

RELION-3: new tools for automated high-resolution cryo-EM structure determination

Jasenko Zivanov*¹, Takanori Nakane*¹, Björn Forsberg*², Dari Kimanius², Wim J.H. Hagen³, Erik Lindahl^{†2,4}, and Sjors H.W. Scheres^{‡1}

¹MRC Laboratory of Molecular Biology, Cambridge, United Kingdom

²Department of Biochemistry and Biophysics, Science for Life Laboratory, Stockholm University, Stockholm, Sweden

³Structural and Computational Biology Unit and Cryo-Electron Microscopy Service Platform, EMBL, Heidelberg, Germany

⁴Department of Applied Physics, Swedish e-Science Research Center, KTH Royal Institute of Technology, Stockholm, Sweden

September 19, 2018

Abstract

Here, we describe the third major release of RELION. CPU-based vector acceleration has been added in addition to GPU support, which provides flexibility in use of resources and avoids memory limitations. Reference-free autopicking with Laplacian-of-Gaussian filtering and execution of jobs from python allows non-interactive processing during acquisition, including 2D-classification, *de novo* model generation and 3D-classification. Per-particle refinement of CTF parameters and correction of estimated beam tilt provides higher-resolution reconstructions when particles are at different heights in the ice, and/or coma-free alignment has not been optimal. Ewald sphere curvature correction improves resolution for large particles. We illustrate these developments with publicly available data sets: together with a Bayesian approach to beam-induced motion correction it leads to resolution improvements of 0.2-0.7 Å compared to previous RELION versions.

*These authors contributed equally

[†]erik.lindahl@dbb.su.se

[‡]scheres@mrc-lmb.cam.ac.uk

1 Introduction

Macromolecular structure determination by single-particle analysis of electron cryo-microscopy (cryo-EM) images has undergone rapid progress in recent years. Before 2010, cryo-EM structures to resolutions beyond 4 Å had only been obtained for large icosahedral viruses, e.g. [Jiang et al., 2008, Yu et al., 2008, Zhang et al., 2008]. Since then, cryo-EM structures to resolutions that are sufficient for *de novo* atomic modelling have been achieved for a wide range of samples [Fernandez-Leiro and Scheres, 2016], and several structures beyond 2 Å resolution have been reported [Merk et al., 2016, Bartesaghi et al., 2018]. In 2012, the availability of the first prototypes of direct electron detectors, which recorded images with unprecedented signal-to-noise ratios [McMullan et al., 2009], represented a crucial step forward. This advance was closely followed and partially overlapped with important improvements in image processing software. Together, they led to what was termed the 'resolution revolution' in cryo-EM structure determination [Kühlbrandt, 2014].

Since the start of this revolution, many new software solutions have been introduced. Some new programs exploited specific opportunities provided by the new detectors, like the correction of beam-induced motions through movie processing [Li et al., 2013, Rubinstein and Brubaker, 2015, Abrishami et al., 2015, McLeod et al., 2017, Zheng et al., 2017]. Other programs tackled individual steps in the single-particle workflow, like fast and robust estimation of the contrast transfer function (CTF) [Rohou and Grigorieff, 2015, Zhang, 2016], local-resolution estimation [Kucukelbir et al., 2014], or on-the-fly processing of images while they are being acquired at the microscope [de la Rosa-Trevín et al., 2016, Biyani et al., 2017]. In the mean time, also general-purpose programs that already existed, like EMAN2 [Tang et al., 2007], FREALIGN [Grigorieff, 2007] or IMAGIC [van Heel et al., 1996], were further improved and used for high-resolution structure determination. Both old and new programs benefitted from the increase in image signal-to-noise ratios provided by the new detectors.

The availability of the first direct-electron detectors coincided with the introduction of the RELION software [Scheres, 2012b]. This computer program introduced an empirical Bayesian approach, in which optimal Fourier filters for alignment and reconstruction are derived from the data in a fully automated manner [Scheres, 2012a]. This strongly reduced the need for user expertise in tuning *ad-hoc* filter parameters compared to existing programs, and probably contributed to the rapid uptake of cryo-EM structure determination by non-expert groups. Nowadays, RELION is widely used [Patwardhan, 2017], and similar statistical approaches have been implemented in software packages that were developed since: cryoSPARC [Punjani et al., 2017] uses the same regularised likelihood optimisation target as RELION, and SPHIRE [Moriya et al., 2017] and cisTEM [Grant et al., 2018a] implement likelihood-based classification methods.

