Published in final edited form as:

Nat Protoc. 2016 November ; 11(11): 2054–2065. doi:10.1038/nprot.2016.124.

Resolving macromolecular structures from electron cryo-tomography data using sub-tomogram averaging in RELION

Tanmay A.M. Bharat* and Sjors H.W. Scheres*

Structural Studies Division, MRC Laboratory of Molecular Biology, Cambridge CB2 0QH, United Kingdom

Abstract

Electron cryo-tomography (cryo-ET) is a technique that is used to produce three-dimensional pictures (tomograms) of complex objects like asymmetric viruses, cellular organelles or whole cells from a series of tilted electron cryo-microscopy (cryo-EM) images. Averaging of macromolecular complexes found within tomograms is known as sub-tomogram averaging, and this technique allows structure determination of macromolecular complexes *in situ*. Sub-tomogram averaging is also gaining in popularity for the calculation of initial models for single-particle analysis. We describe herein, a protocol for sub-tomogram averaging from cryo-ET data using the RELION software (<http://www2.mrc-lmb.cam.ac.uk/relion>). RELION was originally developed for cryo-EM single-particle analysis and the sub-tomogram averaging approach presented in this protocol has been implemented in the existing workflow for single-particle analysis so that users may conveniently tap into existing capabilities of the RELION software. We describe how to calculate three-dimensional models for the contrast transfer function (CTF), which describe the transfer of information in the imaging process, and we illustrate the results of classification and subtomogram averaging refinement for cryo-ET data of purified hepatitis B capsid particles and *S. cerevisiae* 80S ribosomes. Using the steps described in this protocol, along with the troubleshooting and optimisation guidelines, high-resolution maps can be obtained where secondary structure elements are resolved.

Introduction

Imaging macromolecular complexes that are frozen in a thin layer of vitreous ice using an electron microscope (cryo-EM) is rapidly gaining in popularity. A single cryo-EM image may contain 2D projections of many copies of the same complex in different orientations, and these 2D projections may be combined in a 3D reconstruction of its scattering potential.

*Correspondence to: tbharat@mrc-lmb.cam.ac.uk or scheres@mrc-lmb.cam.ac.uk.

Competing Interests:

The authors declare that they have no competing financial interests.

Author Contributions

T.A.M.B. performed tomographic data acquisition and data processing, and developed the python script for pre-processing of the sub-tomograms. S.H.W.S. developed the sub-tomogram procedures inside RELION. Both authors contributed to writing the manuscript.

Data Deposition

The data set of 7 tomograms of 80S ribosomes from *S. cerevisiae* has been deposited at the EMPIAR database (EMPIAR-10045) and the final map from 3D auto-refinement of these data has been submitted to the EMDB (EMD-3228).

This technique, which is known as single-particle analysis, has recently undergone significant progress with the development of highly efficient direct-electron detectors and improved image processing software. Notably, this technique now allows near-atomic resolution structures to be calculated without the need for crystallisation and from as little as 10-100 µg of purified material 1, 2.

In electron tomography (ET) multiple images are taken of the same sample region at different tilt angles in the microscope. From such a series of tilted images, a 3D reconstruction, or tomogram, of a single 3D object such as an entire cell 3 may be obtained. Thereby, this technique provides the unique possibility to image complexes in their native environment. Moreover, if many copies of a complex of interest are present in tomograms, then the reconstructed 3D density corresponding to each complex may be computationally extracted, and the resulting 3D 'sub-tomograms' may be averaged together to increase the signal-to-noise ratio and thereby produce a higher resolution 3D structure 4. This technique is called sub-tomogram averaging, and it has been successfully applied in numerous cases to reveal biological structures *in situ* or in environments that are otherwise not amenable to single-particle analysis 5–9.

To date, the use of sub-tomogram averaging is not as widespread as that of single-particle analysis. An important limitation of sub-tomogram averaging is that the best resolved structures by this technique are markedly lower in resolution than those from single-particle analysis 4. Tomographic data collection is slower, and sub-tomogram averaging requires more complicated image processing, since tomographic reconstruction needs to be followed by alignment and classification of the sub-tomograms. Furthermore, due to increased effective specimen thickness at high tilt angles the sample cannot be imaged at high tilt angles, which leads to a wedge-shaped region in the Fourier domain where data is absent. This 'missing-wedge' leads to blurring artefacts in tomograms. Still, the advantage of being able to study macromolecules *in situ* (e.g. inside an entire cell) remains extremely attractive. This is powerfully illustrated by the recent application of sub-nanometer resolution cryo-ET sub-tomogram averaging to the HIV-1 capsid 10 and to membrane-bound ribosomes 11. Further developments of both experimental data acquisition procedures 12 and image processing algorithms 13 will continue to drive this technique towards higher resolutions and wider applicability.

Recently, we introduced a new image processing approach to sub-tomogram averaging 14 that is based on a regularized likelihood optimization algorithm in the RELION program 15, 16. This program was originally designed for single-particle analysis and has been used to calculate numerous near-atomic resolution structures 1. Because the sub-tomogram averaging approach in RELION was modelled on the single-particle analysis workflow, existing RELION users will find many similarities (Figure 1). The main deviation from the single-particle workflow lies in the generation of a 3D model for the information transfer in each sub-tomogram, which is used to compensate for both the missing wedge as well as the effects of the contrast transfer function (CTF) in the tomogram 14. A significant effort was made to build on existing tools inside RELION, rather than writing new tools specifically for sub-tomogram averaging. This facilitates transitioning between sub-tomogram averaging

and single-particle analysis, and thus naturally supports a hybrid approach of combining cryo-EM and cryo-ET data 17–19.

In this protocol we describe the practical use of RELION for sub-tomogram averaging. Our approach complements various single-particle analysis software packages that also offer functionalities for sub-tomogram averaging 20, 21, as well as multiple specialized packages for sub-tomogram averaging 8, 9, 13, 22–24. As many structure determination projects in practice resort to a combination of different software packages, we will explicitly indicate those points in the workflow that are likely points of conversion between alternative approaches. Recommended procedures for single-particle analysis in RELION are described in detail in the online documentation (the RELION wiki: <http://www2.mrc-lmb.cam.ac.uk/relion>). Using the protocol described here together with the online documentation, novice users should be able to conduct sub-tomogram averaging for their own project using RELION. We assume basic familiarity with Unix/Linux based systems, and the ability to run provided scripts from the command-line.

