

Research Article

Subtomogram averaging for biophysical analysis and supramolecular context

Lauren Ann Metskas^{a,b,*}, Rosalie Wilfong^a, Grant J. Jensen^{c,d,*}^a Department of Biological Sciences, Purdue University, United States^b Department of Chemistry, Purdue University, United States^c Division of Biology and Biological Engineering, California Institute of Technology, United States^d Department of Chemistry and Biochemistry, Brigham Young University, United States

ARTICLE INFO

Edited by "Bauke W. Dijkstra"

Keywords:

Cryo-electron tomography
Subtomogram averaging

ABSTRACT

Recent advances in hardware, software and computing power have led to increasingly ambitious applications of cryo-electron tomography and subtomogram averaging. It is now possible to reveal both structures and biophysical relationships like protein binding partners and small molecule occupancy in these experiments. However, some data processing choices require the user to prioritize structure or biophysical context. Here, we present a modified subtomogram averaging approach that preserves both capabilities. By increasing the accuracy of particle-picking, performing alignment and averaging on all subtomograms, and decreasing reliance on symmetry and tight masks, the usability of tomography and subtomogram averaging data for biophysical analyses is greatly increased without negatively impacting structural refinements.

Introduction

Over the past decade subtomogram averaging (STA) has been used to assess not only the structure of proteins of interest, but also the architecture of their supramolecular assemblies. This is possible because STA provides position and orientation information for each analyzed sub-volume. With the advent of high-resolution data processing (Schur et al., 2016), STA is providing increasingly precise information, allowing for robust analyses of *in situ* drug occupancy (Schur et al., 2016; Tegunov et al., 2021), binding interactions (Metskas et al., 2022), and lattice organization (Mattei et al., 2016; Qu et al., 2021). These studies show the potential for STA results to be used for *in vivo* K_D calculations, simulations of small molecule diffusion, and protein-protein interaction studies in complex environments - here referred to as biophysical context.

STA has been extensively discussed elsewhere, with many protocol papers and software programs available for structure determination approaches (Burt et al., 2021; Chen et al., 2019; Himes and Zhang, 2018; Obr et al., 2022; Tegunov et al., 2021; Turoňová et al., 2017; Zivanov et al., 2022). However, many STA projects have an additional goal of biophysical analysis, which is typically addressed following structural analysis and using the structural data output. Decisions made during

structure-focused STA can degrade the quality of biophysical data. For example, discarding less-ideal particles during refinement can lead to incomplete analysis of the full dataset, extensive per-particle fitting can blur relationships between particles, and refinement of a particular conformation can result in poorer fitting and higher errors in positions and orientations of other conformations. Many of these effects are stronger when proteins of interest are not arranged in a regularly-spaced lattice.

Here, we describe an STA approach designed to maximize organizational information at the supramolecular level. A typical structure-focused STA approach picks particles with substantial error rates, and then proceeds to refine the resolution using a subset of the best-aligning particles; here, we present a context-focused STA pipeline that maximizes the quality of the initial tomograms, uses a high-fidelity approach to particle picking, and then treats particles equally during alignment and averaging. In this way, location and orientation of each particle of interest is refined to the best ability of the full, unclassified dataset, and the output dataset remains compatible with structure-focused STA approaches to allow later resolution refinement.

* Corresponding authors at: Department of Biological Sciences, Purdue University, United States (L.A. Metskas). Department of Chemistry and Biochemistry, Brigham Young University, United States (G.J. Jensen).

E-mail addresses: metskas@purdue.edu (L.A. Metskas), grant_jensen@byu.edu (G.J. Jensen).

<https://doi.org/10.1016/j.yjsbx.2022.100076>

Received 17 August 2022; Received in revised form 11 October 2022; Accepted 17 October 2022

Available online 18 October 2022

2590-1524/© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Test case: Rubisco in carboxysomes

In our recent work studying the alpha-carboxysome of *H. neapolitanus* (CB), we sought to characterize the packing of Rubisco complexes inside the microcompartment (Metskas et al., 2022). Our initial attempt at the analysis used a traditional structure-focused approach to STA, and went to an estimated 4.5 Å resolution (Fig. 1, structure-focused STA pipeline). We then attempted to use the Rubisco positions and orientations for biophysical context analyses, including occupied volume, binding, and nearest-neighbor alignment.

We began with a typical structure-focused particle picking approach, which resulted in a dataset with roughly 15–20 % false negative and false positive rates (Fig. 2). The final average was calculated using the best-correlating two-thirds of the particles. This approach is typical and is not considered problematic for structure determination: misidentified particles align poorly and are removed from consideration, while false negatives merely decrease total particle counts. However, biophysical context depends upon correct and complete particle identification as well as equal accuracy in the alignment of identified particles. Analysis of Rubisco packing using the results of the structure-focused STA pipeline yielded a number of biophysical results that were inconsistent with both visual inspection of the data and literature estimates (Dai et al., 2018; Iancu et al., 2010) (Fig. 3).

We refined our approach to STA to increase the accuracy of our biophysical analyses. Errors in particle picking were brought down to an estimated 2 % compared to manual picking, and we avoided symmetry, classification, tight masking and cross-correlation thresholding in the initial STA. These refinements allowed all particles to attain the best alignment according to the average of all data regardless of the presence of binding partners, alternative conformations, or other variability in the dataset; otherwise, alignment could favor a subset of particles leading to conformation-specific biases. Once the initial alignment was completed to the best resolution possible for the full dataset, subsets of particles were aligned and averaged further to target specific behaviors, without injuring the alignment of less-common species.

For the Rubisco test case, these changes had little effect on initial structure determination (Fig. 1) but strong, statistically significant differences in the biophysical analysis (Fig. 3, statistics in Supplemental Table 1). Rubisco in carboxysomes polymerizes into fibrils, with limited concentration dependence (Metskas et al., 2022); analyzing this behavior thus requires accurate carboxysome volumes and linkage

determination between bound Rubisco particles. Both of these analyses had extensive errors with a structure-focused STA pipeline. Using the Rubisco to estimate occupied volume was highly susceptible to the inclusion of false positives outside the CB shell. Many volume and count estimates from the structure-focused pipeline were therefore larger than previous measurements of purified CB volumes, while the context-focused pipeline was more consistent (Dai et al., 2018; Sun et al., 2022). Conversely, false negative identifications and poorer alignment in subclasses had the strongest effect on measuring the number of subunits in a fibril because missing or misaligned subunits interrupted chain lengths (Fig. 3A). Analyses that combine volume and protein–protein interaction, such as binding curves, were severely affected (Fig. 3C).

