba of a BSC-1 Cell Incubated with Non-Enveloped**  
597 **Rotavirus.**

598 This figure illustrates the effects of CryoSamba denoising on a representative cryo-electron  
599 tomography image of a BSC-1 cell incubated with rotavirus.

600 **(A)** Representative xy plane from a 3D cryo-electron tomogram obtained at 2.62 Å/voxel  
601 resolution and reconstructed to 7.86 Å/voxel resolution after 3x binning. The comparison shows  
602 the raw image before (left) and after (right) denoising. Selected regions of interest are indicated.

603 **(B)** Enlarged xy and zx views of the selected regions before and after denoising, showing  
604 improved clarity in the images post-denoising. These enhancements are particularly noticeable in  
605 the cross-section views, showing the double-track appearance of membranes and the  
606 substructure within the virus particles. Scale bars: 100 nm (A) and 50 nm (B).

607

608 **Figure 4. Comparison of CryoSamba, Topaz-Denoise, and CryoCARE in the 2D Fourier**  
609 **Plane.**

610 **(A)** Representative image from a tomogram of a BSC1 cell at 7.86 Å voxel resolution (3x binning),  
611 depicting the middle section of an organelle surrounded by two sets of membranes. The left  
612 column shows a selected xy view of a single plane of the raw image or of ones denoised using  
613 CryoSamba, Topaz-Denoise, or CryoCARE. The right column shows the (logarithm of) magnitude  
614 of the 3D Fourier transform for the corresponding regions on the left, averaged over 64 planes  
615 along the Z axis to reduce noise. The more populated higher frequency regions observed after  
616 denoising with CryoSamba are consistent with relatively better preservation of high-frequency  
617 information and the clearer double-track appearance of the membrane cross sections.

618 **(B)** Orthogonal view of (A) corresponding to the xz plane. Inspection of the 2D Fourier transforms  
619 illustrates the extent to which high-frequency information is retained in the sample denoised using

620 CryoSamba, the modest suppression of high spatial frequencies by Topaz-Denoise, and the  
621 dramatic loss of high spatial frequencies after denoising by CryoCARE.

622 Scale bars: 50 nm (A, B).

623

624 **Figure 5. Comparison of CryoSamba, Topaz-Denoise, and CryoCARE Performance at**  
625 **Different Image Resolutions.**

626 Visual comparison of the same xy plane selected from the 3D reconstructed tomogram of a BSC1  
627 cell incubated with rotaviruses, shown at voxel resolutions of 2.62 Å (left column), 7.86 Å (middle  
628 column), and 15.72 Å (right column). The images are presented before denoising (raw) and after  
629 denoising using CryoSamba, Topaz-Denoise, or CryoCARE. Scale bars: 100 nm. The images  
630 denoised with CryoSamba appear less blurred, particularly at 2.62 Å resolution and highlight the  
631 preservation of the bilayer appearance of the membrane, spaced 4-5 nm and the fine structure  
632 within the rotavirus particles.

633

634 **Figure 6. Comparison of CryoSamba, Topaz-Denoise, and CryoCARE in Information**  
635 **Content.**

636 **(A)** Comparison of the same xy plane selected from the 3D reconstructed tomogram of a rotavirus  
637 particle present in the sample of a BSC1 cell incubated with rotaviruses, shown at voxel  
638 resolutions of 2.62 Å (no binning) and denoised with CryoSamba, Topaz-Denoise, or CryoCARE.  
639 Spikes are visible in the particle denoised using CryoSamba and Topaz but absent in the sample  
640 denoised with CryoCARE. **Scale bar:** 35 nm.

641 **(B)** Comparison of subtomogram averaging results obtained using raw and denoised images. The  
642 far-left column shows projected averaged maps, and the middle columns show surface renditions  
643 of the corresponding electron density maps obtained using ChimeraX. The renditions were  
644 adjusted to equivalent intensity levels of the top and middle slices. The enlarged region is centered  
645 on the location of VP1 and highlights its detection in the raw and denoised images using  
646 CryoSamba or Topaz-Denoise, but its absence in the image denoised using CryoCARE.

647 **(C)** Comparison of the Fourier Shell Correlation calculated for the averaged subtomograms of the  
648 rotavirus before (raw) or after denoising with CryoSamba, Topaz-Denoise, or CryoCARE. The  
649 comparison highlights the similar behavior of the Fourier Shell Correlations of the raw and  
650 CryoSamba-denoised images, in contrast to the increasing values at high resolution observed for

651 the denoised data obtained using Topaz-Denoise or CryoCARE, which appear to create high-  
652 frequency "signal" from the high-frequency noise.

653

654 **Figure S1. Effect of Maximum Plane Gap  $L_{\max}$  on CryoSamba Denoised Images.**

655 2D xy plane (top rows) and corresponding 2D Fourier transform magnitudes (bottom rows) of a  
656 3D reconstructed tomogram of a BSC1 cell at voxel resolutions of (A) 15.72 Å, (B) 7.86 Å, and  
657 (C) 2.62 Å, displaying the cross-section of a defective rotavirus without spikes. Each series of  
658 images starts with the raw xy plane and is followed by the same plane after CryoSamba denoising  
659 with increasing values of the maximum plane gap hyperparameter  $L_{\max}$ . As  $L_{\max}$  increases, both  
660 the contrast enhancement (desirable) and blurriness (undesirable) effects increase. An optimal  
661 "balancing" point (image boxed in red) is chosen based on visual inspection of these images over  
662 a reasonable range of values. The blurriness effect, which limits  $L_{\max}$ , is more pronounced when  
663 the relative size of structures of interest is smaller with respect to the voxel size (which dictates  
664 the "thickness" of each 2D plane), leading to smaller values of  $L_{\max}$  for higher pixel binning's.  
665 Scale bars: 50 nm.

666

667 **Figure S2. Fourier Plane Modulation Effects of Simple Averaging, Direct Interpolation, and  
668 the Samba Procedure.**

669 We apply three different denoising procedures to the xy planes of a small 3D region of a BSC1  
670 cell cryoET tomogram and show the resulting xz Fourier transform magnitudes. The raw data  
671 image on the left displays a bundle of tilted lines characteristic of cryoET reconstructions.

672 For the first row of data, each xy plane at its z position was directly replaced by the average of all  
673 planes between  $z-L_{\max}$  and  $z+L_{\max}$ , for  $L_{\max}$  ranging from 1 to 5. This "naïve" denoising approach  
674 enhances the imaging contrast at the cost of severe distortions, evidenced by the cosine  
675 modulations along the Z axis of the Fourier images.

676

677 In the second row, the data was denoised with our trained deep learning model using a direct  
678 interpolation approach, where planes  $z-L$  and  $z+L$  are directly passed through the EBME module,  
679 and the outputs for  $L=1\dots, L_{\max}$  are averaged. Depending on the trained model's performance,  
680 the result could range from perfect denoising to merely averaging its inputs. The result lies

681 between these extremes, with denoising coupled with undesirable modulation reminiscent of the  
682 averaging process.

683

684 In the third row, we show the results of the CryoSamba module applied to this 3D region. The  
685 Samba procedure is designed to suppress the modulations seen in the second row. By including  
686 the original plane z in the interpolation process, we ensure that information not easily inferred  
687 from neighboring planes is propagated to the result, avoiding most of the information loss that  
688 causes the xz Fourier plane modulations. Drawings from yeast, monkey, ER and mitochondria  
689 are based on images from Servier Medical Art. Servier Medical Art by Servier is licensed under a  
690 Creative Commons Attribution 3.0 Unported License  
691 (<https://creativecommons.org/licenses/by/3.0/>).



692 **MOVIE LEGENDS**

693 **Movie S1. CryoSamba denoising of a representative plane in a 3D tomogram from a BSC-**  
694 **1 cell.**

695 Comparison of the same image corresponding to a xy plane in a 3D tomogram volume of a BSC-  
696 1 cell incubated with rotaviruses, before and after CryoSamba denoising. The data were acquired  
697 with a per-voxel resolution of 2.62 Å and subsequently 3x pixel binned, resulting in a final  
698 resolution of 7.86 Å per voxel.