RELION underwent a major change in 2016 with the release of its second version, which comprised a pipelined approach for the entire single-particle workflow [Fernandez-Leiro and Scheres, 2016] and used acceleration on graphical processing units (GPUs) to reduce its computational costs [Kimanius et al.,

2016]. In addition, its functionality was expanded with the incorporation of sub-tomogram averaging [Bharat et al., 2015] and helical reconstruction [He and Scheres, 2017]. In this paper, we introduce version 3 of RELION. In the Methods section, we describe algorithmic improvements and changes in implementation that have been made throughout the single-particle pipeline in RELION-3. In the Results and Discussion section we illustrate how these changes allow for improved automation, fast accelerated reconstruction on normal hardware (CPUs) in addition to GPUs, and how the new features enable higher resolution reconstructions compared to previous releases.

2 Methods

2.1 Laplacian-of-Gaussian auto-picking

The recommended procedure for the selection of particles from micrographs in previous versions of RELION was to calculate templates for reference-based auto-picking from particles that were manually picked in a subset of the micrographs [Scheres, 2015]. To allow non-interactive, on-the-fly processing, we implemented a reference-free particle picking method based on the Laplacian-of-Gaussian (LoG) filter. Our method is related to the approach in DoGpicker [Voss et al., 2009], which uses a Difference-of-Gaussian (DoG) filter. Similar to the implementation in swarmPS [Woolford et al., 2007], we apply a LoG-filter with an optimum response for blobs of size d by multiplying the Fourier transform of a micrograph by

$$\frac{|k|^2}{\sigma^2} \exp\left(\frac{|k|^2}{-2\sigma^2}\right) \quad (1)$$

where $k \in \mathbb{R}^2$ is the image frequency, $\sigma = 2/d$, and the division by σ^2 outside the exponentiation provides a normalised response of the filter for different values of d .

In order to deal with elongated particles, the user specifies a minimum (d_{min}) and maximum (d_{max}) particle diameter. The program then applies four LoG-filters to detect areas in the micrograph with blobs that are smaller than d_{min} (with $d = d_{min}/n$ and $n = 2, 3, 4$ or 5); four filters to detect blobs that are larger than d_{max} (with $d = n * d_{max}$ and $n = 2, 3, 4$ or 5); and three filters to detect blobs within the user-specified particle size (with $d = d_{min}$, $d = (d_{min} + d_{max})/2$ and $d = d_{max}$).

Only pixels for which the highest LoG-filtered value over all eleven filtered micrographs (LoG_{best}) is above a user-controlled threshold and the corresponding blob size d_{best} is in the user-specified particle diameter range are considered for particle picking. Particles are picked iteratively at the highest remaining LoG_{best} value, and around each picked particle one discards all pixels within a circle with diameter d_{best} . This process is repeated until no suitable pixels remain. To normalise the threshold for LoG_{best} over many different data sets,

the default threshold is set to the mean of all LoG_{best} values for which the corresponding d_{best} is in the user-specified particle diameter range, and deviations from the default are expressed in standard deviations of those same values.

2.2 CTF refinement

In RELION, per-micrograph CTF parameters are determined through wrappers to CTFFIND [Rohou and Grigorieff, 2015] or GCTF [Zhang, 2016]. These programs fit CTFs to the Thon rings visible in the power spectra of (patches of) micrographs. In RELION-3, we have implemented a program to refine the CTF parameters, i.e. to re-estimate defocus and astigmatism, using a 3D reference structure. This allows CTF estimation to exploit both the phases and the amplitudes of the experimental images, instead of having to rely exclusively on their power spectra. This approach thus produces significantly more reliable estimates, due to higher signal-to-noise ratios and also because it does not require separation of the Thon rings from the background intensity of the power spectrum. This is illustrated in Fig. 1. The increased stability of the CTF fitting can be leveraged to estimate independent defoci for individual particles. Similar functionality also exists in FREALIGN [Grigorieff, 2007] and CISTEM [Grant et al., 2018b].

Formally, CTF refinement in RELION-3 works as follows. Let V be the (complex-valued) Fourier-space representation of the reference structure. We assume the viewing direction and particle position to be known for each observed particle p . This allows us to project and shift the 3D reference, yielding V_p , the corresponding reference projection in 2D. The Fourier space representation of experimental particle image X_p is now assumed equal to V_p , multiplied by the CTF and distorted by noise. We can then estimate the optimal CTF from V_p and X_p by determining a set of CTF-parameters θ that minimises the squared difference between them:

$$\theta = \operatorname{argmin}_{\theta'} \left(\sum_{k,p} g_k (X_{p,k} - \operatorname{CTF}(\theta', k) V_{p,k})^2 \right) \quad (2)$$

$$= \operatorname{argmax}_{\theta'} \left(\sum_k g_k \operatorname{CTF}(\theta', k) \sum_p \operatorname{Re}(V_{p,k}^* X_{p,k}) \right), \quad (3)$$

where the asterisk indicates complex conjugation, k is the 2D Fourier-space frequency and per-frequency weights g_k are given by the FSC between two independently refined half-maps – this a heuristic that serves to attenuate the effect of frequencies where the reference is less reliable.