In general, high false positive and negative rates result in greater errors for simple and local analyses, but do not obscure trends (Fig. 3A). Nearest-neighbor analyses of relative orientation between Rubisco particles show this well (Fig. 3B). In these plots, each line is a trace of the histogram of pairwise angles between near Rubisco complexes within one CB, and the color is scaled according to the concentration of Rubisco complexes in the CB (excluding those closest to the shell, which have unique properties). There is a strong increase in alignment of near neighbors as the concentration of Rubisco within the CB rises (right, color scale). With less-accurate particle picking and alignment of subclasses, the local effects are preserved but the analysis degrades outside this radius (left, loss of color scaling but preservation of the mode value).

A Context-Focused Subtomogram averaging pipeline

Most STA projects follow a similar approach: following data acquisition, frames are aligned, defocus is estimated, and dose weighting is performed to generate a tilt stack. The tilt stack is aligned and the tomogram is reconstructed with CTF correction, typically using weighted back projection (SIRT-like filters or denoised tomograms are frequently used for particle picking). Particles are identified and divided into half-sets, then aligned and averaged until the resolution plateaus.

Structural processing often begins with initial particle picking performed on a lower-quality tomogram reconstruction, followed by local refinement of frame alignment, CTF correction, and other parameters iteratively and automatically (Chen et al., 2019; Tegunov et al., 2021). However, lower-quality initial tomograms can make accurate particle selection more difficult. We therefore make the initial reconstruction as high a quality as possible, facilitating particle selection and allowing us




Fig. 1. Subtomogram averaging comparison of structure-focused and context-focused STA pipelines. A: 1 nm slab views of the final post-processed map and 1 SVD crystal structure. Structure-focused pipeline, grey; context-focused pipeline, blue. B: Fourier Shell Correlation curves for mask-corrected and loose mask maps for both datasets. The subtomogram averaging resolutions and map quality are comparable within sampling error. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 2. Comparison of structure-focused and context-focused particle picking. White markers are particles present in both datasets. Blue markers, false negatives; red markers, false positives for structure-focused picking. 3.5 nm orthoslices, only particles centered in the 10 pixels above the central plane are marked. A: Top view. B: Center view. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

to perform STA and biophysical analyses directly from the initial reconstruction. In the case of our Rubisco work, we were able to achieve 4.5 Å resolution without post-processing based refinements of the tomograms and with inclusion of 100 % of the particles (Fig. 1).

This approach may initially sacrifice resolution in order to gain biophysical context; for example, a recent structure-focused STA approach for Rubisco in CBs reached similar structural conclusions while attaining a higher resolution than our test case (Ni et al., 2022; a comparison of the papers can be found in Metskas et al., 2022). However, the drastic changes we found in the biophysical analyses emphasize the importance of context-focused STA if non-structural information is also wanted (Fig. 3). Furthermore, because context-focused STA does not discard data, the results can be subsequently sent to structure-focused pipelines for higher resolution processing.

Our refinements to a typical STA pipeline include:

Start with high quality data

Optimize sample preparation

It is critical to choose the sample that will provide the most data with the thinnest ice and least contamination - both quality and quantity impact biophysical analyses. In our test case, we needed intact CBs to calculate volume; therefore, we purified the CBs to ensure thin ice with full data inclusion. Because purification can perturb samples, we add *in vivo* data collection to confirm major findings.

Grid freezing can be optimized similarly to single-particle approaches, with buffers adjusted to reduce sample aggregation and osmolarity shifts during blotting. We use 10 nm gold fiducials for tilt stack alignment, which are buffer-exchanged into the sample buffer to avoid sample effects from mixing. We also choose thin-carbon grids, which seem to allow better initial tilt stack alignment and modestly reduce the impact of carbon tilting into the field of view.

Prioritize data quality in the tilt stack acquisition

Data were collected with dose-symmetric tilting (Hagen et al., 2017), optimizing acquisition parameters for the behavior of our microscope in terms of drift, defocus estimation, aperture placements, etc. While STA

compensates for incomplete tilt angle coverage (the missing wedge effect), particle identification occurs before this step. Therefore, we chose to extend the tilt angle collection range to 66 degrees in an effort to modestly decrease the missing wedge effect on protein localization. This increased the time per tomogram by roughly-five minutes, but had no detrimental effect on tomograms because the extra dose is at the extra angles. This approach has decreasing benefits with increasing ice thickness, and was not used for cellular tomography.

Optimize the initial tomogram reconstruction

Perform dose-weighting prior to tilt stack alignment

Dose-weighting can be done at many stages in the tomogram reconstruction pipeline, but for best effect should be done prior to tilt stack alignment and after CTF estimation. While the primary purpose of dose-weighting is to limit the contributions of Fourier space containing the most electron damage (Grant and Grigorieff, 2015), when combined with a dose-symmetric tilt scheme it effectively functions as a low-pass filter in the noise-ridden high tilts (Fig. 4). This improves overall alignment for these tilts, along with a subtle improvement in signal to noise ratio in the reconstructed tomogram.

Fix large defocus estimation errors

Defocus estimation by CTF fitting is often poor at high tilt, particularly for software initially designed for single particle averaging such as ctfind4 (Rohou and Grigorieff, 2015) (Fig. 5A). Because the high tilts do not contribute high-resolution information due to thickness and electron dose, in practice this is not detrimental unless the effect is severe. Therefore, poor fits can be automatically corrected by replacing the erroneous estimate with a weighted average of the nearest good estimates (Fig. 5B). The difference between the weighted average and the true value will be low enough to be tolerated, provided that deviations beyond ~ 500 nm in defocus are not typical in the dataset (generally true for high-quality data acquisition). The defocus estimation can be further refined in per-particle iterative fitting if the dataset reaches high resolutions.