699

700 **Movie S2. CryoSamba denoising of a representative plane in a 3D tomogram from a yeast**  
701 **cell.**

702 Comparison of the same image corresponding to a xy plane in a 3D tomogram volume of a yeast  
703 cell highlighting the cross section of a portion of endoplasmic reticulum (ER) and ribosomes, free  
704 in the cytosol or associated with the ER, before and after CryoSamba denoising. The data were  
705 acquired with a per-voxel resolution of 2.62 Å.

706

707 **Movie S3. CryoSamba denoising of a 3D tomogram from a yeast cell.**

708 Comparison of the same tomogram volume of a yeast cell highlighting the cross section of a  
709 portion of endoplasmic reticulum (ER) and ribosomes, free in the cytosol or associated with the  
710 ER, before and after CryoSamba denoising. The data were acquired with a per-voxel resolution  
711 of 2.62 Å.

## 712 **TABLE LEGENDS**

### 713 **Table 1. SNR for images determined using alternate tilt planes and SNR ratios between raw** 714 **and images denoised using CryoSamba or Topaz-Denoise.**

715 Data for two raw tomograms per experimental condition were obtained using independent sets of  
716 alternate xy tilt plane images acquired with an unbinned voxel size of 2.62 Å. SNR data for each  
717 tomogram (in dB) were calculated by correlating each xy plane throughout the volume of the two  
718 tomogram subsets (see Methods). Values in parentheses represent SNR ratios derived from the  
719 SNR values of the raw images and the corresponding SNR values of the same images denoised  
720 using CryoSamba or Topaz-Denoise (see Methods). The BSC-1 cell sample data sets capture  
721 two different rotavirus particles selected from different tomograms. The yeast sample data sets  
722 are complete tomograms obtained from three lamellae. These results, which compare raw and  
723 denoised images with the same information content, demonstrate the similar extent of SNR  
724 improvement achieved using both denoising methods.

### 725 726 **Table 2. SNR for images determined at different resolutions using all tilt planes and SNR** 727 **ratios between raw and images denoised using CryoSamba, Topaz-Denoise or Cryo-CARE.**

728 Data from the tomograms obtained using all xy tilt plane images of the BSC-1 cell depicted in  
729 Fig.5 acquired with voxel size of 2.62 Å and processed at voxel size of 2.62 Å (unbinned), 7.86 Å  
730 (3x binned) and 15.72 Å (6x binned). SNRs were determined (see Methods) by correlating alternate  
731 xy planes (with slightly different information content) along the z axis of a given cryo-ET  
732 (Even/Odd plane) or by correlating the same xy plane (with the same information content) along  
733 the z axis determined between the raw image and the same image denoised using CryoSamba,  
734 Topaz-Denoise or Cryo-CARE (Denoise/Raw); these values are trivially infinity when the  
735 comparison is between raw and itself. The highest SNR value within a given column is highlighted  
736 in red. The results, obtained by relating raw and denoise images with the same information  
737 content highlight the extent of SNR improvement obtained using CryoSamba.

### 738 739 **Table 3. Comparison of training and inference times and VRAM usage for two** 740 **workstations.**

741 The table presents a comparison of typical training and inference times, as well as VRAM usage,  
742 for two workstations employed to denoise the same tomogram using CryoSamba. The DGX-A  
743 workstation, equipped with 8 Nvidia A-100 GPUs (40GB VRAM each), utilized PyTorch's  
744 distributed data parallel (DDP) protocol to split the training across all GPUs, and the PyTorch  
745 Compiler function to speed up computations. The LENOVO Legion Pro 5 laptop, containing a

746 single Nvidia Geforce RTX 3070 GPU card (8GB VRAM), operated under Windows OS and could  
747 not use the PyTorch Compiler due to OS constraints. Batch sizes were adjusted to fit the available  
748 GPU memory for each workstation.

749

750 Denoising calculations from a tomogram of 682 x 960 x 333 voxels, with a voxel resolution of  
751 15.72 Å (6x binning) of a BSC-1 cell incubated with rotavirus were performed with Lmax gaps of  
752 3 and 6 for training and inference, respectively. Times for training (per epoch) and total inference  
753 were approximately linearly proportional to the volume dimensions and maximum frame gap.  
754 Total training time is dependent on the convergence of the loss functions and tended to stabilize  
755 after around 30 epochs of training.

756 **ACKNOWLEDGMENTS**

757 The authors thank Stephen C. Harrison for insight and editorial help, Conny Leistner for  
758 assistance in obtaining all the tomograms, Arkash Jain for assistance in the code release, Justin  
759 O'Connor for maintaining the IT infrastructure in the TK lab, and members of the TK lab for their  
760 help and support. This work was supported by a National Institute of General Medical Sciences  
761 Maximizing Investigators' Research Award GM130386 to T.K. J.I.C.F. was supported in part by a  
762 grant from IONIS to T.K. L.M.T. acknowledges a National Science Foundation Graduate  
763 Research Fellowship under Grant DGE 2140743. M.S. was supported by a NIH Grant  
764 R01CA013202-51 to Stephen C. Harrison. The facilities of the Harvard Cryo-EM Center for  
765 Structural Biology were essential for FIB-milling the yeast cells and for recording tilt series for all  
766 the tomograms. Acquisition of the computing hardware including the DGX's GPU-based  
767 computers, CPU clusters, fast access memory, archival servers and workstations that made  
768 possible this study were supported by generous grants from the Massachusetts Life Sciences  
769 Center to T.K. and an equipment supplement to the National Institute of General Medical Sciences  
770 Maximizing Investigators' Research Award GM130386. Construction of the server room housing  
771 the computing hardware was made possible with generous support from the PCMM Program at  
772 Boston Children's Hospital.

773

774 The authors declare no competing financial interests.

775

776 **AUTHOR CONTRIBUTIONS**

777 J.I.C.F. and T.K. conceived the project. J.I.C.F. performed the denoising computational work with  
778 input from T.K. and L.T. L.T. and M.D.S prepared the samples and collected the cryoET data.  
779 L.T. calculated the subtomogram averages. J.I.C.F, L.T. and T.K. participated in data analysis  
780 and interpretation. J.I.C.F., L.T. and T. K. wrote the manuscript. All authors provided feedback  
781 and agreed on the final manuscript.

782

783 **CODE AND DATA AVAILABILITY**

784 The code of CryoSamba and subtomogram examples used in this study are available upon  
785 request. They will be freely available in GitHub upon approval of the peer reviewed version of the  
786 manuscript.

787

788 **REFERENCES**

789 Bepler, T., Kelley, K., Noble, A.J. et al. Topaz-Denoise: general deep denoising models for  
790 cryoEM and cryoET. Nat Communications 11, 5208 (2020). [https://doi.org/10.1038/s41467-020-](https://doi.org/10.1038/s41467-020-18952-1)  
791 18952-1.

792  
793 Buchholz, T-O., Krull, A., Shahidi, R., Pigino, G., Jékely, G., & Jug, F. (2019). 'Content-aware  
794 Image Restoration for Electron Microscopy'. Methods in Cell Biology, 152, pp. 277-289

795  
796 Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T. and Ronneberger, O., 2016, October. 3D U-  
797 Net: learning dense volumetric segmentation from sparse annotation. In International conference  
798 on medical image computing and computer-assisted intervention (pp. 424-432). Springer, Cham.

799  
800 Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z.,  
801 Gimelshein, N., Antiga, L. and Desmaison, A., 2019. Pytorch: An imperative style, high-  
802 performance deep learning library. Advances in neural information processing systems, 32.

803  
804 Jin, X., Wu, L., Shen, G., Chen, Y., Chen, J., Koo, J., & Hahm, C-H. (2023). 'Enhanced Bi-  
805 directional Motion Estimation for Video Frame Interpolation', 2023 IEEE/CVF Winter Conference  
806 on Applications of Computer Vision (WACV), pp. 5038-5046. doi:  
807 10.1109/WACV56688.2023.00502.

808  
809 Ronneberger, O., Fischer, P., & Brox, T. (2015). 'U-Net: Convolutional Networks for Biomedical  
810 Image Segmentation'. In: Navab, N., Hornegger, J., Wells, W., & Frangi, A. (eds) Medical Image  
811 Computing and Computer-Assisted Intervention – MICCAI 2015, Lecture Notes in Computer  
812 Science, vol. 9351. Springer, Cham. doi: 10.1007/978-3-319-24574-4\_28.

813  
814 LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. Nature 521, 436–444 (2015).  
815 <https://doi.org/10.1038/nature14539>.

