

CRYO-CARE: CONTENT-AWARE IMAGE RESTORATION FOR CRYO-TRANSMISSION ELECTRON MICROSCOPY DATA*

Tim-Oliver Buchholz^{1,2}, Mareike Jordan², Gaia Pigino², Florian Jug^{1,2}

¹Center for Systems Biology Dresden (CSBD)

²Max-Planck Institute of Molecular Cell Biology and Genetics

ABSTRACT

Multiple approaches to use deep learning for image restoration have recently been proposed. Training such approaches requires well registered pairs of high and low quality images. While this is easily achievable for many imaging modalities, *e.g.* fluorescence light microscopy, for others it is not. Cryo-transmission electron microscopy (cryo-TEM) could profoundly benefit from improved denoising methods, unfortunately it is one of the latter. Here we show how recent advances in network training for image restoration tasks, *i.e.* denoising, can be applied to cryo-TEM data. We describe our proposed method and show how it can be applied to single cryo-TEM projections and whole cryo-tomographic image volumes. Our proposed restoration method dramatically increases contrast in cryo-TEM images, which improves the interpretability of the acquired data. Furthermore we show that automated downstream processing on restored image data, demonstrated on a dense segmentation task, leads to improved results.

Index Terms— image restoration, cryo-electron microscopy, deep learning, denoising

1. INTRODUCTION

Modern cryo-transmission electron microscopy (cryo-TEM) enables the observation of biological structures in their native state at high resolution. In order to prevent sample destruction during image acquisition, the total electron dose needs to be restricted [1]. This restriction results in noisy, low contrast acquisitions. In practice, electron microscopists typically acquire defocused images to trade resolution for increased contrast. Hence, the existence of better performing image restoration methods would enable image acquisitions at low electron dose with reduced defocus and therefore also at elevated resolution.

For fluorescence microscopy data, deep learning can be used for content-aware image restoration (CARE) [2]. Data for training CARE networks requires adequately imaged or synthetically generated pairs at low and high quality. The ideas presented in [2] do not translate to cryo-TEM data, where the before mentioned electron dose prevents the acquisition of non-noisy ground truth images.

Here we present cryo-CARE, a way to train content-aware restoration networks for cryo-TEM data. Cryo-CARE can be trained by using registered pairs of noisy images, an idea that was recently introduced in the context of real-world RGB and MRI images [3]. More concretely, we show how single TEM projections and whole tomographic volumes can be denoised using a strong, learned, and content-aware prior.

We compare our results to simple baseline methods such as median-filtering or NAD [4]. Despite their simplicity, these methods are widely used by cryo-TEM experts to improve the interpretability of their data. Furthermore we show that automated downstream processing on restored image data leads to significantly improved results.

2. APPROACH AND METHODS

Since cryo-TEM acquisitions suffer from very low signal-to-noise ratio (SNR) due to the limited electron dose, recording high SNR ground truth images is not possible. Recently it was demonstrated that image denoising networks can be trained without having high SNR images available. Instead, only well registered pairs of low SNR images are required [3]. Hence, for cryo-TEM data a combination of CARE and the ideas from [3] enable us, for the first time, to train cryo-CARE networks.

Network Architecture and Training Procedure: For all experiments we used a U-Net [5] of depth two, a convolution kernel size of three, and a linear activation function at the last layer. Moreover we used a per-pixel mean squared error loss. In all experiments we keep 10% of extracted training patches (see below) as validation set. Note that we use the open-source CARE framework [2] for all experiments.

2.1. Restoring Single cryo-TEM Projections (P2P)

Here we describe three ways to train cryo-CARE networks on adequately prepared pairs of cryo-TEM projections. Hence, we call this approach PROJECTION2PROJECTION or P2P.

P2P-*ip*, *training using acquired image pairs:* The most straight forward way to combining CARE [2] and NOISE2NOISE [3] is to acquire pairs of images for which the noise is independent.

* This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.




Fig. 1. Cryo-CARE results on a 2D cryo-TEM projection. Subfigures and insets show: raw input data (a), median filtered restoration baseline (b), Cryo-CARE results when trained on tomographic tilt-angle pairs (c), on acquired image pairs (d), and on dose-fractionated movie frames (e).