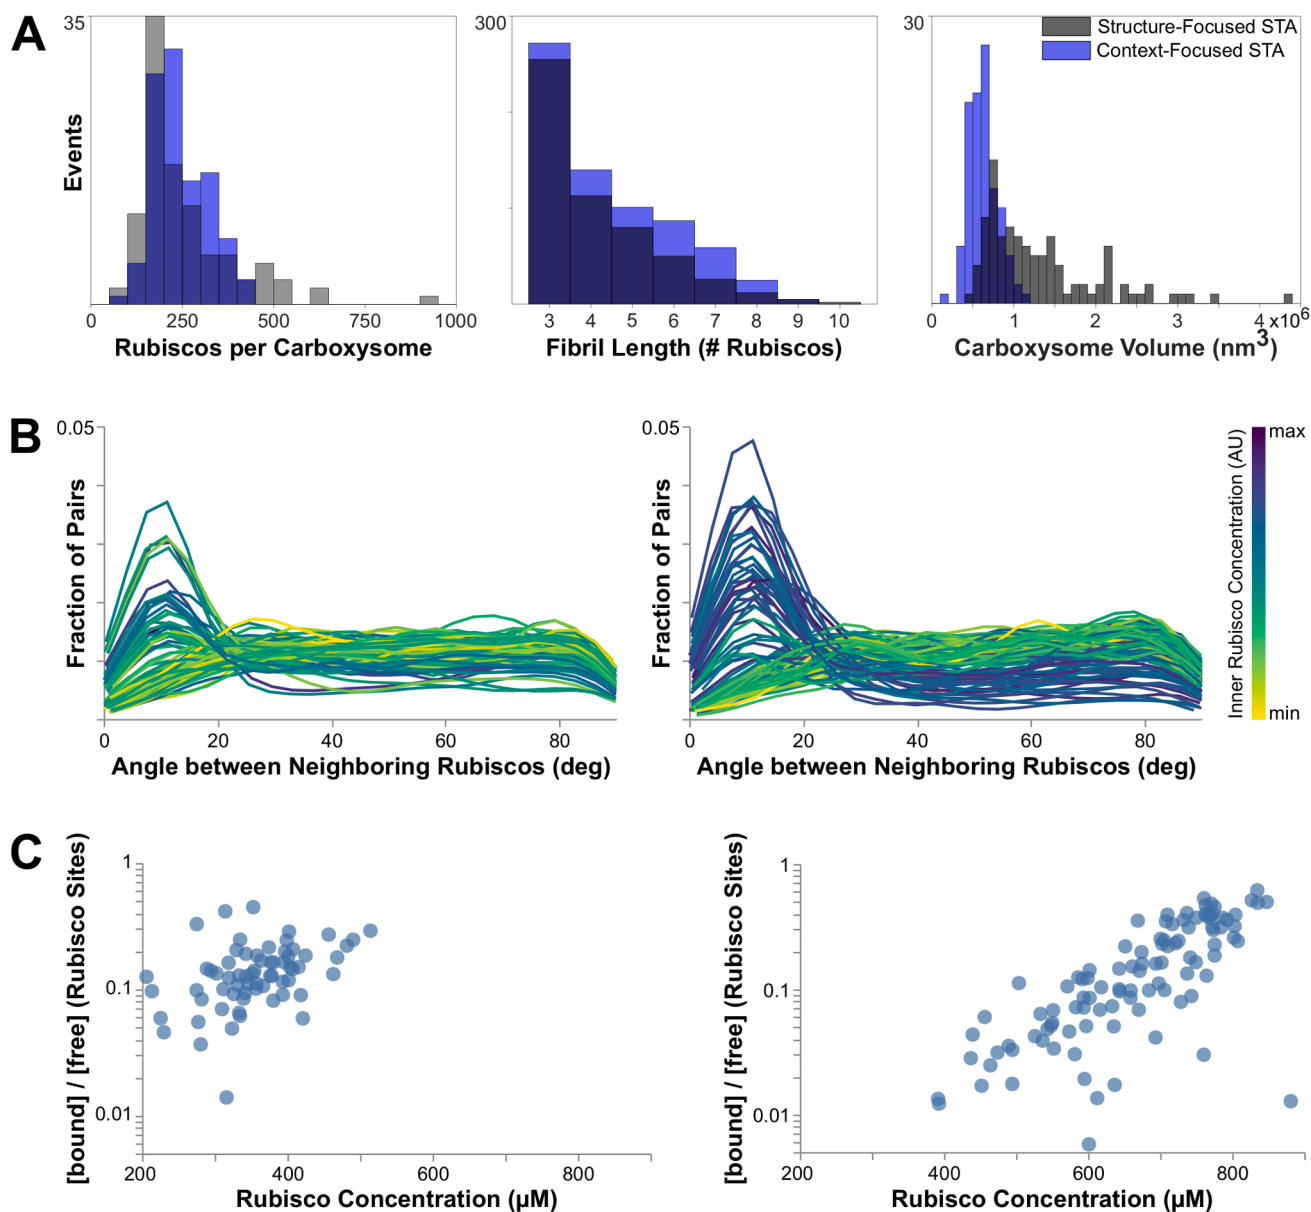


Fig. 3. Comparison of biophysical data from a structure-focused dataset and a context-focused dataset of the purified carboxysomes displayed in Fig. 2 (Metskas et al., 2022). A, from left: histograms showing the difference in the number of particles identified per carboxysome, calculated fibril length based on particle position and orientation, and estimated carboxysome volume. B: histogram traces of nearest-neighbor Rubisco orientations show similar trends, but more robust behavior and stronger concentration dependence in context-focused picking (right) compared to structure-focused picking (left). Each trace is a histogram for one carboxysome. C: plots of percent occupied binding sites compared to Rubisco concentration in each carboxysome. Left, structure-focused pipeline; right, context-focused pipeline. All analyses are significantly different between the datasets ($p < 0.01$) except 3A left.

Align the tilt stack well

Tilt stack alignment is one of the largest contributors to resolution loss for our pipeline and others (Tegunov et al., 2021). We prefer fiducial-based tracking for reliable tilt stack alignment, with roughly 20 fiducials in the field of view. We do not allow microscope-based parameters to vary in our initial fitting of the tilt stack (microscope tilt axis, tilt angle, magnification distortion), as we believe that our Krios performance in these parameters is likely more accurate than the quality of the fit at this step.

For fiducial-based alignment, we implement a semi-automated approach in IMOD. Briefly, every fiducial in the hole is used to seed the fiducial model, and fiducials are then removed completely if they pass outside the field of view or have consistently high residuals compared to the other fiducials in the sample. The fit is fine-tuned until the average residual falls below roughly 2 \AA . Any software or automated

approach could be used provided the residuals are sufficiently low and parameters do not deviate from the known performance of the microscope.