816  
817 Reda, F.A., Sun, D., Dundar, A., Shoeybi, M., Liu, G., Shih, K.J., Tao, A., Kautz, J., & Catanzaro,  
818 B. (2019). 'Unsupervised Video Interpolation Using Cycle Consistency'. In ICCV 2019.

819  
820 Guillermprieto, A. (1991). Samba. 1st Vintage Departures ed. Vintage Books, New York, NY..

821

- 822 Hua, T., et al. (2021). 'On Feature Decorrelation in Self-Supervised Learning'. In 2021 IEEE/CVF  
823 International Conference on Computer Vision (ICCV), Montreal, QC, Canada, pp. 9578-9588. doi:  
824 10.1109/ICCV48922.2021.00946  
825
- 826 Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., & Courville,  
827 A. (2019). 'On the Spectral Bias of Neural Networks'. In Proceedings of the 36th International  
828 Conference on Machine Learning, edited by Chaudhuri, K. & Salakhutdinov, R., vol. 97,  
829 Proceedings of Machine Learning Research, PMLR, pp. 5301-5310.  
830
- 831 Horn, B.K.P. & Schunck, B.G. (1981). 'Determining Optical Flow'. *Artificial Intelligence*, 17(1-3),  
832 pp. 185-203. doi: 10.1016/0004-3702(81)90024-2  
833
- 834 Niklaus, S. & Liu, F. (2020). 'Softmax Splatting for Video Frame Interpolation'. In IEEE Conference  
835 on Computer Vision and Pattern Recognition (CVPR).  
836
- 837 Xu, X., Li, S., Sun, W., Yin, Q., & Yang, M-H. (2019). 'Quadratic Video Interpolation'. In NeurIPS.  
838 Alsallakh, B., Kokhlikyan, N., Miglani, V., Yuan, J., & Reblitz-Richardson, O. (2021). 'Mind the Pad  
839 — CNNs Can Develop Blind Spots'. In International Conference on Learning Representations.  
840
- 841 Li, S., Zhao, Y., Varma, R., Salpekar, O., Noordhuis, P., Li, T., Paszke, A., Smith, J., Vaughan,  
842 B., Damania, P., & Chintala, S. (2020). 'PyTorch Distributed: Experiences on Accelerating Data  
843 Parallel Training'. *Proc. VLDB Endow.*, 13(12), pp. 3005–3018. doi: 10.14778/3415478.3415530  
844
- 845 Kingma, D. & Ba, J. (2015). 'Adam: A Method for Stochastic Optimization'. In Proceedings of the  
846 3rd International Conference on Learning Representations (ICLR 2015).  
847
- 848 Charbonnier, P., Blanc-Féraud, L., Aubert, G., & Barlaud, M. (1994). 'Two Deterministic Half-  
849 Quadratic Regularization Algorithms for Computed Imaging'. In ICIP.  
850
- 851 Meister, S., Hur, J., & Roth, S. (2018). 'UnFlow: Unsupervised Learning of Optical Flow with a  
852 Bidirectional Census Loss'. In AAAI, New Orleans, Louisiana, February.  
853
- 854 Shorten, C. & Khoshgoftaar, T.M. (2019). 'A Survey on Image Data Augmentation for Deep  
855 Learning'. *Journal of Big Data*, 6(60). doi: 10.1186/s40537-019-0197-0

856  
857 Ayhan, M.S. & Berens, P. (2018). 'Test-Time Data Augmentation for Estimation of  
858 Heteroscedastic Aleatoric Uncertainty in Deep Neural Networks'.  
859  
860 Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., Bao, B., Bell, P., Berard,  
861 D., Burovski, E., Chauhan, G., Chourdia, A., Constable, W., Desmason, A., DeVito, Z., Ellison,  
862 E., Feng, W., Gong, J., Gschwind, M., Hirsh, B., Huang, S., Kalambarkar, K., Kirsch, L., Lazos,  
863 M., Lezcano, M., Liang, Y., Liang, J., Lu, Y., Luk, C.K., Maher, B., Pan, Y., Puhrsch, C., Reso,  
864 M., Saroufim, M., Siraichi, M.Y., Suk, H., Zhang, S., Suo, M., Tillet, P., Zhao, X., Wang, E., Zhou,  
865 K., Zou, R., Wang, X., Mathews, A., Wen, W., Chanan, G., Wu, P., & Chintala, S. (2024). 'PyTorch  
866 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph  
867 Compilation'. In Proceedings of the 29th ACM International Conference on Architectural Support  
868 for Programming Languages and Operating Systems, Volume 2, ASPLOS '24, La Jolla, CA, USA,  
869 pp. 929–947. doi: 10.1145/3620665.3640366  
870  
871 Micikevicius, P., Narang, S., Alben, J., Diamos, G.F., Elsen, E., García, D., Ginsburg, B., Houston,  
872 M., Kuchaiev, O., Venkatesh, G., & Wu, H. (2017). 'Mixed Precision Training'. CoRR,  
873 abs/1710.03740.  
874  
875 Frank, J. & Al-Ali, L. (1975). 'Signal-to-Noise Ratio of Electron Micrographs Obtained by Cross  
876 Correlation'. Nature, 256, pp. 376–379. doi: 10.1038/256376a0  
877  
878 González-Ruiz, V., Fernández-Fernández, M.R. & Fernández, J.J. (2023). 'Structure-preserving  
879 Gaussian Denoising of FIB-SEM Volumes'. Ultramicroscopy, 246, 113674. doi:  
880 10.1016/j.ultramic.2022.113674  
881  
882 Harauz, G. & van Heel, M. (1986). 'Exact Filters for General Geometry Three Dimensional  
883 Reconstruction'. Optik, 73, pp. 146-156.  
884  
885 Verbeke, E.J., Gilles, M.A., Bendory, T. et al. (2024). 'Self Fourier Shell Correlation: Properties  
886 and Application to Cryo-ET'. Communications Biology, 7, 101. doi: 10.1038/s42003-023-05724-y  
887  
888 Scheres, S. & Chen, S. (2012). 'Prevention of Overfitting in Cryo-EM Structure Determination'.  
889 Nature Methods, 9, pp. 853–854. doi: 10.1038/nmeth.2115

890  
891 Liu, Y.T., Zhang, H., Wang, H. et al. (2022). 'Isotropic Reconstruction for Electron Tomography  
892 with Deep Learning'. *Nature Communications*, 13, 6482. doi: 10.1038/s41467-022-33957-8  
893  
894 Baumeister, W. & Grimm, R. (1999). 'Electron Tomography of Molecules and Cells'. *Trends in*  
895 *Cell Biology*, 9(2), pp. 81-85.  
896  
897 Medalia, O., Weber, I., Frangakis, A.S., Nicastro, D., Gerisch, G. and Baumeister, W., 2002.  
898 Macromolecular architecture in eukaryotic cells visualized by cryoelectron tomography. *Science*,  
899 298(5596), pp.1209-1213.  
900  
901 Gan, L. & Jensen, G.J. (2012). 'Electron Tomography of Cells'. *Quarterly Reviews of Biophysics*,  
902 45, pp. 27–56. doi: 10.1017/S0033583511000102  
903  
904 Wan, W. and Briggs, J.A., 2016. Cryo-electron tomography and subtomogram averaging.  
905 *Methods in enzymology*, 579, pp.329-367.  
906  
907 Mosalaganti, S., Obarska-Kosinska, A., Siggel, M., Taniguchi, R., Turoňová, B., Zimmerli, C.E.,  
908 Buczak, K., Schmidt, F.H., Margiotta, E., Mackmull, M.T. and Hagen, W.J., 2022. AI-based  
909 structure prediction empowers integrative structural analysis of human nuclear pores. *Science*,  
910 376(6598), p.eabm9506.  
911  
912 de Sautu, M., Herrmann, T., Scanavachi, G., Jenni, S. & Harrison, S.C. (2024). 'The Rotavirus  
913 VP5\*/VP8\* Conformational Transition Permeabilizes Membranes to Ca<sup>2+</sup>'. *PLOS Pathogens*,  
914 20(4), e1011750. doi: 10.1371/journal.ppat.1011750  
915  
916 Erdmann, P.S., Hou, Z., Klumpe, S. et al. (2021). 'In Situ Cryo-Electron Tomography Reveals  
917 Gradient Organization of Ribosome Biogenesis in Intact Nucleoli'. *Nature Communications*, 12,  
918 5364. doi: 10.1038/s41467-021-25413-w  
919  
920 Frangakis, A.S. et al. (2002). 'Identification of Macromolecular Complexes in Cryoelectron  
921 Tomograms of Phantom Cells'. *Proceedings of the National Academy of Sciences of the United*  
922 *States of America*, 99, pp. 14153–14158.  
923