To this end, we acquired such image pairs on a 300 kV Thermo Fisher cryo-TEM Titan Halo that is equipped with a K2 direct electron detector from Gatan. More precisely, we acquired images of *Chlamydomonas reinhardtii* cilia in dose-fractionation mode (movie mode) [6], splitted the frames in two halves, and averaged them without additional alignment, giving us the equivalent of two independently acquired images at half the available electron dose each¹. From such pairs of images we extracted 1000 randomly selected patch-pairs of size 128×128 which are used to train a CARE network in the NOISE2NOISE regime. After training, we use the trained network to restore all image pairs and retrieve the final result by per-pixel averaging the two individual restorations (see Fig. 1(d)).

P2P-tap, using tomographic tilt-angle pairs: For readily acquired, not dose-fractionated data, the previously described scheme cannot be applied. Archived data for which only single acquisition exist can therefore not be used for training cryo-CARE networks. For existing tilt-series, acquired for tomographic reconstruction, we asked ourselves if pairs of neighboring tilt-angles could be used for training. We used IMOD [7] to align and register all acquired tilt-angles. As be-

¹Note that each image in such a pair has an even lower SNR due to the halved electron dose.

fore, training was performed on 1000 randomly selected patch pairs of size 128×128 taken from adjacent tilt-angle projections. Final restorations are retrieved by applying the trained network to both tilt-angles individually (see Fig. 1(c)).

P2P-df, using dose-fractionated movie frames: Since our data was acquired on a Gatan K2 direct detector, we were able to go an additional step further. Instead of using two acquired images, as described initially, we can leverage the fact to have many more frames acquired. As it is usually done during dose-fractionation, we can additionally correct for motion-blur of the sample by registering the individual frames using MotionCor2 [8]. We then sum all even and odd frames to retrieve two images with independent noise. This interleaved frame-splitting is advantageous because induced beam damage will be equally shared in both independent images. Again we trained on 1000 randomly selected patch pairs of size 128×128 , and created the final restored projection by applying the network to both images followed by per-pixel averaging (see Fig. 1(e)).

2.2. Restoration of cryo-TEM Tomograms (T2T)

Also here we present three ways to restore cryo-TEM tomograms. One obvious way is to first use P2P cryo-CARE to restore all tilt-angles individually and then reconstruct a tomogram from them. All tomographic reconstructions were performed with ETOMO, which is part of IMOD [7]. We will later see that this leads to clearly visible artifacts that can be mitigated by first reconstructing two tomograms, each from half the available data, and then train a 3D cryo-CARE network in a NOISE2NOISE regime. We call these approaches collectively TOMO2TOMO, or T2T in short.

T2T-*eea*, using even-odd acquisitions: This protocol is designed to work for conventionally acquired tilt-series, when no direct detector is available. Here we split all tilted projections in two sets based on their acquisition number. From all tilt angles with an even/odd acquisition number, we reconstructed two data-independent tomograms and used those to train a 3D cryo-CARE network on 1200 randomly selected 3D sub-volumes of size $64 \times 64 \times 64$. We created the final restored tomogram by applying the trained network to both tomograms followed by per-pixel averaging (see Fig. 2(d)).

T2T-df, using dose-fractionated movie frames: In case the available data was acquired in dose-fractionation mode (movie mode), we propose a slightly different protocol. For each tilt-angle, similar to our P2P approach on dose-fractionated data, we can equally split all frames, align and sum them. The two sets of independent tilt-angle projections can then be used to reconstruct two independent tomograms. We trained as before and created the final restored tomogram by applying the trained network to both tomograms followed by per-pixel averaging. The advantage of this approach is that the angular sampling for both tomograms is denser and consistent, hence leading to better results (see Fig. 2(d)).