Higher-accuracy particle picking

For biophysical analysis, it is important to accurately identify every particle, driving both false positive and false negative rates as low as possible. False positive and false negative rates can be estimated by comparing particle identifications with manual picking in a high-defocus and lower-defocus tomogram. We find that false positives are more easily identified by looking at lower-cc particles in the tomogram context, while false negatives are best revealed by subtracting identified particles from the tomogram to allow searching of the remaining density.




Fig. 4. The effect of dose-weighting on high-angle tilt images. A: 60-degree tilted projection of *H. neapolitanus*, 3 electron/Å² dose and 120 electron/Å² dose accumulation. B: the same projection after dose-weighting. Pixel size is 2.153 Å.

Many high-resolution STA projects with biophysical analyses are on lattices, with consistent particle spacing and position that allow high accuracy and completeness (Erlendsson et al., 2020; Hutchings et al., 2018; Peukes et al., 2020; Schur et al., 2016; von Kügelgen et al., 2020). When proteins are not in spatially ordered assemblies, particle picking becomes a more difficult problem. In our hands, template-matching and neural network-based approaches have false positive and/or false negative rates of at least 15–25 %, or require extensive manual training to improve the identifications (see Supplemental Table 2 and accompanying discussion). We therefore design semi-automated approaches capitalizing on relationships in our specific datasets (Supplemental Fig. 1).

In our test case, Rubisco is inside CBs; therefore, we segmented the CB shell to confine the search. Although accurate shell segmentation is difficult and time-consuming, a simple 2-plane trace takes only 20 s to provide a rough area of interest (Fig. 6A). We then generated an over-sampled grid of points inside the area of interest (Fig. 6B), and centered each point on the highest-density object within its box. We removed duplicate particles (within 1 pixel) and separated Rubisco, shell and noise through 3D classification (Fig. 6C) (Scheres, 2012). This approach yielded false positive and false negative rates of roughly 20 % each, similar to other methods in our hands, and with only the two-plane segmentation for manual input. The dataset of roughly 42,000 particles in 62 tomograms reached high resolution with the 66 % best-correlating particles (Fig. 1), but was not accurate enough for biophysical analysis (Fig. 3).

For our context-focused workflow, we therefore drove the accuracy of this approach higher through rule-based dataset cleaning. These rules are designed according to the behavior of the specific dataset. For Rubisco in CBs, we first established what worked well: the “shell” class

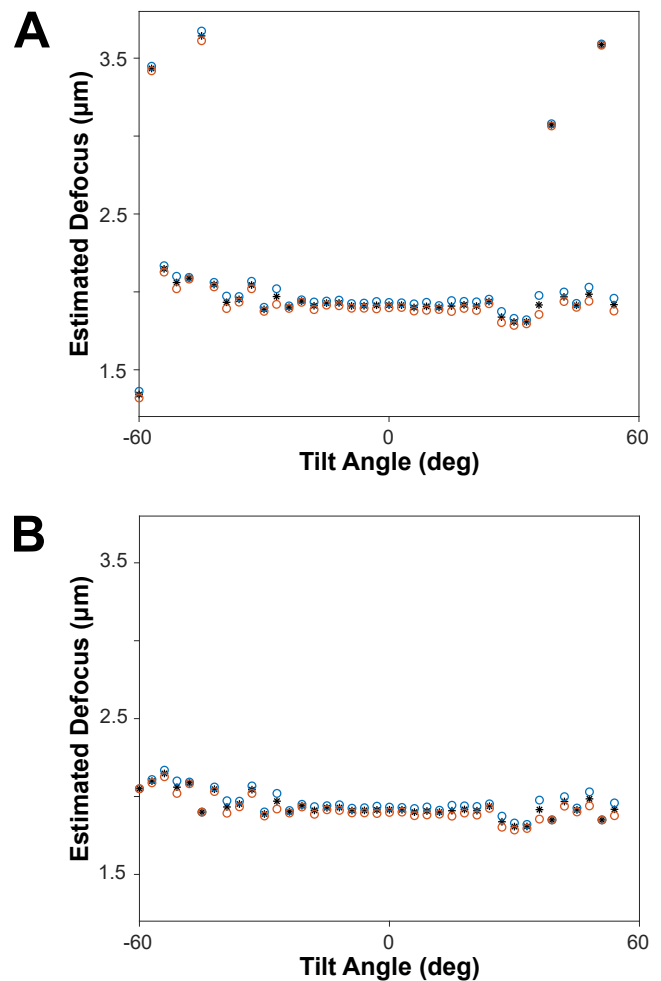


Fig. 5. Estimated defoci in ctffind4 can fail at high tilt. A: ctffind4 raw output. B: Corrected output. Black stars, average defocus; orange and blue circles, defocus with astigmatism. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

had very low false positives in all tomograms (Fig. 7, center), allowing those results to be accepted outright. Removal of the particles in the initial shell class made subsequent classification more accurate, presumably by increasing the percentage of Rubisco particles in the dataset. Therefore, we removed the shell particles, removed particles within 30 pixels of each other, performed a second classification, and again removed the shell class from the data.

At this point, we examined the “noise” and “Rubisco” classes (Fig. 7, right and left). Both had a high percentage of accurate identifications, but contained false positives and false negatives that led to context loss. Looking at the tomograms, some particles could be easily corrected using a rule-based approach. Noise identifications outside the shell were automatically accepted, and Rubisco identifications outside the shell were automatically rejected. The combination of shell class removal with these simple rules brought false positive and negative rates below roughly 8 %, after which point manual screening was required. This approach returned 32,930 particles for context-focused STA (Fig. 1).