- 924 Lučić, V., Förster, F. and Baumeister, W., 2005. Structural studies by electron tomography: from  
925 cells to molecules. *Annu. Rev. Biochem.*, 74, pp.833-865.  
926
- 927 Lehtinen, J. et al. Noise2Noise: Learning Image Restoration without Clean Data. In *Proc. 35th*  
928 *International Conference on Machine Learning* (eds Dy, J. & Krause, A.) 2965–2974 (PMLR,  
929 2018)  
930
- 931 Zeng, X., Ding, Y., Zhang, Y., Uddin, M.R., Dabouei, A. & Xu, M. (2024). 'DUAL: Deep  
932 Unsupervised Simultaneous Simulation and Denoising for Cryo-Electron Tomography'. *bioRxiv*,  
933 2024.03.02.583135. doi: 10.1101/2024.03.02.583135  
934
- 935 Li, H., Zhang, H., Wan, X., Yang, Z., Li, C., Li, J., Han, R., Zhu, P. & Zhang, F. (2022). 'Noise-  
936 Transfer2Clean: Denoising Cryo-EM Images Based on Noise Modeling and Transfer'.  
937 *Bioinformatics*, 38(7), pp. 2022–2029. doi: 10.1093/bioinformatics/btac052  
938
- 939 Turoňová, B., Schur, F.K.M., Wan, W. & Briggs, J.A.G. (2017). 'Efficient 3D-CTF Correction for  
940 Cryo-Electron Tomography Using NovaCTF Improves Subtomogram Averaging Resolution to 3.4  
941 Å'. *Journal of Structural Biology*, 199(3), pp. 187-195. doi: 10.1016/j.jsb.2017.07.007  
942
- 943 Abdelhakim AH, Salgado EN, Fu X, Pasham M, Nicastro D, Kirchhausen T, et al. Structural  
944 correlates of rotavirus cell entry. *PLoS Pathog.* 2014;10(9):e1004355.  
945
- 946 Zheng, S., Palovcak, E., Armache, JP. et al. MotionCor2: anisotropic correction of beam-induced  
947 motion for improved cryo-electron microscopy. *Nat Methods* 14, 331–332 (2017).  
948 <https://doi.org/10.1038/nmeth.4193>  
949
- 950 Kremer, J.R., Mastronarde, D.N. & McIntosh, J.R. (1996). 'Computer Visualization of Three-  
951 Dimensional Image Data Using IMOD'. *Journal of Structural Biology*, 116(1), pp. 71-76. doi:  
952 10.1006/jsbi.1996.0013  
953
- 954 Rohou, A. & Grigorieff, N. (2015). 'CTFFIND4: Fast and Accurate Defocus Estimation from  
955 Electron Micrographs'. *Journal of Structural Biology*, 192(2), pp. 216-221. doi:  
956 10.1016/j.jsb.2015.08.008  
957

- 958 Zheng, S., Wolff, G., Greenan, G., Chen, Z., Faas, F.G.A., Bárcena, M., Koster, A.J., Cheng, Y.,  
959 & Agard, D.A. (2022). 'AreTomo: An Integrated Software Package for Automated Marker-Free,  
960 Motion-Corrected Cryo-Electron Tomographic Alignment and Reconstruction'. *Journal of*  
961 *Structural Biology: X*, 6, 100068. doi: 10.1016/j.yjsbx.2022.100068  
962
- 963 Pettersen, E.F., Goddard, T.D., Huang, C.C., et al. (2021). 'UCSF ChimeraX: Structure  
964 Visualization for Researchers, Educators, and Developers'. *Protein Science*, 30, pp. 70–82. doi:  
965 10.1002/pro.3943  
966
- 967 Tang, G., Peng, L., Baldwin, P.R., Mann, D.S., Jiang, W., Rees, I., & Ludtke, S.J. (2007). 'EMAN2:  
968 An Extensible Image Processing Suite for Electron Microscopy'. *Journal of Structural Biology*,  
969 157(1), pp. 38-46. doi: 10.1016/j.jsb.2006.05.009  
970
- 971 Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R. & Carroll, R. (2018). 'Burst Denoising  
972 with Kernel Prediction Networks'. In 2018 IEEE/CVF Conference on Computer Vision and Pattern  
973 Recognition (CVPR), Salt Lake City, UT, USA, pp. 2502-2510. doi: 10.1109/CVPR.2018.00265  
974
- 975 Eisenstein, F., Yanagisawa, H., Kashihara, H. et al. (2023). 'Parallel Cryo Electron Tomography  
976 on In Situ Lamellae'. *Nature Methods*, 20, pp. 131–138. doi: 10.1038/s41592-022-01690-1  
977
- 978 Mastronarde, D.N. (2005). 'Automated Electron Microscope Tomography Using Robust Prediction  
979 of Specimen Movements'. *Journal of Structural Biology*, 152(1), pp. 36-51. doi:  
980 10.1016/j.jsb.2005.07.007  
981
- 982 Schur, F.K.M., Obr, M., Hagen, W.J.H., Wan, W., Jakobi, A.J., Kirkpatrick, J.M., Sachse, C.,  
983 Kräusslich, H-G. & Briggs, J.A.G. (2016). 'An Atomic Model of HIV-1 Capsid-SP1 Reveals  
984 Structures Regulating Assembly and Maturation'. *Science*, 353(6298), pp. 506-508. doi:  
985 10.1126/science.aaf9620  
986

**Table 1**

SNR for the images (in dB) and its ratio between raw and denoised data			
Source	Raw	Topaz-Denoise	CryoSamba
Rotavirus in BSC1 cell (1)	-32.4 (0.0)	-22.9 (+9.4)	-25.7 (+6.7)
Rotavirus in BSC1 cell (2)	-26.3 (0.0)	-16.4 (+9.9)	-16.3 (+10.0)
Yeast lamella (1)	-20.2 (0.0)	-14.7 (+5.5)	-14.3 (+5.9)
Yeast lamella (2)	-22.9 (0.0)	-16.9 (+6.0)	-17.3 (+5.6)
Yeast lamella (3)	-36.2 (0.0)	-34.4 (+1.8)	-34.2 (+2.0)
<b>Total average</b>	<b>-27.6 (0.0)</b>	<b>-21.1 (+6.5)</b>	<b>-21.6 (+6.0)</b>

**Table 2**

SNR (in dB) for different voxel resolutions							
	Even/Odd plane SNR				Denoised/Raw SNR		
	2.62 Å	7.86 Å	15.72 Å		2.62 Å	7.86 Å	15.72 Å
Raw	-1.07	1.74	-0.26		N/A	N/A	N/A
CryoCARE	-1.72	11.98	17.77		-28.73	-21.83	-22.95
Topaz	8.86	12.49	13.21		0.87	-1.08	0.38
CryoSamba	7.29	8.10	8.85		3.02	2.64	6.91

**Table 3**

Time and VRAM usage for two workstations		
	<b>DGX workstation (batch size 32, 8 GPUS)</b>	<b>Personal laptop (batch size 4, 1 GPU)</b>
Training VRAM (per GPU)	31 GB	6 GB
Training time (per epoch)	24 sec	742 sec (12 min 22 sec)
Training time (total: 30 epochs)	720 sec (12 min)	22260 sec (6 h 11 min)
Inference VRAM (per GPU)	8 GB	1 GB
Inference time (total)	95 sec (1 min 35 sec)	3973 sec (1 h 6 min 13 sec)