Materials

Equipment setup

- A computer or a computing cluster with at least 64–128 processors, 120 Gb of shared memory, and 1 Tb of disk space.
- A RELION installation (version 1.4; available for free from <http://www2.mrc-lmb.cam.ac.uk/relion>).
- A Python installation (version 2.6 or newer; available from <https://www.python.org>)
- An IMOD installation (version 4.7 or newer; available for free from <http://bio3d.colorado.edu/imod/>)
- A CTFFIND installation (version 3 or newer; available for free from <http://grigoriefflab.janelia.org/ctffind4>)

Data

- Tomograms: To enter the RELION workflow, we assume that tomograms have been generated in MRC format 25. Tomogram calculation is not done inside RELION, but relies on software packages like IMOD 24, Tomo3D 26, pyTOM 13 or Bsoft 21. In the examples below, we used the IMOD package for tilt series alignment, and we used Tomo3D for tomographic reconstruction.
- Sub-tomogram coordinates: To find the positions of macromolecular complexes within tomograms we either used the MolMatch software for template matching 8 or manually picked particles using 3dmod 24 and converted the resulting model into a text file using the IMOD model2point command.

- Aligned tilt series: The final aligned tilt series in MRC format is required for CTF estimation in RELION for 3D classification and refinement.

Procedure

Arrangement of input files and directories

1| The directory from which all RELION commands and the graphical-user-interface (GUI) is launched will be known as the project directory (/). Make a folder inside the project directory that is called ./Tomograms/ where all the input data is stored.

2| Generate a sub-directory in the ./Tomograms/ directory with the name of each tomogram in the data set, and copy all tomograms to their respective sub-directories. For example, if there are two tomograms in the data set, then the locations of those tomograms should be as follows -

./Tomograms/tomogram1/tomogram1.mrc

./Tomograms/tomogram2/tomogram2.mrc

3| Save the coordinates of the centres of all macromolecular complexes, or particles, in a tomogram in the same sub-directories as the corresponding .mrc file. The coordinate files should have the suffix .coords and the prefix should be the same as the prefix of the tomogram name. In our two-tomogram example the name and location of the coordinate files should be -

./Tomograms/tomogram1/tomogram1.coords

./Tomograms/tomogram2/tomogram2.coords

The coordinates within each file should be written out in a three-column ASCII format, corresponding to the X,Y,Z position of the centre of each sub-tomogram in pixels. The origin of the tomogram is in the lower left corner (when displayed in IMOD). Each coordinate file should contain as many lines as there are particles in the tomogram. For example -

100.0 355.0 200.0

2034.0 1100.0 561.0

3011.0 2539.0 321.0

4| For CTF estimation, place the aligned tilt series in the same sub directories. For each tomogram, this should be a single MRC stack, and the suffix of the files should be .mrscs . It is important that this is the exact same aligned tilt series as the one used for tomographic reconstruction. Again, in our example these should be called -

./Tomograms/tomogram1/tomogram1.mrscs

./Tomograms/tomogram2/tomogram2.mrscs

5| Along with the aligned stack, also provide the final tilt angles used for tomographic reconstruction (e.g. from IMOD) in a separate text file. Each line in these text files should contain one number corresponding to the final tilt angle assigned to that image during alignment. The order of the lines should correspond to the order of the images in the aligned stack.

59.44

56.44

53.43

Copy these angles files into the same directory.

`./Tomograms/tomogram1/tomogram1.tlt`

`./Tomograms/tomogram2/tomogram2.tlt`

If .tlt files with angle values are not provided, then the tilt angles will be read from the extended header of the .mrcs file, if this exists, in step 9 (see below).

6| For each aligned tilt series, create a text file that lists the tilt angles and the accumulated radiation for each image in the tilt series. If the dose on the detector was calibrated then the values for the accumulated dose could be read from output log files of the microscope data acquisition software, for example the .mdoc file written out by SerialEM 27. This information will be used to calculate a dose-dependent 3D CTF model (Figure 2A-B) that also accounts for radiation-induced damage of the specimen. This text file should have as many lines as there are images in the tilt series. The text file should have two columns: the first one for the refined tilt angle (after tilt-series alignment), and the second one for the total accumulated dose in $e^-/\text{\AA}^2$ prior to collecting that image. If you provide the nominal tilt angles rather than the refined tilt angles, the python setup script in step 9 will assign accumulated dose values to the closest refined tilt angle from the .tlt file. For example –

61.1 0.0

57.3 2.0

54.2 4.0

Save these files in the same sub-directories with the suffix .order -

`./Tomograms/tomogram1/tomogram1.order`

`./Tomograms/tomogram2/tomogram2.order`

7| Generate a RELION-type metadata file in the STAR format 16, 28, called `all_tomograms.star`, using the command line -

`relion_star_loopheader rlnMicrographName > all_tomograms.star`

```
ls ./Tomograms/*/*.mrc >> all_tomograms.star
```

The output file lists all the tomograms in the data set, and should have the format -

```
data_
```

```
loop_
```

```
_rlnMicrographName
```

```
./Tomograms/tomogram1/tomogram1.mrc
```

```
./Tomograms/tomogram1/tomogram2.mrc
```

Calculation of 3D CTF models for each sub-tomogram

<CRITICAL> For routine application of the steps in this part of the protocol you can use a single python script called `relion_prepare_subtomograms.py` on the RELION wiki. This script takes the `all_tomograms.star` file as input, and carries out steps 8-11, depending on the options that the user specifies in its header. For specialized applications or in case of troubleshooting, any of the steps 8-11 may be carried out independently.

8| CTF correction depends on the estimation of a defocus value for every image of the tilt series (i.e. for all the individual images in the .mrcs stacks mentioned above). We do this by calling CTFFIND 2.9 through the wrapper provided in RELION: `relion_run_ctffind`. Executing this program will be done automatically by the `relion_prepare_subtomograms.py` script. Change the input parameters for CTFFIND in the header of the python script. Specify inputs like voltage, spherical aberration and detector pixel size based on your data collection experiment.