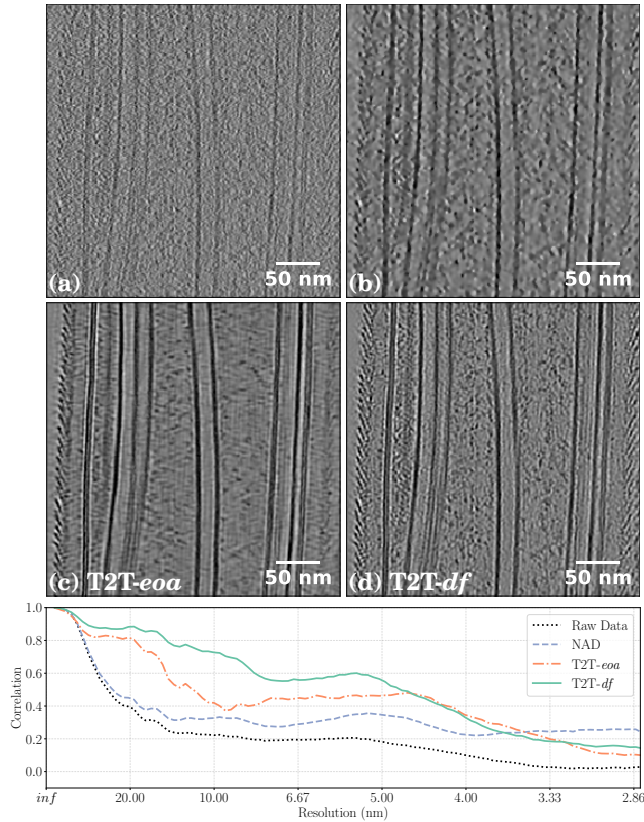


Fig. 2. Cryo-CARE results on a 3D cryo-TEM tomogram. Subfigures show: a section through the raw tomogram (a), the non-linear anisotropic diffusion filtered baseline (b), cryo-CARE results when trained via our proposed T2T-eoa (c) and T2T-df (d). The graph shows the Fourier shell correlation (FSC) curves of the raw tomogram, the baseline, and our proposed methods.

2.3. Automated Downstream Analysis

We train a dense segmentation and detection workflow for *Chlamydomonas reinhardtii* outer dynein arm (ODA). Also here we use a U-Net [5] trained on manually generated and with PEET [9, 10] refined ground truth. The predicted segmentation is then normalized and Otsu thresholded [11]. Each connected component is then filtered according to its size in voxels. Remaining components are treated as one detected ODA. Since we had only one hand annotated tomogram, we split annotations in train and test set, counting 383 and 712, respectively. We did not use data augmentation.

3. RESULTS

We evaluated all variations of our proposed methods on two datasets. The first one, called TOMO110, was acquired by ourselves on a 300 kV Thermo Fisher cryo-TEM Titan Halo with a Gatan K2 direct electron detector. This enables us to

test cryo-CARE variations that require image acquisitions in dose-fractionation mode (movie mode). The second dataset, EMPIAR-10110², is publicly available via the EMPIAR database [12] and consists of a complete tomographic series of tilted projections. We did not find publicly available raw dose-fractionated data but will make our TOMO110 data available upon publication.

P2P results are computed on unbinned data and T2T results are computed on six times binned data.

3.1. P2P Restoration Experiments

We first tested P2P-tap, which is applicable to both available datasets (see Fig. 1(c) and Fig. 3(c)). Using adjacent tilt-angles pairs during training leads to restored images that appear blurry. This is due to the relative displacement of individual structures in the imaged volume when projected at adjacent tilt-angles.

To circumvent this problem, we train cryo-CARE networks on specifically acquired pairs of images (P2P-ip) and dose-fractionated movie frames (P2P-df). Restorations using these networks are shown in Fig. 1(d,e).

Tomographic reconstruction from P2P restored tilt-angles: A canonical idea to reconstruct denoised tomograms is to use restored P2P-df tilt angles (like in Fig. 1(e)). This does, unfortunately, amplify the missing wedge artifacts at high-gradient locations (see Fig. 4). Since neural networks are complex non-linear filters and tilt-angle reconstructions are performed independently, the predicted intensities for a given structure is not necessarily consistent across restored tilt-angles. These inconsistent amplitudes are likely the reason for the amplification of the observed missing wedge artifacts. Nevertheless, this problem can be addressed with the T2T network training regimes described in Section 2.