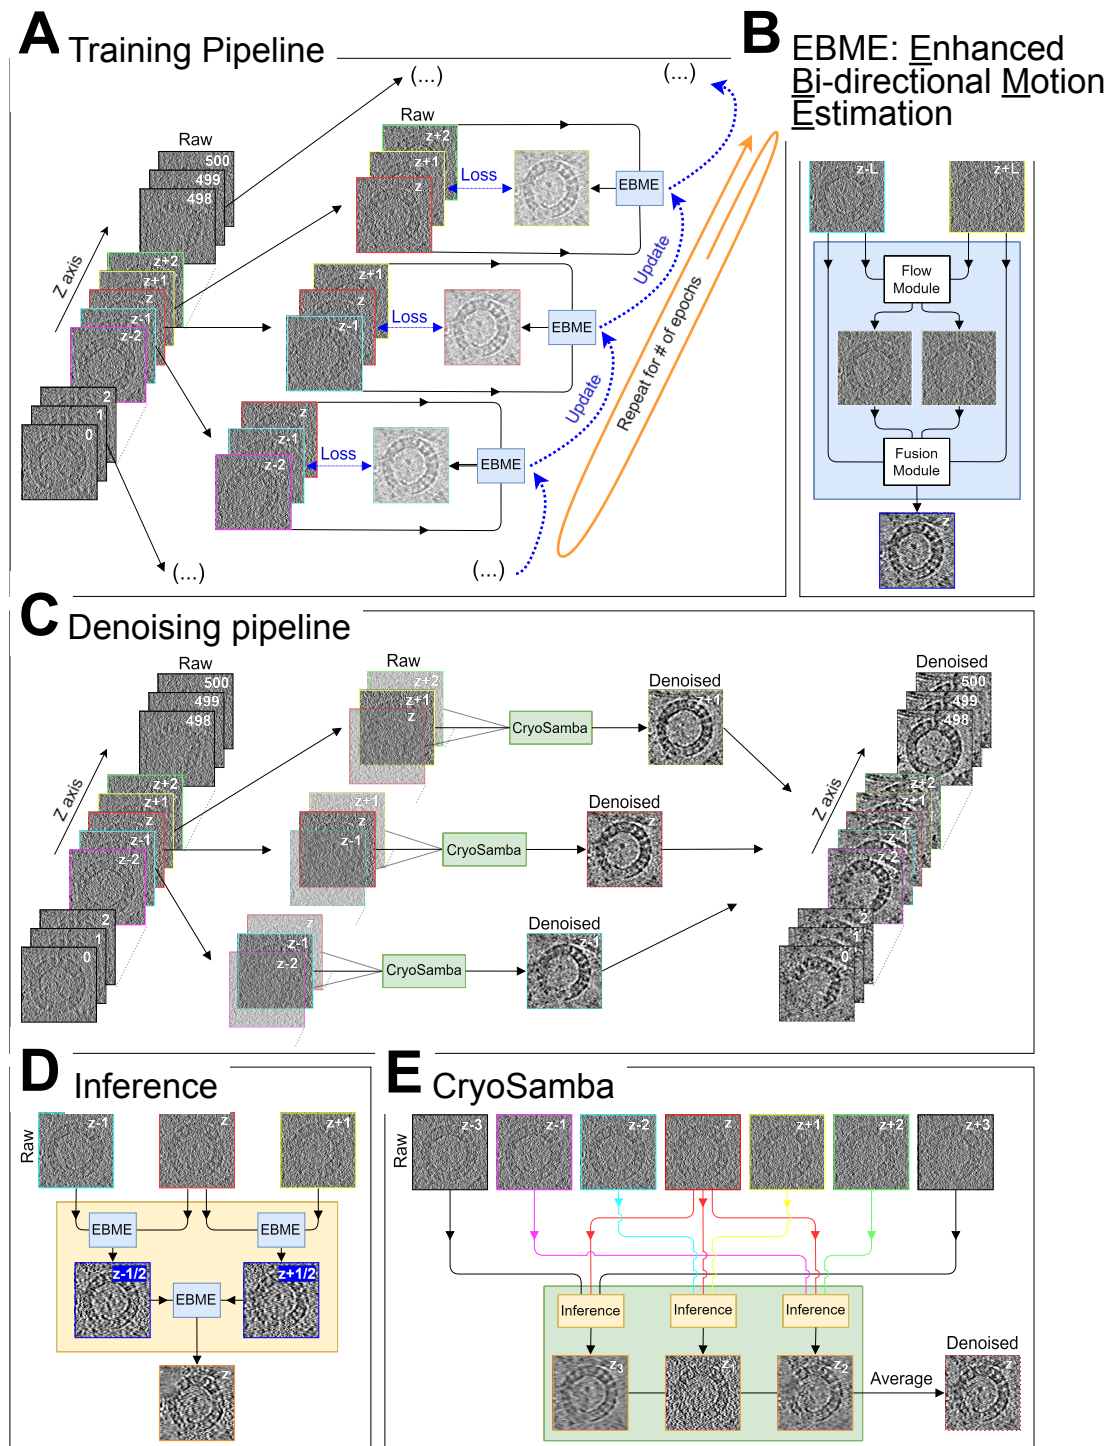


Fig. 1

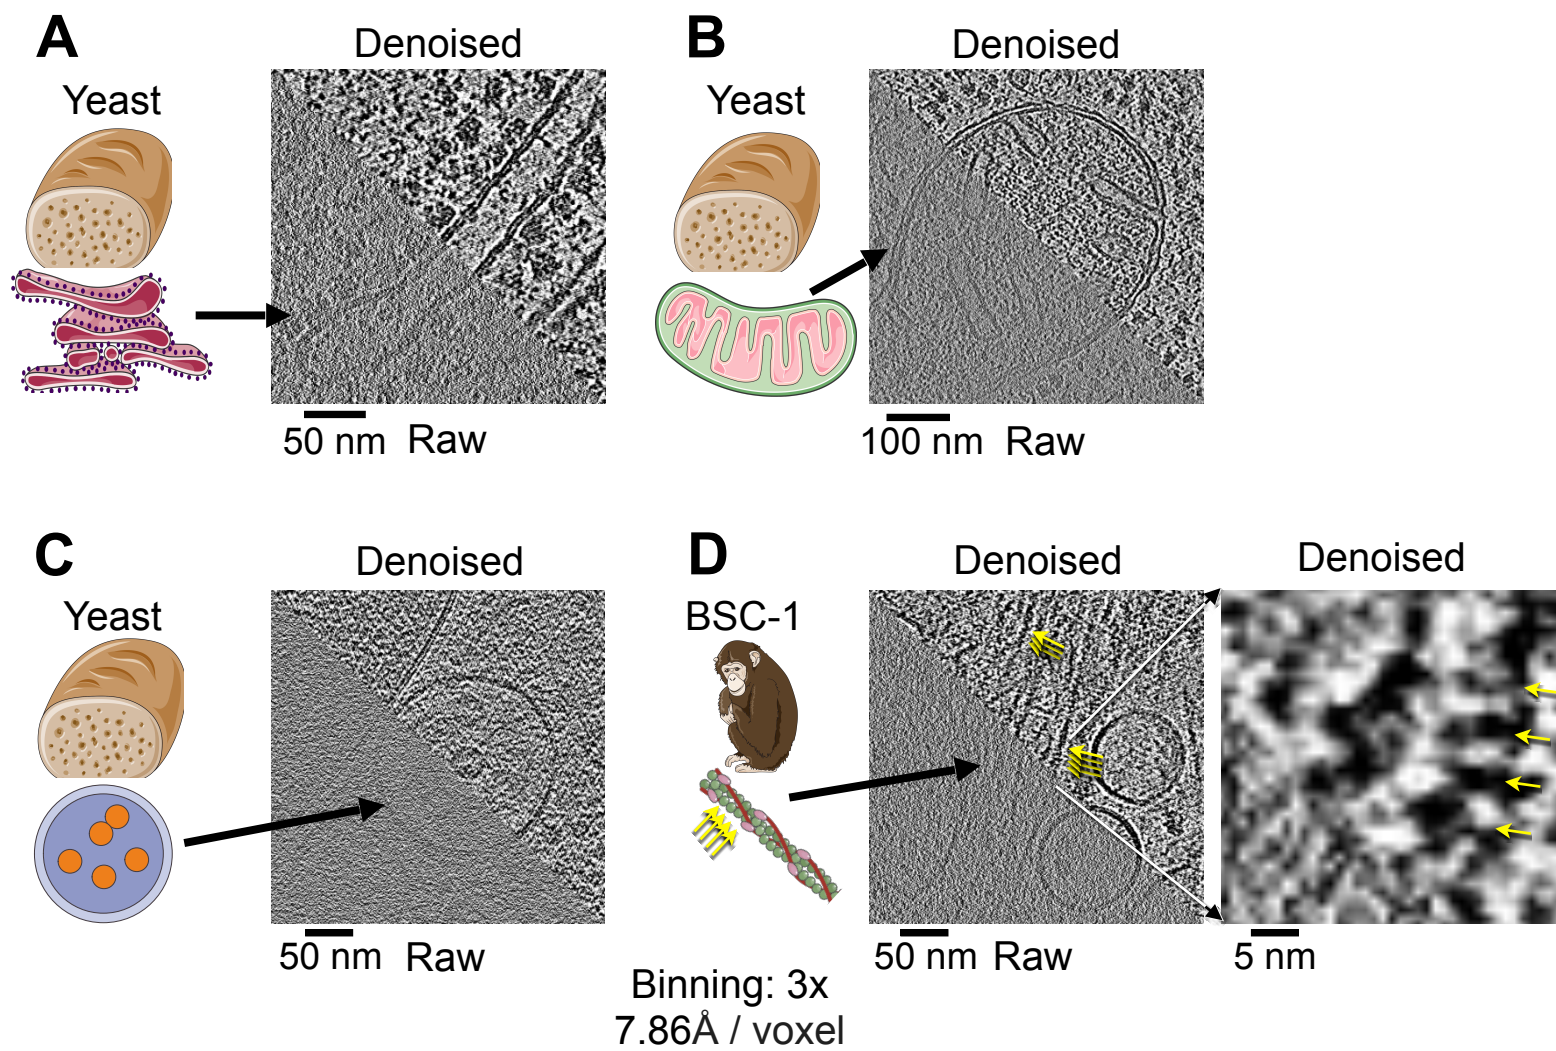