Subtomogram averaging

We typically choose a standard cross correlation (cc)-based translation and rotation search for the STA, using Dynamo to extract particles from reconstructed tomograms and perform STA with an adaptive bandpass search (Castaño-Díez et al., 2012). Because the goal is to align each particle as accurately as possible, we do not use strong cc-based




Fig. 6. The particle-picking workflow. A: A simple two-plane segmentation is performed in IMOD on reconstructed tomograms. Green, manual segmentation; scale bar 50 nm. B: The two planes are extended to the maximum third dimension found in the other plane, and the intersection of these volumes is filled with a grid of points for particle searches. C: Cropped particles are centered on the nearest object and classified in Relion into shell (top two classes), Rubisco (bottom left), and noise (bottom right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

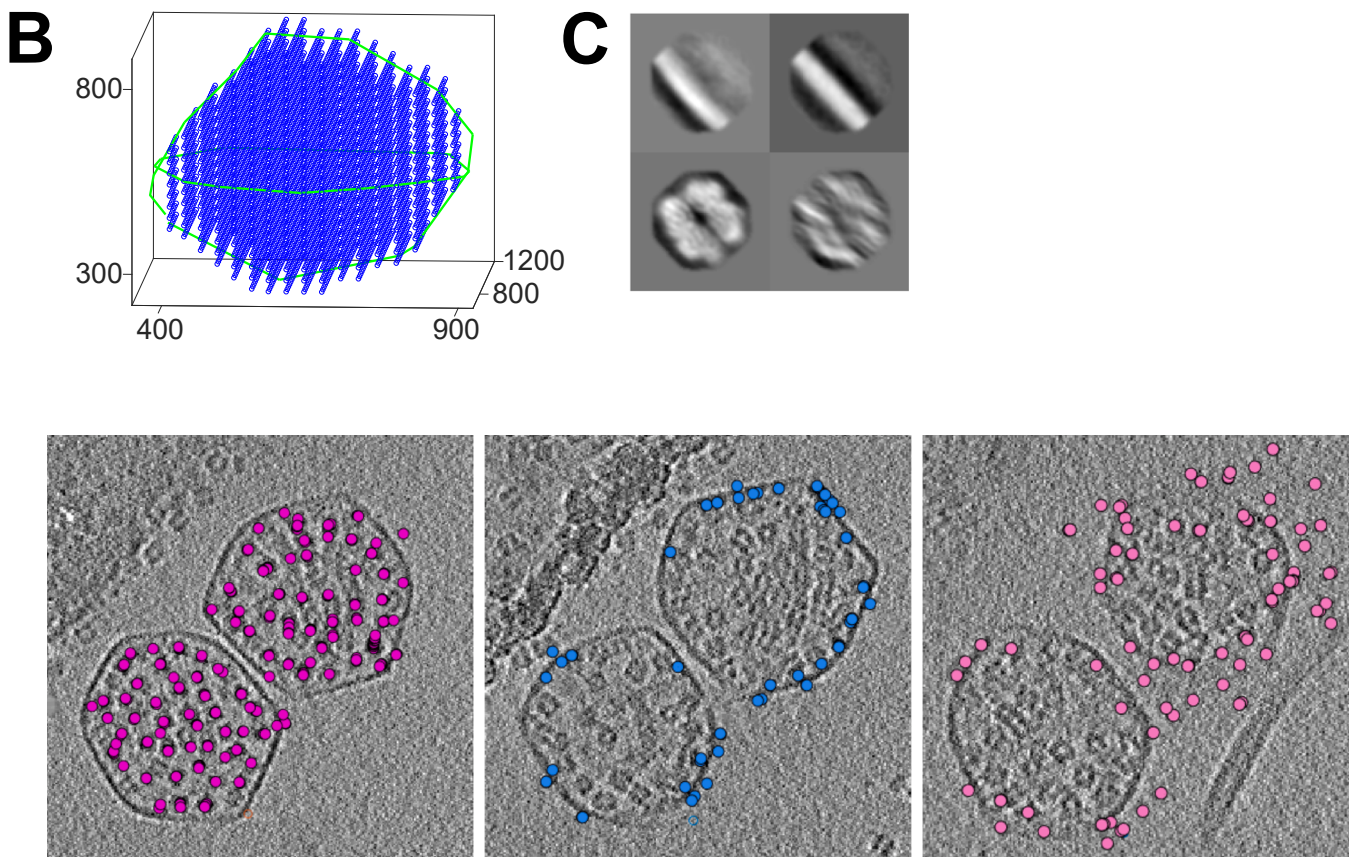


Fig. 7. Unguided classification results. From left: Rubisco classes, shell classes, and noise classes displayed on their parent tomograms.

thresholding for alignment. If a subset of particles were consistently excluded from the average and resolution calculations, this could cause larger errors in the position and orientation of the excluded particles even as the map resolution improves. We therefore monitor early alignment of the dataset using 100 % of particles and a loose mask for FSC calculations.

Following the initial subtomogram average using 100 % of the particles and the initial tomogram reconstructions, we perform biophysical analyses. If desired, additional structural processing can continue from this stage, using the best-aligning subset of particles and iterative refinement of reconstruction parameters to drive resolutions higher.

Biophysical analysis

Tomograms and subtomograms are inherently noisy. We often find

that analyses are best performed from the numerical data of subtomogram position and orientation rather than from the images. This has the additional benefit of being easily adjusted for symmetry, as subunit registrations can be interchangeable in the reference frame of the particle but have contextual differences in the larger tomogram. Because low-defocus tomograms are noisy, we only include tomograms with defoci of 2 μm or greater in our biophysical analysis (after aligning and averaging with the full defocus range). The exact analyses will vary according to the dataset; two are presented here for our test case.

We identified particles within fibrils by searching the position space above and below each Rubisco particle, moving along the C4 axis. The precise distance and angular thresholds were empirically determined by comparing to manual identifications in our “ground truth” CBs used for estimating particle identification accuracy. This approach provided connectivity between particles, which was then used to calculate fibril

length and count the number of bound and unbound Rubisco-Rubisco interaction sites (Fig. 3A and C). Finally, we used these data to perform functional classifications for fibril participation and position within lattices, allowing clean calculations of features such as bend, twist and tilt (Metskas et al., 2022).

To calculate true concentrations of Rubisco inside the CB, we needed to calculate CB volumes. In the CB, two phases existed within a confined space, so both local and global concentrations were needed for analysis. Shell segmentation failed due to the missing wedge effect. However, Rubisco filled an outer layer adjacent to the shell with consistent spacing, allowing us to use the Rubisco positions as an occupied volume model (Metskas et al., 2022). We used a convex hull calculation to identify the outermost layer of Rubisco, and calculated the volume contained within that layer.