3.2. T2T Restoration Experiments

In Fig. 2 we compare T2T even-odd acquisitions and T2T dose-fractionated to the reconstructed raw tomogram and a




Fig. 3. Cryo-CARE restoration on the publicly available EMPIAR-10110 dataset. (a) Raw projection (single tilt angle). (b) Median filtered baseline. (c) Our P2P-tap results. (d) Raw tomogram. (e) NAD filtered baseline. (d) Our even-odd T2T restoration.

²<http://dx.doi.org/10.6019/EMPIAR-10110>




Fig. 4. Tomogram reconstruction artifacts. Tomograms reconstructed from P2P restored tilt-angles lead to strong missing-wedge artifacts (a). This problem is reduced using our proposed T2T training scheme (b).

non-linear anisotropic diffusion filtered (NAD) baseline. For a quantitative comparison we computed the Fourier shell correlation (FSC) for all reconstructions. We computed the FSC for the raw data directly from the two tomograms we created for T2T dose-fractionated training. For the NAD filtered tomogram and our T2T dose-fractionated results, we computed the FSC on the filtered and restored data, respectively. For the T2T even-odd acquisitions we split, as commonly done, the available tilt-angles in order to compute the FSC. The slices we show in Fig. 2(a-d) allow a qualitative assessment of the power of cryo-CARE, while the plotted FSC curves quantify the gained image quality. Still, the question remains if cryo-CARE is improving the automated analysis of acquired cryo-TEM data.

3.3. Cryo-CARE Facilitates Automated Analysis

In order to test if the improved SNR and contrast in restored cryo-TEM tomograms benefits automated downstream analysis, we implemented a segmentation and detection workflow (described in Section 2.3). Fig. 5 shows results of the exact same automated analysis performed on raw tomograms and on cryo-CARE restored data. In Fig. 5 we can appreciate a significant increase in precision and recall when data is first restored and then analyzed.

4. DISCUSSION

In this publication we show how content-aware image restoration can successfully be applied to cryo-TEM data. EM experts are currently using relatively simple filtering techniques, *i.e.* NAD, before manually investigating acquired data. Cryo-CARE, as we have shown, leads to highly contrasted and well resolved 2D and 3D data. Our experiments also show that P2P reconstructions are not ideal for tomographic reconstructions. Nevertheless, with T2T we can offer a simple and powerful tool for content-aware tomographic restorations. We therefore believe that cryo-CARE will facilitate manual data browsing, a step that can hardly be underestimated when many and/or large volumes have to be browsed for regions of interest.

Additionally we showed that cryo-CARE restorations can lead to highly improved automated analysis results. An essen-