Fig. 2

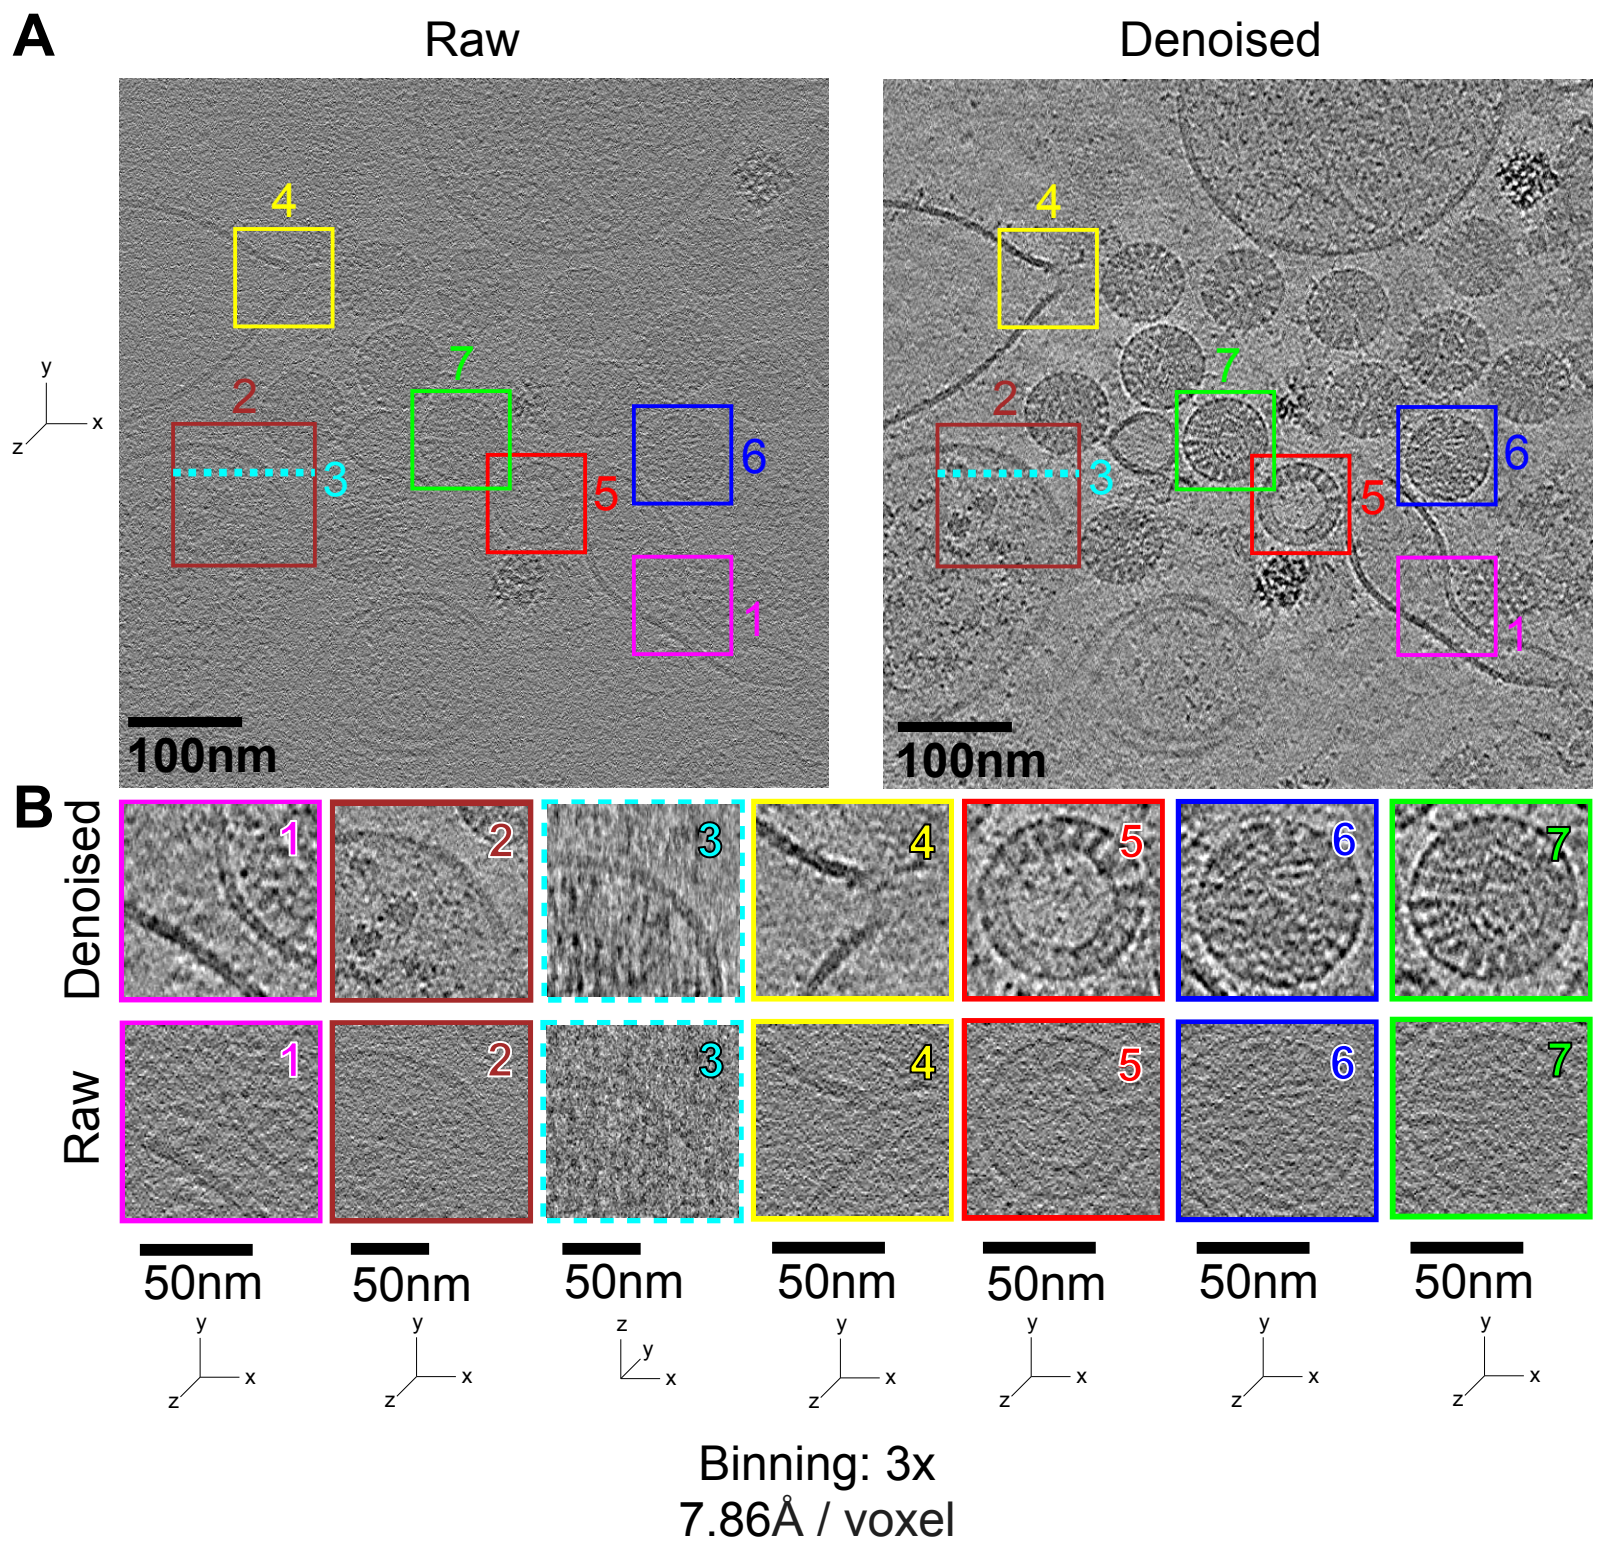


Fig. 3



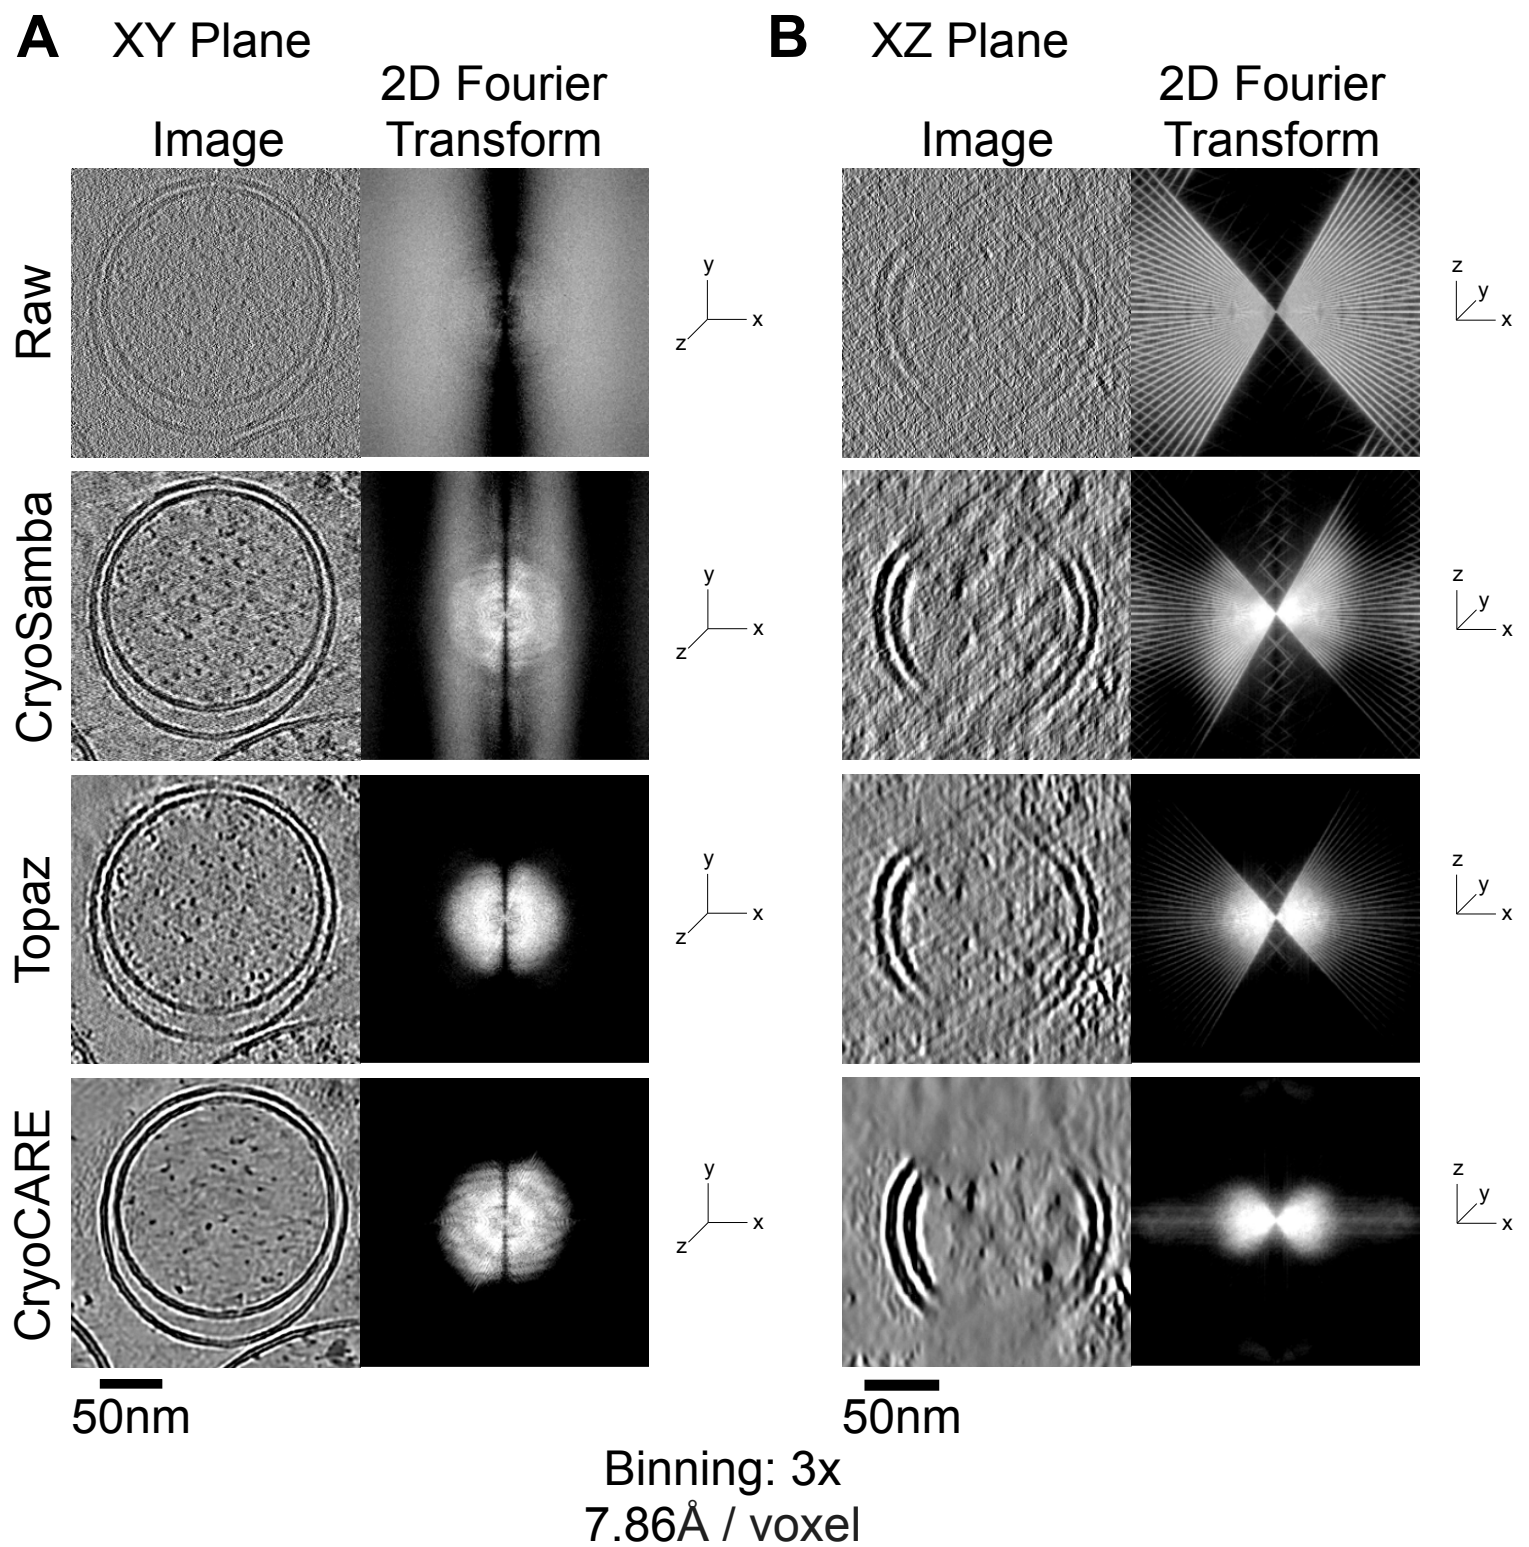