? TROUBLESHOOTING

9| Run the provided python script from the RELION project directory by typing the following text at the command line -

```
python relion_prepare_subtomograms.py
```

As in single-particle analysis, this script will create a `./Particles/Tomograms/` directory with one sub-directory per tomogram. Following CTF estimation using CTFFIND, in each tomogram sub-directory .star files corresponding to the CTF model for each sub-tomogram will be written out. A `particles_subtomo.star` file that lists all the sub-tomograms as well as their corresponding 3D CTF models will also be written to the project directory. This file will be the input for 3D classification and 3D auto-refinement in RELION.

? TROUBLESHOOTING

10| Inspect the diagnostic output of CTFFIND for every image of every tilt series in the data set (Figure 2C-E). This may be conveniently done using the **'Display'** button in the RELION

GUI (Figure 3), and selecting the output `_ctffind.star` file for each tomogram, e.g. `./Tomograms/tomogram1/ctffind/tomogram1_ctffind.star`.

<CRITICAL STEP> If CTF estimation failed for some images (see Figure 2E), then `reliion_run_ctffind` should be executed again for that image with different parameters. We recommend that either the entire python script should be run again (step 9), or users can look in the `reliion_subtomo_commands.txt` file for the relevant commands for individual tilt series and run them again. Alternatively, one could use an external CTF estimation program, possibly one that allows manual steering like EMAN2's `e2ctf.py` 20, to find a defocus that fits the observed power spectrum of the image. In that case, the resulting defocus value should be inserted manually into the output `_ctffind.star` file by changing the corresponding values using a text file editor. In fact, a complete set of defocus values (for all images in the tilt series) pre-estimated using another software could be manually entered into this `_ctffind.star` file.

? TROUBLESHOOTING

11 | At this point there are three extra options that can be used to improve CTF estimation accuracy.

For flat samples, the effective thickness of the ice that the electron beam passes through increases with increasing tilt angle. This results in lower signal-to-noise ratios for the higher tilt images, which may preclude reliable CTF estimation. If this is the case, the average defocus measured for the lower tilt images may be applied to the higher tilt images, especially if the applied defocus is stable throughout the tilt series 10. To do this, set the `UseOnlyLowerTiltDefoci` variable to `True` in the header of the `reliion_prepare_subtomograms.py` script and provide a threshold value for the lower tilt.

Another way to estimate the CTF parameters of higher tilt images more accurately is to acquire two extra images for each setting of the stage in the tilt series. We collected these images along the direction of the tilt axis, and spaced equally on either side 14, 30 of the region of interest. If such images were collected, then save the extra images in MRC format with a `.trial` suffix in the relevant tomogram sub-directories -

```
./Tomograms/tomogram1/tomogram1.trial
```

```
./Tomograms/tomogram2/tomogram2.trial
```

The number of images in the `.trial` stack should be exactly double the number of images in the original tilt-series stack and the order of the images should be the same as the aligned tilt series ($60^\circ, 60^\circ, 57^\circ, 57^\circ$ and so on). We do not recommend collecting a single extra image on one side of the region of interest because there could be a systematic focus offset between the extra image and the tilt series image. In order to run CTFFIND on these extra images, set the `UseTrialsForCtffind` variable in the `reliion_prepare_subtomograms.py` script to `True`.

The last method to improve the 3D CTF model is to apply a linear, dose-dependent B-factor to the data (also see 14). Based on observations made for single-particle data sets 31, we increased the B-factor by $4 \text{ \AA}^2$ for each $1 \text{ e}^-/\text{\AA}^2$ of accumulated dose for both examples

described in this paper. Users may want to select a different value depending on the radiation sensitivity of their specimen. Change the Bfactor variable in the header of the python script. Note that the tilt-dependent scale factor and the position-dependent defocus of each particle will be calculated automatically for every image in the tilt series.

12| The command to reconstruct each 3D CTF model is written out in a run script called `do_all_reconstruct_ctfs.sh`, which is automatically generated in step 9. Run this script from the RELION project directory –

```
./do_all_reconstruct_ctfs.sh 200
```

In the above command, the parameter '200' determines the size of the generated CTF volumes in pixels. This should be set to the same values as the "Particle box size" in step 13, i.e. the size of the extracted sub-tomograms. The above script is a text file containing single commands on each line, and may thus be split into shorter text files for convenient parallelization, for example using the `split` command in Unix.

<CRITICAL> CTF models generated above will follow the IMOD convention with the tilt axis of the aligned tilt series being the Y-axis of the image. If the tilt axis is in a different orientation, then the Euler angles used to generate CTF models will need to be changed accordingly.

Particle Extraction

<CRITICAL> Particle extraction could also be performed using an external program like EMAN2, Bsoft or Xmipp. In this case, steps 13 may be skipped. However, steps 1-12 would still need to be adequately completed and a `.coords` text file would still need to be provided to generate position-dependent 3D CTF models for each sub-tomogram.

13| Extract every particle into a sub-tomogram by performing a 3D window operation on the corresponding tomogram. On the "**General**" job-type of the RELION GUI (Figure 3), provide the tomogram pixel size in Å, and the diameter of a spherical mask (in Å) that will be used for calculating the average and standard deviation of the background pixel values. Make sure this mask is slightly larger than the longest diameter of the particle. On the "**Particle extraction**" job-type provide the `all_tomograms.star` file (from step 7) as "micrograph STAR file"; set the "Coordinate-file suffix" to `.coords`; and the "rootname" to the same as the `RootName` variable in the `relion_prepare_subtomograms.py` script. Set the "Particle box size" (which is given in pixels) to reflect 150-200% of the particle's longest dimension. If the particles in the tomogram are black, invert the contrast of the sub-tomograms. If the tomograms were taken with a smaller pixel size than necessary for the target resolution, use the re-scaling option to downscale the extracted sub-tomograms, in order to save computational resources. Running this job (by pressing the "**Run!**" button on the GUI) will write out the sub-tomograms as individual `.mrc` files in the `./Particles/Tomograms/tomogram?` directories that were created by the python setup script in step 9.

14| Along with ordinary sub-tomograms (which are 3D volumes), 2D projections of all sub-tomograms (along the Z-axis) may also be calculated by again using the '**Particle**

extraction' job-type. To do this, run the same job as in step 13, but provide the extra option `--project3d` on the '**Additional arguments**' line of the '**Running**' tab. By setting the 'Extract rootname' on the 'I/O' tab to `subtomo_proj2d`, an additional `.star` file called `./subtomo_proj2d.star` will be generated, which lists the 2D projections of all sub-tomograms.