The value of the CTF for a spatial frequency k is defined as follows:

$$\operatorname{CTF}(k) = -\sin \left(\pi \lambda (\delta_0 + \delta_A \cos(2(\phi_k - \phi_A))) |k|^2 + \frac{\pi}{2} C_s \lambda^3 |k|^4 - \chi \right), \quad (4)$$

where λ is the wavelength of the electrons, δ_0 the average defocus, δ_A half the difference between the two effective defoci resulting from astigmatism, ϕ_k and ϕ_A are the azimuthal angles of k and the major astigmatism axis, respectively,

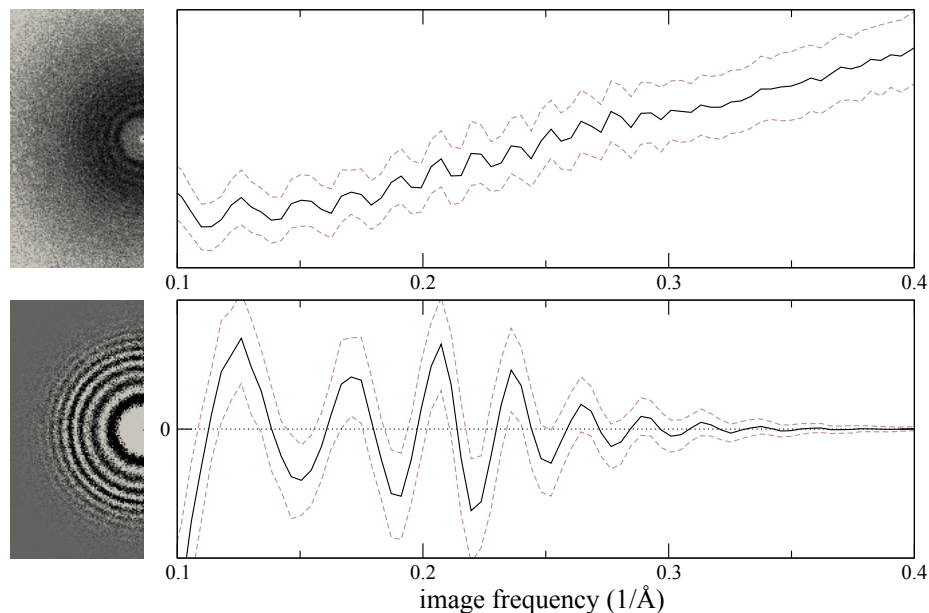


Figure 1: Comparison of the traditional power-spectrum based CTF estimation technique with reference-map based CTF refinement. **Top:** Sum over the power spectra $\|X_k\|$ of all polished particle images in micrograph EMD-2984_0000 of EMPIAR-10061 (left) and the averages and std. deviations of that sum over Fourier rings between 10\AA and 2.5\AA (right). The dashed grey curves are located one std. deviation of each ring above and below the average. The overall rise in spectral power as a function of frequency is a consequence of particle polishing. **Bottom:** The corresponding plot for the sum over the real components of $V_k^* X_k$, instead of the power spectra. Note that the latter term is not only centered around zero, obviating the need to estimate the background intensity, but it also offers a higher signal-to-noise ratio.

C_s is the spherical aberration constant of the microscope and χ the overall phase shift resulting from amplitude contrast and, if one is used, a phase plate.

The expression in Eq. 3 allows the CTF-parameters $\theta = \{\delta_0, \delta_A, \phi_A, C_s, \chi\}$ to be determined either as constant for all particles in an entire micrograph or for every particle separately. In the latter case, the second sum consists of one single term corresponding to the index p of the particle in question. Our implementation allows for the estimation of δ_0 on a per-particle basis, δ_A and ϕ_A on a per-particle or per-micrograph basis, and C_s and χ on a per-micrograph basis. In our experiments, we have found that for many data sets the best results are obtained by estimating only δ_0 for individual particles, while leaving all the other parameters unchanged.

Because the large number of CTF-parameters for a given dataset introduces the risk of overfitting, it is important to maintain a strict separation between the two independently refined half-sets. For that reason, the reference image V_p used to estimate the parameters of particle p is always obtained by projecting the reference map computed from the half-set to which particle p belongs.