Confirmation

All sample systems are biased. Purified systems can suffer from compression artefacts, shearing forces and/or chemical differences compared to *in vivo* preparations. Thick *in vivo* tomograms have increased noise and decreased information content at high tilt, which increases false negatives and/or noise alignment. Finally, FIB-milling samples can sacrifice context; for example, a CB may not be complete inside a lamella. Therefore, where possible a combination of analyses should be performed, identifying which types of artefacts are present and confirming the behavior of interest with complementary approaches.

We typically purify samples when possible, to allow manipulation of variables such as chemical environment, particle concentrations, etc. Purification also provides thinner ice and more numerous imaging locations on the grid, which aid accuracy of particle and environment identifications by decreasing noise and increasing sample sizes. The purified sample is then validated with *in vivo* data to identify any purification-specific effects.

Discussion and conclusion

The subtomogram averaging test case presented here contains a high concentration of particles within well-defined boundaries, which allows the rare generation of a “ground truth” dataset. Such high accuracy in particle identification is not currently possible in other systems, such as a dispersed protein within a cellular environment. However, the illustration of this test case and the separate effects of false negatives and false positives can guide experimentalists in analyzing these data and establishing the most robust analysis possible. Software development in the field is ongoing, and improved particle picking may well allow higher-accuracy workflows in such systems in the future.

While this work focuses only on maximizing information from STA, the field of single-particle cryo-electron microscopy (SPA) has a parallel situation. In SPA, a large number of particles are routinely discarded in an effort to attain the highest resolution possible in the final average of a structural class. Many of these discarded particles are at the air-water interface and thus undesirable due to damage or preferred orientation (Noble et al., 2018), but it is likely that this approach also discards relevant data such as alternate conformations. It is possible that an approach similar to our design for STA could be useful for conformational ensemble analyses in SPA prior to solving of high-resolution structures.

STA is a powerful tool for analyzing the arrangement of proteins in supramolecular assemblies. Here we present and discuss a modified STA pipeline designed to reveal biophysical context, which remains compatible with additional structure-focused processing. In this way, experimentalists can use a context-focused pipeline to generate both biophysical data and an initial target protein structure, and then send the results to established software pipelines for iterative per-particle refinement and resolution improvement.

Detailed step-by-step protocol

Here we present a detailed, step-by-step protocol for a range of potential samples and targets, with specific details for our CB test case to show one possible workflow. Explanations for unique components of the workflow can be found above; other elements are standard for STA and have been described extensively in other publications.

Tomography

1. Prepare sample and freeze onto thin carbon grids (we use CF-2/2–300).
2. Perform dose-symmetric tomogram acquisition with a pixel size $\sim 1\text{--}1.5\text{ \AA}$ and a range ~ 66 degrees in 3-degree increments, 8–10 frames per tilt, using a total dose of 120 electrons/ \AA^2 or as suits the sample.
 - a. Our SerialEM acquisition script is closely based upon the original dose-symmetric Serial EM acquisition script (Hagen et al., 2017), but was modified to optimize performance for the Caltech Titan Krios (see script repository). However, newer SerialEM versions now support dose-symmetric acquisition, streamlining the process and allowing easier optimization for different targets. Relevant parameters from our acquisition are a 1 $\text{\AA}/\text{s}$ drift limit in Trial mode and two-tilt grouping throughout the tilt scheme.
3. Perform frame alignment in the software of choice, and assemble ordered tomograms. We use IMOD alignframes and newstack with default settings, with a script that automatically monitors the residuals and re-runs poor frame alignments with settings optimized for our high-tilt data (see script repository).
4. Remove any unusable tilts (off-target, large chunks of ice, severe drift, etc.) from the tilt stacks using the IMOD newstack command. While tilts can be excluded later in etomo reconstructions, software handling of bad tilts can vary between packages and wastes processing time in automated pipelines.
5. Perform defocus estimation. We use ctfind4 for its easy bulk processing and integration with other software, with a script to correct bad estimates as needed (see script repository). Ensure that the signal driving the defocus estimation is at the same z height as the center of the object of interest. For example, CBs tend to be centered in the ice of a hole, such that a carbon edge contributes to the CTF signal at roughly the same defocus as the center of the CB; however, this is not true for cells sitting atop a carbon surface.
6. Perform dose-weighting and write a dose-weighted tilt stack using software of choice. Pseudocode for dose-weighting has been included in the script package for this dataset. For datasets with thin ice where high tilt angles contribute more information, we adjust the pixel size according to the tilt angle; however, this is less useful for thick samples.
7. Remove X-rays and hot pixels, align the tilt stack, and identify fiducial positions for erasing gold. We perform these functions in IMOD using the dose-weighted tilt stack (Mastronarde and Held, 2017), but any software tool with good results can be used. We optimize fine alignments until mean residuals drop below 2 \AA for a non-iterative STA resolution of 4.5 \AA .
8. Reconstruct the tomogram. We use novaCTF according to its documentation, incorporating gold bead erasing, edge tapering and other features in the indicated places (see novaCTF documentation and our script repository) (Turoňová et al., 2017). We generate both a bin1 WBP and a bin2 SIRT-like filtered tomogram at this step; the WBP is for STA, while the SIRT-like filtered tomogram is used for particle picking.