Sub-tomogram classification and refinement

15| The 2D projections may be used in reference-free 2D classification of the data, much like one would use in single-particle analysis, which is a computationally cheaper alternative to classification of the sub-tomograms 14, 32. On the '**I/O**' tab of the '**2D classification**' job-type, provide the `./subtomo_proj2d.star` as the input `.star` file, and set the 'Output rootname' for example to `Class2D/run1`. The number of classes will depend on the number of sub-tomograms and the expected heterogeneity in the data. As a rule of thumb, we typically use at least on average 30 sub-tomograms per class, and we hardly ever use more than 50-100 classes. Because the 2D projections do not have a CTF model, switch CTF-correction off on the '**CTF**' tab. Go to the '**Optimisation**' tab to select parameters for the classification. We typically perform 25 iterations, and use a regularisation parameter of 2 for 2D classification 15. Typically, we mask particles with zeros, and provide a limit on the resolution of around 10-15 Å to include in the expectation (E-) step of the algorithm. On the '**Sampling**' tab, the default parameters are suitable for most cases. Change angular and translational search ranges, if desired. Execute this job by pressing the '**Run!**' button. This will result in multiple output files for each iteration in the `./Class2D` directory. The final class averages are stored in a `.mrcs` stack file called `./Class2D/run1_it025_classes.mrcs`. The `./Class2D/run1_it025_model.star` file contains information about the final classes, such as their relative size and the estimated resolution of each class average.

16| Visualize the resulting 2D class averages from step 15 using the '**Display**' button on the GUI and selecting the `Class2D/run1_it025_model.star` file. In the subsequent pop-up window, it is useful to reverse-sort the classes on `rlnClassDistribution`, which will place the largest class averages on the top of the display window. Select good-looking class averages (which reveal recognizable protein features, e.g. see Figure 4A-B) by double-clicking the class averages, which will put a red border around the image. Save the new `.star` file containing only the particles that correspond to the good classes by right-clicking the display window, and selecting the option 'Save STAR with particles from the selected classes' option from the pop-up menu as `subtomo_proj2d_sel.star`.

17| Use the selected 2D classes to write a `.star` file containing only the corresponding sub-tomograms. We again provide a python script on the RELION wiki to facilitate this. Run this script from the project directory –

```
python relion_2Dto3D_star.py subtomo_proj2d_sel.star particles_subtomo.star
```

This script will take two input `.star` files; the first is the `subtomo_proj2d_sel.star` file containing selected particles from step 16 and the second is the `particles_subtomo.star` file containing all sub-tomograms from step 9. This script will write out the subset of the sub-tomograms selected in step 16 into `output_2Dto3D.star`.

18| Use the particle subset generated in step 17 to get a sub-tomogram averaging structure using the '**3D auto-refine**' job-type. The output rootname on the '**I/O**' tab could for example be `./Refine3D/run1`. Whereas single-particle analysis requires an input 3D reference map, sub-tomogram averaging in either the 3D auto-refine or 3D-classification (see next step) job-types may be performed without an initial reference by stating None at the 'Reference map' entry on the '**Reference**' tab. If you are using an external 3D reference, indicate whether it is on the correct absolute greyscale. In general, maps created by RELION from the same data will be on the correct greyscale, whereas maps coming from elsewhere may not be. Back on the '**I/O**' tab, provide the resolution of an initial low-pass filter that will be applied to the input 3D map. To prevent model bias, we typically use relatively harsh filters, e.g. in the range of 40-100 Å. On the '**CTF**' tab, turn 'on' CTF correction (provided steps 8-12 were performed). If CTF-correction is desired and a 3D map is provided as initial reference, then indicate whether the input map has been CTF-corrected. For maps from RELION that have been CTF-corrected (depending on whether CTF correction was performed in previous runs), and also for maps that were generated from an atomic model (i.e. reflecting that they do not suffer from CTF artefacts), answer "Yes" to the question "Has reference been CTF corrected?". Go to the "**Auto-sampling**" tab to set the initial search step and range. The angular and offset sampling rates and ranges will only be used in the first several iterations. After that the auto-sampling algorithm will automatically use finer samplings and smaller ranges until the refinement converges 16. The default parameters will be suitable for most projects, perhaps with the exception of particles with icosahedral symmetry, for which initial angular sampling rates of 3.7 degrees, and local angular searches from 0.9 degrees may yield better results. Note that the auto-refinement will divide the data into two random half-sets, each of which will be refined independently, so that in step 23 "gold-standard" resolution estimates may be calculated 33, 34. Executing this job by pressing the '**Run!**' button will output multiple files for each iteration in the `./Refine3D` directory. The final `./Refine3D/run1_half[1,2]_class001_unfil.mrc` files will be used in the post-processing as described in step 23.

? TROUBLESHOOTING

19| Once a 3D reconstruction has been obtained using sub-tomogram averaging, classify the sub-tomograms using the '**3D classification**' job-type on the GUI to detect different conformational states of the specimen. On the '**I/O**' tab select the original `particles_subtomo.star` file, or the `output_2Dto3D.star` from the 2D classification described in step 17. The output rootname could for example be `Class3D/run1`. Because of computational costs, we often use fewer classes for 3D classification than for 2D classification, with typical values in the range of 3-10. On the '**Reference**' tab, use the final reconstruction from step 18 (`./Refine3D/run1_class001.mrc`) as the reference map. This map is now on the correct absolute greyscale, and a similar initial low-pass filter as in step 18 may be applied. On the '**CTF**' tab, indicate that the reference has been CTF-corrected (if this was indeed done in step 18) and select 'Do CTF correction' to use the combined missing wedge and CTF model (Figure 2A-B). On the '**Sampling**' tab, the default parameters are again suitable for most projects. Edit the entries in this tab if any changes are desired. Go to the '**Optimisation**' tab to enter parameters for the refinement. We typically use similar

options as for 2D classification runs, with the exception of the regularisation parameter, which is set to 4 for 3D classification. Execute this job by pressing the '**Run!**' button. This will result in multiple output files for each iteration in the `./Class3D` directory. The final reconstructions for each class are stored in `.mrc` files called `./Class3D/run1_it025_class001.mrc`. The `./Class3D/run1_it025_model.star` file contains information about the final classes, such as their relative size and the estimated resolution of each reconstruction.