Fig. 4




Fig. 5

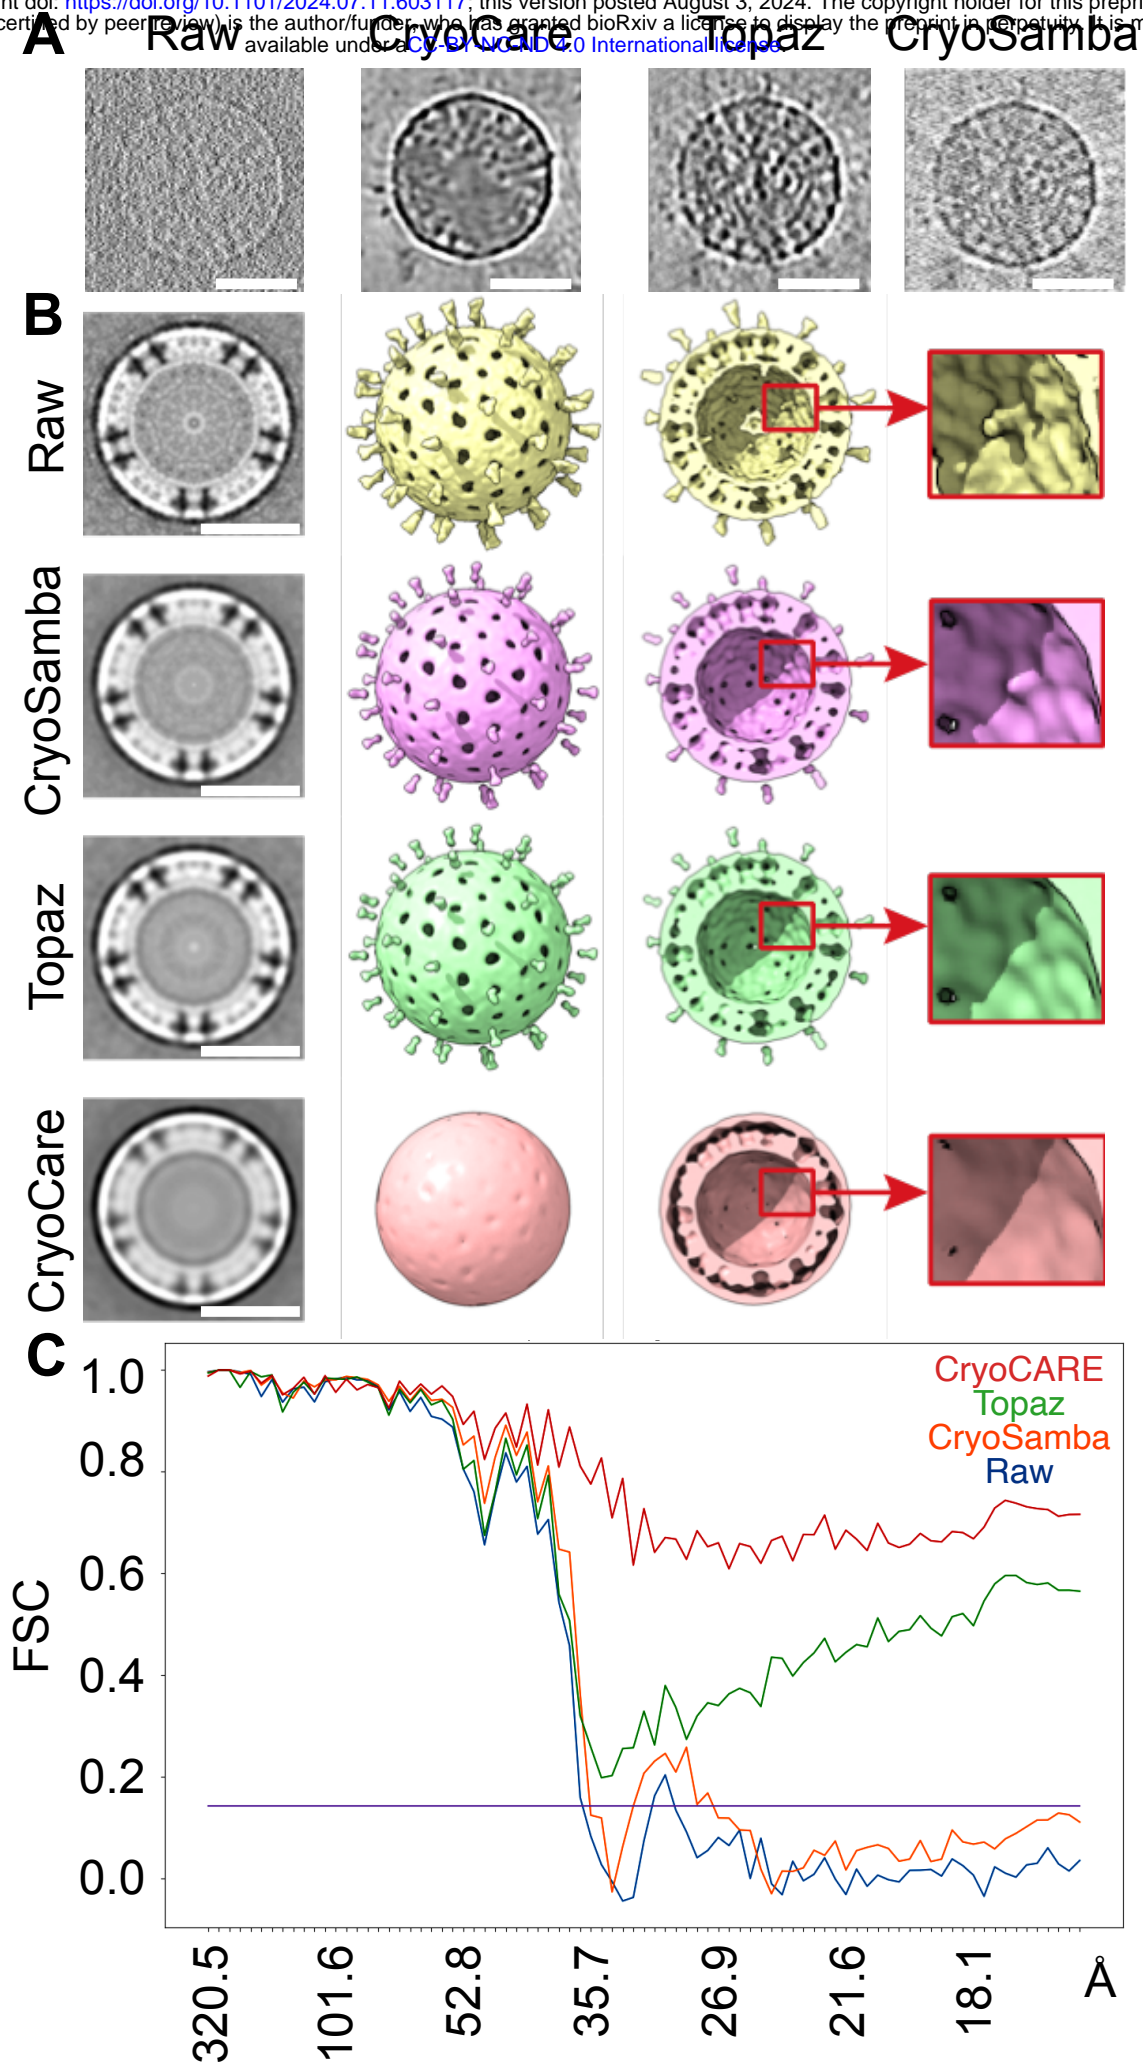




Fig. 6

### A 1 x Binning (2.62Å / voxel)



### B 3 x Binning (7.86Å / voxel)



### C 6 x Binning (15.72Å / voxel)