Subtomogram averaging

9. Define the region of interest for particle picking in the bin2 SIRT-like filtered tomograms. This can be a segmented surface or volume, a geometrical approximation, or a simple exclusion zone.
10. Identify particles within the region of interest. Particle picking is highly target-dependent, so users should use the approach providing the highest accuracy. For samples such as membrane lattices, this can be done with a simple geometric search and oversampled grid (Scaramuzza and Castaño-Díez, 2021). Sparse samples may be picked manually, or with neural networks or template-matching algorithms. For densely packed samples in an irregular geometry, we use the following approach:
 - a. Fill the region of interest with an oversampled grid (see script repository).
 - b. Use 1–2 iterations of alignment and averaging to center each box onto the nearest dense object, using a low-resolution reference and maximal possible symmetry. Remove duplicate particles.
 - c. Perform 3D classification in Relion to separate particles of interest from other features. (The repository contains a script for transferring results from Dynamo v1.1.532 to Relion v3.1.)
 - d. Follow dataset-specific rules to drive false negative and false positive rates as low as possible.
11. Divide the final dataset into half-sets, dividing by object within tomograms.
12. Recrop the particles within the halfsets from the bin1 WBP tomograms, converting the table positions to bin1.
13. Proceed with STA, using adaptive bandpass filtering and as few iterations as possible to reduce noise alignment. The workflow will vary according to the project, and numerous STA workflows have been described in the literature. Briefly, for Rubisco tomograms with 1.104 Å unbinned pixel size, we use the following approach in Dynamo:
 - a. Perform six alignment iterations using a box size of 160 pixels, spherical alignment mask of radius 70 with a 5-pixel gaussian edge, and binning by two. Constrain shift limits to 10 pixels (particles should already be roughly centered from particle picking). All iterations include five refinement steps, cutting the search space in half each time.
 - i. The first iteration is to center the recropped particles in the boxes. D4 symmetry, 24 degree angular search space, 2.2 nm low-pass filter, reference a simple initial average of the half-set.
 - ii. Two global search iterations are performed in C1 symmetry to preserve any asymmetric interactions. The cone search can be confined to 90 degrees if needed to stop particles from losing orientation.
 - iii. Two iterations are performed in C4 symmetry to improve resolution and help lock on to subunit registration. The angular search space is confined to 21 degrees but allows a cone flip.
 - iv. The final iteration is in C4 symmetry, 21 degree search without a cone flip. Resolution reached 6.5 Å at this point, roughly 2/3 Nyquist for the bin2 pixel size.
 - b. Generate an alignment mask following the contours of the particle, keeping it loose (at least 8 pixels from the surface of the density) and softening the edges by 5 pixels. A crystal structure can be used, or thresholded average density if no structure is available.
 - c. Perform fine alignment iterations with C4 symmetry, maintaining the shift limits from above and using the unbinned pixel size.
 - i. Begin with a 15-degree angular search and 6.5 Å low-pass filter, and decrease the search space and the filter together until the resolution ceases to improve by at least 1 Fourier pixel per iteration. In our hands, this took four iterations and resulted in a 4.7 Å resolution using the loose mask.
14. Post-process using a tighter mask that does not cut into any protein density, and sharpen and filter the final map.

- a. We prefer to do STA in Dynamo but post-processing in Relion. To do this, save the half-maps as mrc files following Relion naming conventions and update the headers with the correct pixel size. In Relion, post-process using default values and automated B-factor estimation.

Biophysical analysis

15. Recenter the particles in their boxes by applying any accumulated shifts to the tomogram coordinates.
16. Perform geometrical analyses and classifications as desired; ours are described in the above section and in a recent publication (Metskas et al., 2022).
17. Perform an orthogonal analysis to validate key findings.

Additional structural processing

18. If desired, continue refining the resolution using an iterative per-particle refinement approach. From this point the dataset can be culled to include only well-resolving particles.

Data availability

Scripts, pseudocode and data tables are available in our script repository at https://github.com/LAMetskas/2022_STA. All scripts are provided as is and may require revision according to directory structure, naming conventions, or software. For our presented CB/Rubisco test case, sample tomograms, particle locations and orientations for the context-focused STA, and the set of all Rubisco particles were deposited in the Electron Microscopy Public Image Archive under accession code EMPIAR-11125 (<https://www.ebi.ac.uk/empiar/EMPIAR-11125/>). Rubisco half-maps and a full filtered map of the Rubisco subtomogram average were deposited in the Electron Microscopy Data Bank under accession code EMD-27654 (<https://www.ebi.ac.uk/emdb/EMD-27654/>). Biophysical analysis code is at <https://observablehq.com/collection/@lametskas/cbpaper>. Data will be made available on request.

CRedit authorship contribution statement

Lauren Ann Metskas: Conceptualization, Investigation, Software, Methodology, Supervision. **Rosalie Wilfong:** Investigation. **Grant J. Jensen:** Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank S. Chen and A. Malyutin for assistance with the tomography data collection script, and A. Burt for pseudocode to transition between Dynamo and Relion software packages. This work was supported by a Ruth L. Kirschstein NRSA Individual Postdoctoral Fellowship F32 1F32GM135994-01 to LAM, NIAID New Innovators Award 1DP2AI164293-01 to LAM, and NIH R01 AI127401 to GJJ.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.yjsbx.2022.100076>.