20| The user again needs to decide which of the classes look good. Use the '**Display**' button on the GUI to show 2D slices through each of the 3D maps (e.g. select `./Class3D/run1_it025_class001.mrc` to visualize the first class). Visualization of all classes together in UCSF Chimera 35 is also useful. Good classes typically yield expected protein-like features in the 3D maps (e.g. Fig 5E-F), whereas bad classes are often noisy both inside the region of the particle and in its surrounding solvent background (e.g. Fig 5G). Resolution in itself is not necessarily a good indicator of class quality, as smaller classes will be calculated to lower resolution by the intrinsic filtering of the regularised likelihood algorithm.

21| Once good classes have been identified, select the `./Class3D/run1_it025_model.star` file from the '**Display**' button on the GUI to display (central slices through) the maps of all classes. Save a `.star` file with the particles belonging to the good classes as described in step 16.

22| Sometimes, performing steps 19-21 multiple times, where selected classes from one run are re-classified in a next run, helps to obtain more homogeneous classes. Thus, refine each class as identified in steps 19-21 again to high-resolution using the 3D auto-refinement as described in step 18.

Post-processing and resolution estimation

23| After completion of the final 3D auto-refinement for each class, use the '**Post-processing**' job-type to obtain resolution estimates that have been corrected for the influence of a solvent mask 36, to calculate for the modulation transfer function (MTF) of the detector, and to sharpen the final map. On the '**I/O**' tab, one of the two unfiltered half-maps (calculated from the two half data sets described in step 18) should be provided as `./Refine3D/run1_half1_class001_unfil.mrc`; the output rootname could be set to `run1_post`. Go to the '**Mask**' tab, to provide parameters for the calculation of an automated solvent mask (automask). The average of the two half-maps will be used to make a binary mask at the specified initial threshold; extended by the provided number of pixels; and finally a cosine-shaped soft edge with the specified width will be added to the mask. Choose an initial threshold value such that the automask does not contain isolated white regions in the solvent area. Often a good estimate for the initial threshold value is the threshold at which a display of the `./Refine3D/run1_class001.mrc` map in UCSF Chimera is noise-free in the region around the particle. Once a suitable automask has been created, it can also be provided as input on the '**Mask**' tab instead of calculating a new automask in subsequent Post-processing runs. Go to the '**Sharpen**' tab, to provide a MTF curve for the detector. Curves for some detectors may be downloaded from the RELION wiki. For other detectors, the manufacturer may provide MTF curves. If no curve is available, this entry may also be left empty, in

which case MTF-correction will be emulated by additional B-factor sharpening. For maps with resolutions beyond 9-10 Å, automated B-factor sharpening 37 may be performed. Alternatively, a user-defined value may be provided. Go to the '**Filter**' tab if you want to skip FSC-weighting. Typically, this option is not used. Execute the post-processing step by pressing the '**Run!**' button. This will generate the final map (`./Refine3D/run1_post.mrc`), the automask (`./Refine3D/run1_post_automask.mrc`) and a .star file with the applied B-factor, the resolution estimate and the corrected FSC curve (`./Refine3D/run1_post.star`). In addition, the corrected FSC curve will be written out as a file called `./Refine3D/run1_post_fsc.xml`, which can be directly uploaded to the EMDb.

? TROUBLESHOOTING

24| Because many macromolecular complexes are inherently flexible, even after classification data sets will often still contain some extent of structural heterogeneity. This will lead to local variations in resolution in the refined map. To estimate local resolution variations, the '**Local-resolution**' job-type implements a wrapper to the ResMap program 38. On the '**I/O**' tab, again provide one of the two unfiltered half-maps, and provide the range and step size of the resolutions to be tested by ResMap. One typically does not change the default P-value. ResMap will provide much more reliable resolution estimates if one provides a suitable mask. The automask calculated in the previous step typically performs well. Click the '**Run!**' button to launch the GUI from the ResMap programs. In most cases one can just hit '**Continue**' inside ResMap, i.e. without adjusting any of its parameters. The result is a file called `./Refine3D/run1_half1_class001_unfil_resmap.mrc`, which can be used inside UCSF Chimera to colour the `./Refine3D/run1_post.mrc` map, using menu options Tools -> Surface Color -> by volume data value.

25| Sometimes, a refined map still shows signs of large amounts of structural heterogeneity, i.e. it contains regions of relatively low local resolution. In this case, perform a focused classification on that specific region. Provide the `./Refine3D/run1_data.star` file as input for a new 3D classification run as explained under step 19. In this run, one could use a solvent mask that is only white in the disordered region. This mask can be generated within RELION using the command `relion_mask_create` or using an external program such as `beditimg` from Bsoft 21 or `e2proc3d.py` from EMAN2 20. Also, one could skip orientational searches (through the corresponding option on the '**Sampling**' tab), or one could perform relatively fine, but local angular searches around the input orientations. One could iterate two or more times through steps 19-25.

Timing

The time taken for the procedure depends approximately linearly on the number of tomograms, the number of sub-tomograms, and it is inversely proportional to the number of processors used on the cluster. Here, for the HBV capsids, we analyzed a dataset containing 15 tomograms (each occupying ~75 Gb of hard disk space). 1851 capsid particles were extracted in boxes of 240^3 pixels from these tomograms and data was processed on a computing cluster with 4 hyper-threaded 12-core Intel Xeon nodes at 2.9 GHz, each with at least 32 Gb of RAM. The nodes were inter-connected by a 10 Gb Ethernet network. CTF

volume reconstruction (distributed over 15 cores, 1 per tomogram) took 2 hours. Sub-tomogram extraction and projection of the extracted sub-tomograms took 30 minutes using a single core. 2D classification of the data in 20 classes took ~3 hours on 2 nodes. 3D auto-refinement was concluded in 24 hours, using 4 nodes. Postprocessing was performed on a single core in 10 minutes.

For the 80S ribosomes, 7 tomograms were used and 3120 ribosomes were extracted in boxes of 200^3 pixels and projected into 2D in 20 minutes using a single core. CTF volume reconstruction took 30 minutes on all nodes; 2D classification into 20 classes took 30 minutes on 2 nodes; 3D classification into 3 classes took 10 hours on 4 nodes; and 3D auto-refinement lasted 50 hours on 4 nodes.

Troubleshooting

Troubleshooting guidelines can be found in Table 1.