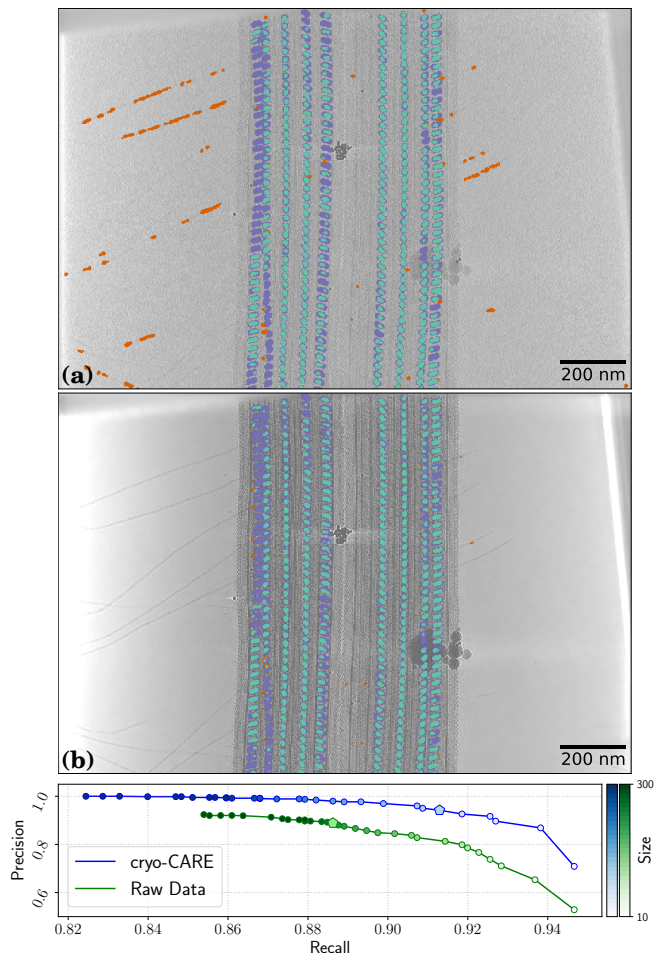


Fig. 5. Automated downstream analysis on raw data (a) and a T2T-*df* restored tomogram (b). Ground-truth voxels are shown in violet, true-positives in turquoise, and false-positives in orange. Precision-recall plots on increasing segment size threshold (see main text) are shown below. The pentagons correspond to sub-figures (a) and (b).

tial feature of cryo-CARE is that training data can be generated by the microscope itself and does not require tedious human labeling. While end-to-end pipelines on raw data might need huge amounts of labeled data to also co-learn to restore the noisy data, cryo-CARE helps to uncouple these two tasks – a preprocessing step that does not need human labels and an analysis stage that is likely to require lesser amounts of training data.

We are confident that cryo-CARE will rapidly find application in the cryo-EM field. It improves data-browsing, creates well contrasted, high SNR images for improved visualization of single projections/tomograms, and improves the performance of automated analysis pipelines, hence it enables to work efficiently on much larger bodies of data.

5. REFERENCES

- [1] E Knapek and J Dubochet, “Beam damage to organic material is considerably reduced in cryo-electron microscopy,” *Journal of molecular biology*, vol. 141, no. 2, pp. 147–161, 1980.
- [2] Martin Weigert, Uwe Schmidt, Tobias Boothe, M Andreas, Alexander Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Cullley, et al., “Content-aware image restoration: pushing the limits of fluorescence microscopy,” *bioRxiv*, p. 236463, 2017.
- [3] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila, “Noise2noise: Learning image restoration without clean data,” *arXiv preprint arXiv:1803.04189*, 2018.
- [4] Achilleas S Frangakis and Reiner Hegerl, “Noise reduction in electron tomographic reconstructions using nonlinear anisotropic diffusion,” *Journal of structural biology*, vol. 135, no. 3, pp. 239–250, 2001.
- [5] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [6] Xueming Li, Paul Mooney, Shawn Zheng, Christopher R Booth, Michael B Braunfeld, Sander Gubbens, David A Agard, and Yifan Cheng, “Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-em,” *Nature methods*, vol. 10, no. 6, pp. 584, 2013.
- [7] James R Kremer, David N Mastronarde, and J Richard McIntosh, “Computer visualization of three-dimensional image data using imod,” *Journal of structural biology*, vol. 116, no. 1, pp. 71–76, 1996.
- [8] Shawn Zheng, Eugene Palovcak, Jean-Paul Armache, Yifan Cheng, and David Agard, “Anisotropic correction of beam-induced motion for improved single-particle electron cryo-microscopy,” *bioRxiv*, 2016.
- [9] Daniela Nicastro, Cindi Schwartz, Jason Pierson, Richard Gaudette, Mary E. Porter, and J Richard McIntosh, “The molecular architecture of axonemes revealed by cryoelectron tomography,” *Science*, vol. 313, no. 5789, pp. 944–948, 2006.
- [10] John M. Heumann, Andreas Hoenger, and David N. Mastronarde, “Clustering and variance maps for cryo-electron tomography using wedge-masked differences,” *J Struct Biol*, vol. 175, no. 3, pp. 288–299, 2011.
- [11] Nobuyuki Otsu, “A threshold selection method from gray-level histograms,” *IEEE transactions on systems, man, and cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [12] Andrii Iudin, Paul K Korir, José Salavert-Torres, Gerard J Kleywegt, and Ardan Patwardhan, “Empiar: a public archive for raw electron microscopy image data,” *Nature methods*, vol. 13, no. 5, pp. 387, 2016.