References

- Burt, A., Gaifas, L., Dendooven, T., Gutsche, I., 2021. A flexible framework for multi-particle refinement in cryo-electron tomography. *PLoS Biol.* 19, e3001319 <https://doi.org/10.1371/journal.pbio.3001319>.
- Castaño-Díez, D., Kudryashev, M., Arheit, M., Stahlberg, H., 2012. Dynamo: a flexible, user-friendly development tool for subtomogram averaging of cryo-EM data in high-performance computing environments. *J. Struct. Biol.* 178, 139–151. <https://doi.org/10.1016/j.jsb.2011.12.017>.
- Chen, M., Bell, J.M., Shi, X., Sun, S.Y., Wang, Z., Ludtke, S.J., 2019. A complete data processing workflow for cryo-ET and subtomogram averaging. *Nat. Methods* 16, 1161–1168. <https://doi.org/10.1038/s41592-019-0591-8>.
- Dai, W., Chen, M., Myers, C., Ludtke, S.J., Pettitt, B.M., King, J.A., Schmid, M.F., Chiu, W., 2018. Visualizing Individual RuBisCO and Its Assembly into Carboxysomes in Marine Cyanobacteria by Cryo-Electron Tomography. *J. Mol. Biol.* 430, 4156–4167. <https://doi.org/10.1016/j.jmb.2018.08.013>.
- Erlendsson, S., Morado, D.R., Cullen, H.B., Feschotte, C., Shepherd, J.D., Briggs, J.A.G., 2020. Structures of virus-like capsids formed by the *Drosophila* neuronal Arc proteins. *Nat. Neurosci.* 23, 172–175. <https://doi.org/10.1038/s41593-019-0569-y>.
- Grant, T., Grigorieff, N., 2015. Measuring the optimal exposure for single particle cryo-EM using a 2.6 Å reconstruction of rotavirus VP6. *Elife* 4, e06980. <https://doi.org/10.7554/eLife.06980>.
- Hagen, W.J.H., Wan, W., Briggs, J.A.G., 2017. Implementation of a cryo-electron tomography tilt-scheme optimized for high resolution subtomogram averaging. *J. Struct. Biol.* 197, 191–198. <https://doi.org/10.1016/j.jsb.2016.06.007>.
- Himes, B.A., Zhang, P., 2018. emClarity: software for high-resolution cryo-electron tomography and subtomogram averaging. *Nat. Methods* 15, 955–961. <https://doi.org/10.1038/s41592-018-0167-z>.
- Hutchings, J., Stancheva, V., Miller, E.A., Zanetti, G., 2018. Subtomogram averaging of COPII assemblies reveals how coat organization dictates membrane shape. *Nat. Commun.* 9, 4154. <https://doi.org/10.1038/s41467-018-06577-4>.
- Iancu, C.V., Morris, D.M., Dou, Z., Heinhorst, S., Cannon, G.C., Jensen, G.J., 2010. Organization, structure, and assembly of alpha-carboxysomes determined by electron cryotomography of intact cells. *J. Mol. Biol.* 396, 105–117. <https://doi.org/10.1016/j.jmb.2009.11.019>.
- Mastronarde, D.N., Held, S.R., 2017. Automated tilt series alignment and tomographic reconstruction in IMOD. *J. Struct. Biol. Electron Tomogr.* 197, 102–113. <https://doi.org/10.1016/j.jsb.2016.07.011>.
- Mattei, S., Glass, B., Hagen, W.J.H., Kräusslich, H.-G., Briggs, J.A.G., 2016. The structure and flexibility of conical HIV-1 capsids determined within intact virions. *Science* 354, 1434–1437. <https://doi.org/10.1126/science.aah4972>.
- Metskas, L.A., Ortega, D., Oltrogge, L.M., Blikstad, C., Laughlin, T., Savage, D.F., Jensen, G.J., 2022. Rubisco forms a lattice inside alpha-carboxysomes. <https://doi.org/10.1101/2022.01.24.477598>.
- Ni, T., Sun, Y., Seaton-Burn, W., Al-Hazeem, M.M.J., Zhu, Y., Yu, X., Liu, L.-N., Zhang, P., 2022. Tales of Two α -Carboxysomes: the Structure and Assembly of Cargo Rubisco. <https://doi.org/10.1101/2022.03.15.484529>.
- Noble, A.J., Wei, H., Dandey, V.P., Zhang, Z., Tan, Y.Z., Potter, C.S., Carragher, B., 2018. Reducing effects of particle adsorption to the air–water interface in cryo-EM. *Nat. Methods* 15, 793–795. <https://doi.org/10.1038/s41592-018-0139-3>.
- Obr, M., Hagen, W.J.H., Dick, R.A., Yu, L., Kotecha, A., Schur, F.K.M., 2022. Exploring high-resolution cryo-ET and subtomogram averaging capabilities of contemporary DEDs. *J. Struct. Biol.* 214, 107852. <https://doi.org/10.1016/j.jsb.2022.107852>.
- Peukes, J., Xiong, X., Erlendsson, S., Qu, K., Wan, W., Calder, L.J., Schraidt, O., Kummer, S., Freund, S.M.V., Kräusslich, H.-G., Briggs, J.A.G., 2020. The native structure of the assembled matrix protein 1 of influenza A virus. *Nature* 587, 495–498. <https://doi.org/10.1038/s41586-020-2696-8>.
- Qu, K., Ke, Z., Zila, V., Anders-Össwein, M., Glass, B., Mücksch, F., Müller, R., Schultz, C., Müller, B., Kräusslich, H.-G., Briggs, J.A.G., 2021. Maturation of the matrix and viral membrane of HIV-1. *Science* 373, 700–704. <https://doi.org/10.1126/science.abe6821>.
- Rohou, A., Grigorieff, N., 2015. CTFFIND4: Fast and accurate defocus estimation from electron micrographs. *J. Struct. Biol.* 192, 216–221. <https://doi.org/10.1016/j.jsb.2015.08.008>.
- Scaramuzza, S., Castaño-Díez, D., 2021. Step-by-step guide to efficient subtomogram averaging of virus-like particles with Dynamo. *PLoS Biol.* 19, e3001318. <https://doi.org/10.1371/journal.pbio.3001318>.
- Scheres, S.H.W., 2012. RELION: Implementation of a Bayesian approach to cryo-EM structure determination. *J. Struct. Biol.* 180, 519–530. <https://doi.org/10.1016/j.jsb.2012.09.006>.
- Schur, F.K.M., Obr, M., Hagen, W.J.H., Wan, W., Jakobi, A.J., Kirkpatrick, J.M., Sachse, C., Kräusslich, H.-G., Briggs, J.A.G., 2016. An atomic model of HIV-1 capsid-SP1 reveals structures regulating assembly and maturation. *Science* 353, 506–508. <https://doi.org/10.1126/science.aaf9620>.
- Sun, Y., Harman, V.M., Johnson, J.R., Brownridge, P.J., Chen, T., Dykes, G.F., Lin, Y., Beynon, R.J., Liu, L.-N., 2022. Decoding the Absolute Stoichiometric Composition and Structural Plasticity of α -Carboxysomes. *mBio* 13, e0362921. <https://doi.org/10.1128/mbio.03629-21>.
- Tegunov, D., Xue, L., Dienemann, C., Cramer, P., Mahamid, J., 2021. Multi-particle cryo-EM refinement with M visualizes ribosome-antibiotic complex at 3.5 Å in cells. *Nat. Methods* 18, 186–193. <https://doi.org/10.1038/s41592-020-01054-7>.
- Turoňová, B., Schur, F.K.M., Wan, W., Briggs, J.A.G., 2017. Efficient 3D-CTF correction for cryo-electron tomography using NovaCTF improves subtomogram averaging resolution to 3.4 Å. *J. Struct. Biol.* 199, 187–195. <https://doi.org/10.1016/j.jsb.2017.07.007>.
- von Kügelgen, A., Tang, H., Hardy, G.G., Kureisaite-Ciziene, D., Brun, Y.V., Stansfeld, P. J., Robinson, C.V., Bharat, T.A.M., 2020. In Situ Structure of an Intact Lipopolysaccharide-Bound Bacterial Surface Layer. *Cell* 180, 348–358.e15. <https://doi.org/10.1016/j.cell.2019.12.006>.
- Zivanov, J., Otón, J., Ke, Z., Qu, K., Morado, D., Castaño-Díez, D., Kügelgen, A. von, Bharat, T.A.M., Briggs, J.A.G., Scheres, S.H.W., 2022. A Bayesian approach to single-particle electron cryo-tomography in RELION-4.0. <https://doi.org/10.1101/2022.02.28.482229>.