Anticipated results

We illustrate the results of this protocol for two test data sets. The first set comprised 15 tomograms that were collected on a sample of purified hepatitis B virus (HBV) capsids. For this data set, two extra trial images on either side of the region of interest were used for CTF estimation. After optimizing the input parameters for CTFFIND (defocus search range, resolution search range and box size), CTF estimation from these extra images was found to be adequate even at high tilts (Figure 2D). Thus, the arithmetic mean of the defoci were applied to the tilt series image at every tilt. The coordinate files for the HBV capsids were obtained by automated picking using the template matching routines in the program MolMatch 8. In addition to true HBV capsids, this program also picked up 10 nm gold fiducials and other undesirable features. Using both steps 13 and 14, HBV capsid particles were extracted in 240^3 pixel sub-tomograms, as well as the corresponding 2D projections along the Z-axis. 2D classification, as described in steps 15-17 readily separated HBV capsids from false positives of the automated picking procedure (Figure 4A). 3D auto-refinement (step 18), starting from random orientation assignments for all sub-tomograms, followed by post-processing (step 23), yielded a final map to a resolution of 9.4 Å (Figures 4C and 5A) from 1851 particles where secondary structure elements like α -helices were clearly resolved (Figure 5C).

The second test data set of 7 tomograms was collected on a sample of purified 80S ribosomes from *S. cerevisiae*. We have deposited this data set together with the results described here at the EMPIAR data base 39 under accession number EMPIAR-10045. The initial tomograms, the corresponding aligned tilt series, as well as the .tilt, .order, .trial and .coords files as described in steps 1-7 are all stored in the ./Tomograms subdirectory of the EMPIAR entry. The results of all other steps are stored in the ./AnticipatedResults subdirectory of the EMPIAR entry. Thereby, novice users can follow the exact steps described above to replicate, and compare with, the results described here.

CTF estimation was attempted using the same procedures as for the HBV capsid data set. However, even after optimizing the inputs of CTFFIND, CTF estimation from high-tilt

images was inaccurate in some instances. We believe that this is due to increased specimen thickness due to the deposition of a layer of carbon on the grid during sample preparation. Thus, we used the average of the defoci from the lower tilt images and applied this average to all high tilt images in step 11. Manually picked particles were extracted as sub-tomograms, as well as 2D projections along the Z-axis. 2D classification of the projected sub-tomograms revealed multiple different views of the ribosomes, and some small classes that needed to be discarded (Figure 4B). Analysis of the distribution of the ribosome particles in tomograms showed that this sample often contained two layers of ribosomes, one at the top and the other at the bottom of the ice layer. Overlap of the projected densities of some of these ribosomes, which would complicate single-particle analysis, is not a problem in the tomographic approach (Figure 6). In this case, 3D auto-refinement of the input 3120 particles, again starting from random orientations, followed by post-processing, led to a 13 Å reconstruction (Figures 4D and 5B). Typical features like grooves of RNA helices were clearly visible in this map (Figure 5D). Using the output of the 3D auto-refinement we also conducted 3D classification (as in step 19) of the entire data set into 3 classes (Figure 5E-G). Although we could not identify different ratcheted states of the ribosomes, we identified a small subset of particles (Class 3, Figure 5G) that gave rise to a poor average. Removing these particles from the data set and subsequent 3D auto-refinement resulted in a somewhat cleaner output map, albeit at the same measured resolution. To assess local resolution variations in this map, ResMap analysis was performed as explained in step 24. The resulting map shows somewhat lower resolution in the small subunit than in the large subunit (Figure 5H).

Acknowledgements

We thank Xiaochen Bai, Israel Sanchez Fernandez and K. Vinothumar for help with sample preparation; Jake Grimmert and Toby Darling for assistance with high-performance computing; Shaoxia Chen and Christos Savva for assistance with electron microscopy; and Jan Löwe for helpful discussions. This work was supported by funds from EMBO (ALTF 3-2013 and aALTF 778-2015 to T.A.M.B.) and the UK Medical Research Council (MC_UP_A025_1013 to S.H.W.S.).

References

1. Bai XC, McMullan G, Scheres SH. How cryo-EM is revolutionizing structural biology. *Trends Biochem Sci.* 2015; 40:49–57. [PubMed: 25544475]
2. Cheng Y. Single-Particle Cryo-EM at Crystallographic Resolution. *Cell.* 2015; 161:450–457. [PubMed: 25910205]
3. Baumeister W. Electron tomography: towards visualizing the molecular organization of the cytoplasm. *Curr Opin Struct Biol.* 2002; 12:679–684. [PubMed: 12464323]
4. Briggs JA. Structural biology in situ--the potential of subtomogram averaging. *Curr Opin Struct Biol.* 2013; 23:261–267. [PubMed: 23466038]
5. Beck M, Luci V, Förster F, Baumeister W, Medalia O. Snapshots of nuclear pore complexes in action captured by cryo-electron tomography. *Nature.* 2007; 449:611–615. [PubMed: 17851530]
6. Grünewald K, et al. Three-dimensional structure of herpes simplex virus from cryo-electron tomography. *Science.* 2003; 302:1396–1398. [PubMed: 14631040]
7. Briggs JA, et al. Structure and assembly of immature HIV. *Proc Natl Acad Sci U S A.* 2009; 106:11090–11095. [PubMed: 19549863]
8. Förster F, Medalia O, Zauberman N, Baumeister W, Fass D. Retrovirus envelope protein complex structure in situ studied by cryo-electron tomography. *Proc Natl Acad Sci U S A.* 2005; 102:4729–4734. [PubMed: 15774580]

9. Nicastro D, et al. The molecular architecture of axonemes revealed by cryoelectron tomography. *Science*. 2006; 313:944–948. [PubMed: 16917055]
10. Schur FK, et al. Structure of the immature HIV-1 capsid in intact virus particles at 8.8 Å resolution. *Nature*. 2015; 517:505–508. [PubMed: 25363765]
11. Pfeffer S, et al. Structure of the native Sec61 protein-conducting channel. *Nature communications*. 2015; 6:8403.
12. Rigort A, Plitzko JM. Cryo-focused-ion-beam applications in structural biology. *Arch Biochem Biophys*. 2015; 581:122–130. [PubMed: 25703192]
13. Hrabe T, et al. PyTom: a python-based toolbox for localization of macromolecules in cryo-electron tomograms and subtomogram analysis. *J Struct Biol*. 2012; 178:177–188. [PubMed: 22193517]
14. Bharat TA, Russo CJ, Löwe J, Passmore LA, Scheres SH. Advances in Single-Particle Electron Cryomicroscopy Structure Determination applied to Sub-tomogram Averaging. *Structure*. 2015; 23:1743–1753. [PubMed: 26256537]
15. Scheres SH. A Bayesian view on cryo-EM structure determination. *J Mol Biol*. 2012; 415:406–418. [PubMed: 22100448]
16. Scheres SH. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *J Struct Biol*. 2012; 180:519–530. [PubMed: 23000701]
17. Bharat TA, et al. Cryo-electron microscopy of tubular arrays of HIV-1 Gag resolves structures essential for immature virus assembly. *Proc Natl Acad Sci U S A*. 2014; 111:8233–8238. [PubMed: 24843179]
18. Bharat TA, et al. Structure of the immature retroviral capsid at 8 Å resolution by cryo-electron microscopy. *Nature*. 2012; 487:385–389. [PubMed: 22722831]
19. Bartesaghi A, Lecumberry F, Sapiro G, Subramaniam S. Protein secondary structure determination by constrained single-particle cryo-electron tomography. *Structure*. 2012; 20:2003–2013. [PubMed: 23217682]
20. Tang G, et al. EMAN2: an extensible image processing suite for electron microscopy. *J Struct Biol*. 2007; 157:38–46. [PubMed: 16859925]
21. Heymann J, Belnap D. Bsoft: image processing and molecular modeling for electron microscopy. *J Struct Biol*. 2007; 157:3–18. [PubMed: 17011211]
22. Castano-Diez D, Kudryashev M, Arheit M, Stahlberg H. Dynamo: a flexible, user-friendly development tool for subtomogram averaging of cryo-EM data in high-performance computing environments. *J Struct Biol*. 2012; 178:139–151. [PubMed: 22245546]
23. Huiskonen JT, et al. Electron cryotomography of Tula hantavirus suggests a unique assembly paradigm for enveloped viruses. *J Virol*. 2010; 84:4889–4897. [PubMed: 20219926]
24. Kremer JR, Mastronarde DN, McIntosh JR. Computer visualization of three-dimensional image data using IMOD. *J Struct Biol*. 1996; 116:71–76. [PubMed: 8742726]
25. Crowther RA, Henderson R, Smith JM. MRC image processing programs. *J Struct Biol*. 1996; 116:9–16. [PubMed: 8742717]
26. Agulleiro JJ, Fernandez JJ. Fast tomographic reconstruction on multicore computers. *Bioinformatics*. 2011; 27:582–583. [PubMed: 21172911]
27. Mastronarde DN. Automated electron microscope tomography using robust prediction of specimen movements. *J Struct Biol*. 2005; 152:36–51. [PubMed: 16182563]
28. Hall SR. The Star File - a New Format for Electronic Data Transfer and Archiving. *J Chem Inf Comp Sci*. 1991; 31:326–333.
29. Mindell JA, Grigorieff N. Accurate determination of local defocus and specimen tilt in electron microscopy. *J Struct Biol*. 2003; 142:334–347. [PubMed: 12781660]
30. Eibauer M, et al. Unraveling the structure of membrane proteins in situ by transfer function corrected cryo-electron tomography. *J Struct Biol*. 2012; 180:488–496. [PubMed: 23000705]
31. Scheres SH. Beam-induced motion correction for sub-megadalton cryo-EM particles. *Elife*. 2014; 3:e03665. [PubMed: 25122622]
32. Yu L, Snapp RR, Ruiz T, Radermacher M. Projection-based volume alignment. *J Struct Biol*. 2013; 182:93–105. [PubMed: 23410725]

33. Henderson R, et al. Outcome of the first electron microscopy validation task force meeting. *Structure*. 2012; 20:205–214. [PubMed: 22325770]
34. Scheres SH, Chen S. Prevention of overfitting in cryo-EM structure determination. *Nature methods*. 2012; 9:853–854. [PubMed: 22842542]
35. Pettersen EF, et al. UCSF Chimera--a visualization system for exploratory research and analysis. *J Comput Chem*. 2004; 25:1605–1612. [PubMed: 15264254]
36. Chen S, et al. High-resolution noise substitution to measure overfitting and validate resolution in 3D structure determination by single particle electron cryomicroscopy. *Ultramicroscopy*. 2013; 135:24–35. [PubMed: 23872039]
37. Rosenthal PB, Henderson R. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *J Mol Biol*. 2003; 333:721–745. [PubMed: 14568533]
38. Kucukelbir A, Sigworth FJ, Tagare HD. Quantifying the local resolution of cryo-EM density maps. *Nature methods*. 2014; 11:63–65. [PubMed: 24213166]
39. Patwardhan A, et al. A 3D cellular context for the macromolecular world. *Nat Struct Mol Biol*. 2014; 21:841–845. [PubMed: 25289590]

Sub-tomogram analysis in RELION

Figure 1. Workflow of the image processing protocol.

A schematic representation of the recommended workflow for sub-tomogram analysis using RELION presented in this protocol. The main difference between single-particle analysis and subtomogram analysis in RELION is related to CTF estimation and the new 3D CTF model. This 3D CTF model also compensates for the missing wedge, and is used in both 3D classification as well as 3D auto-refinement. Steps highlighted in orange are unchanged from the single-particle analysis workflow.

Figure 2. CTF estimation for the 3D CTF model.

(A) The unweighted 3D CTF model used in RELION. This model is constructed by placing the 2D CTFs of each image in the tilt series into a 3D volume in Fourier-space, with the correct orientation depending on the tilt angle. Therefore this model also compensates for the missing wedge. (B) The weighted 3D CTF model used in RELION. The weighted model accounts for increase in noise at high tilts, and for radiation-induced damage. The volume is coloured from red at low resolution to dark blue at the highest obtainable (Nyquist) resolution. (C) A diagnostic output file of CTFFIND3 from a low-tilt tilt series image. There is no visible radiation-induced motion, and many Thon rings are visible making CTF estimation accurate. (D) Corresponding diagnostic file from a high-tilt image. Fewer Thon rings are visible due to increased specimen thickness. CTF estimation in this case is adequate but not as accurate as C. (E) A diagnostic file from a high-tilt image where no Thon rings are visible. CTF estimation is not possible from this image, and it should either

be removed from the tilt series or the data collection strategy should be modified to include the recording of additional images for CTF estimation on either side of the target region 14.

Figure 3. The RELION-1.4 graphical user interface.

After CTF parameters have been estimated for each particle in each image of the tilt series and 3D CTF models have been reconstructed, the actual tasks of sub-tomogram analysis may all be performed using the RELION graphical user interface. The 3D auto-refine page of this user interface is shown. The white column on the left shows different 'job-types', which are ordered according to the natural workflow from top to bottom. On the main panel, the '3D auto-refine' job-type is shown. This job-type has tabs for "I/O", "Reference", "CTF", "Optimization", "Auto-sampling", "Movies", and "Running" where users should enter the input parameters as described in the main text. The "Display", "Print command" and "Run!" buttons are used to view images, commands and launch jobs, respectively.

Figure 4. 2D Classification and initial model generation.

(A) 2D classification of projected sub-tomograms of the HBV capsid particles. Particles were selected from a template matching procedure and 2D classification helped in removing bad particles, for example ones that correspond to 10 nm gold fiducials. Good classes that were selected for further processing are marked with blue dots. (B) 2D classification of projected sub-tomograms of *S. cerevisiae* 80S ribosomes. These data were picked manually in IMOD. Classes of ribosomes selected for refinement are marked with blue dots. (C) Reference-free refinement of the HBV capsid data set. Sub-tomograms were assigned random Euler angles initially (in iteration 0) and then refinement was commenced. (D) Reference-free refinement of the *S. cerevisiae* 80S ribosome particles, again starting from random orientations. Initial models described in panels (C-D) may then be used to begin 3D refinements within RELION.

Figure 5. 3D auto-refinement and classification using the regularized-likelihood algorithm in RELION.

(A) Output of the 3D auto-refinement procedure from RELION for the HBV capsid dataset. (B) Output of 3D auto-refinement for the 80S ribosome data set (This map has been deposited at the EMDB under the accession number EMD-3228). The scale bar shown applies to panels A-B and E-H. (C) Secondary structure features (α -helices) are resolved in the HBV capsid map. Fitted atomic co-ordinates into the sub-tomogram average highlight the positions of the helices. (D) RNA helices resolved in the 80S ribosome map. The atomic co-ordinates have been fitted into this map as rigid bodies for visualization. (E-G) 3D

classification of the ribosome data set (with the combined 3D missing wedge and CTF model applied) into three classes reveals a subset of particles (~15% of the data set) in G that show a poor sub-tomogram average. (H) Removing these particles in a second 3D auto-refinement leads to a cleaner map. The result of ResMap is plotted onto the final density showing somewhat lower resolution in the small subunit of the ribosome. The colour map is defined from blue (10 Å) to red (18 Å).

Figure 6. Sub-tomogram analysis of particles at different Z-heights.

(A) The *S. cerevisiae* 80S ribosomes particles were found to localize at either the air-water or the carbon-water interface. This panel shows a small population of ribosomes at a Z-slice corresponding to the air – water interface. The scale bar applies to all the panels in the figure. (B) A tomographic slice at the carbon-water interface showing a surface packed with ribosomes. (C) Same slice as panel A is shown (with a transparency) overlaid with a plot of the centres of all ribosomes in the tomogram. A green cone is placed at the centre of a ribosome that is sorted into a good class (see Figure 4B) and a red cone is placed at the

centre of a ribosome sorted into a bad class. This figure shows that in most tilt series images, the signal from ribosomes in the top layer is superimposed with the signal from ribosomes in the bottom layer, therefore tomography and sub-tomogram analysis is ideal for studying this sample. (D) The same picture as panel C rotated to show the 3D arrangement of the sample. (E) A view of the tomogram along the XY plane with the plot of the centres of all ribosomes shown as in panels C-D. The edges of the tomogram have been demarcated with a solid black line.

Table 1
Troubleshooting

Step	Problem	Possible reason	Solution
All	RELION executables not found	RELION was not installed correctly, or the shell environment variables were not set correctly	Install RELION 1.4 on your system, and confirm that the shell environment variables are set correctly as indicated on the RELION wiki.
All	IMOD executables not found	IMOD is not installed on the system.	Install IMOD version 4.7 or newer on your system.
8	CTF correction is not desired.	Target resolution is lower than spatial resolution of the first zero of the CTF, or it is impossible to estimate and compensate for the CTF.	Set the variable SkipCTFCorrection in the <code>relion_prepare_subtomograms.py</code> script to True. In this case a wedge shaped 3D 'CTF' model will be created that is weighted appropriately depending on the tilt angle and the accumulated radiation.
9	The python setup script crashes with the complaint "File not found".	One or more of the input files are missing.	Please ensure that all the files needed for steps 1-7 are present in the correct locations. The .tilt file may be omitted if the .mrcs file has the correct angles stored in its extended header.
9	The python script exits with the warning that the number of tilt angles and the number of images is different.	Some images were omitted in the tomogram reconstruction, but the tilt angle values were not updated.	Use the tilt angles for only the images that are included in the aligned tilt series and thus in tomogram reconstruction.
10	CTF cannot be estimated for some images of the tilt series.	Low signal to noise in the image leading to poor visibility of Thon rings.	Either follow one of the strategies in step 11 for accurate CTF estimation, or remove the offending images from the aligned tilt series stack. If the latter option is followed, the tomogram would have to be regenerated with that image removed.
18-19	The job crashes with an error 'Cannot allocate memory'	The job requires more computer memory than available, or too many MPI processes were run on the same node.	If multiple MPI processes were run on a single computer, run multiple threads instead to share its available memory. More detailed instructions on the hybrid parallelisation scheme are on the RELION wiki.
23	The post-processing program complains that the masked FSC never drops below 0.8.	The automasking procedure yielded an all-black, or otherwise unsuitable, mask.	Adjust the parameters of the automasking procedure, or make a mask outside the post-processing procedure. For the latter, one could use any program outside RELION, or the <code>relion_mask_create</code> program.
23	Even though estimated resolution of the reconstruction is high, the reconstruction looks lowpass filtered with no high-resolution features	The tomogram reconstruction was conducted with a SIRT algorithm.	Try using weighted back projection for tomogram generation and re-extract the sub-tomograms, or apply a negative (ad-hoc) B-factor in the post-processing.