2.3 Beam-tilt estimation and correction

A tilted beam alters the phase delays of the scattered beam components with respect to the unscattered reference beam. The delays ψ_1 and ψ_2 for a pair of corresponding Friedel mates are then no longer identical. We can decompose them into a symmetrical component $\psi_s = (\psi_1 + \psi_2)/2$ and an anti-symmetrical one, $\psi_a = (\psi_1 - \psi_2)/2$. The symmetrical part only induces astigmatism, which is already being modelled reasonably by our CTF model (Eq. 4). The anti-symmetrical part, however, induces axial coma, i.e. a phase shift ϕ in the respective image frequencies. The magnitude of ϕ as a function of tilt $b \in \mathbb{R}^2$ is given by [Glaeser et al., 2011, Zemlin et al., 1978]:

$$\phi(b) = 2\pi C_s \lambda^2 \langle k, b \rangle |k|^2, \quad (5)$$

where C_s is the spherical aberration constant of the microscope, λ the wavelength of the electrons and \langle, \rangle denotes the scalar product in 2D.

We have developed an efficient method to determine the tilt vector b from a set of observations of a particle and the corresponding high-resolution reference. Let $V_{j,k} \in \mathbb{C}$ denote the predicted complex amplitude of image frequency k of particle j (including the effects of the CTF) and $X_{j,k} \in \mathbb{C}$ the corresponding amplitude observed in the experimental image X . We aim to find a vector $b \in \mathbb{R}^2$ that minimises the squares difference between the observations and the prediction:

$$b = \operatorname{argmin}_{\beta} \sum_{j,k} |z_k(\beta) V_{j,k} - X_{j,k}|^2, \quad (6)$$

where the phase shift is expressed as a multiplication with a complex number: $z_k(\beta) = \exp(i\phi(\beta))$. Since the error in Eq. 6 is a sum of quadratic functions,

the equation can be transformed into:

$$b = \operatorname{argmin}_{\beta} \sum_k w_k |z_k(\beta) - q_k|^2 \quad (7)$$

with the per-pixel optima $q_k \in \mathbb{C}$ and weights $w_k \in \mathbb{R}$ given by:

$$q_k = \sum_i (X_{j,k} V_{j,k}^*) / \sum_i |V_{j,k}|^2 \quad (8)$$

$$w_k = \sum_i |V_{j,k}|^2, \quad (9)$$

where the asterisk indicates complex conjugation. The computation of q_k can be thought of as an averaging of the phase differences between the projected reference and the experimental images for Fourier-space pixel k over all the images. Because of the cyclical nature of phase angles, a naive average over the angles themselves would not produce the desired result.

This formulation allows us to first compute q_k and w_k for each pixel by only iterating over the dataset once. This takes less than one hour for most datasets. Once the q_k and w_k are known, we can fit the optimal b to them using an iterative non-linear method, such as Nelder-Mead downhill simplex [Nelder and Mead, 1965]. Due to the small number of k , this iterative optimisation then only takes seconds.

In addition, inspecting the phases of q_k allows us to ascertain that the effects observed are really those of a tilted beam. If we were to merely fit the optimal tilt vector b to the data directly, we would run the risk of fitting b to noise. Example plots of phase angles of q_k are shown in Fig. 6 - figure supplement 1. Furthermore, such phase-error plots make it possible to observe the effects of higher-order anti-symmetrical aberrations, such as three-fold astigmatism [Saxton, 1995, Krivanek and Stadelmann, 1995]. Although no methods are available to correct for these yet, we are currently investigating these.

Once the tilt vector b is known, it can be used to account for the expected phase error in X during refinement and reconstruction. This functionality has already been present in RELION since release 2.0, but no practical methods were available to estimate b until now.

2.4 Ewald sphere correction

We implemented the single-side band algorithm to correct for Ewald sphere curvature as proposed by Russo and Henderson [2018]. This implementation was made in the `relion_reconstruct` program, which was also parallelised using the message-passing interface (MPI) standard.

Briefly, one divides the Fourier transform of each extracted particle into two halves and multiplies each half with one of the following single-side band CTFs:

$$\begin{aligned} CTFP &= e^{+i(\psi+\pi/2)} \text{ and} \\ CTFQ &= e^{-i(\psi+\pi/2)}, \end{aligned} \quad (10)$$

where ψ is the phase shift of the electron wave caused by the objective lens, and the term $\pi/2$ is the phase shift upon scattering by the specimen. These complex multiplications leave the amplitudes of the Fourier components intact, but change the phases, resulting in a breakage of Friedel symmetry. To avoid discontinuities at the edge between the two halves of the Fourier transform, the multiplication is done in sectors by rotating the edge between the two halves. Two Fourier transforms, one for each side-band, are then obtained for each particle by combination of those parts of the Fourier transform that are furthest away from the edge in the different sectors. The number of sectors can be controlled by the user. By default it is set to two.

The two Fourier transforms of each particle sample the 3D transform in positions that are slightly above and below the central section. Both are back-transformed, masked in real space to remove noise in the area around the particle, and then transformed back into Fourier space before being inserted into their correct position in the 3D transform. The user has the option to use either upward or downward curvature of the Ewald sphere. The distance Δz^* from the central section is calculated as:

$$\Delta z^* = \lambda/2|k|^2. \quad (11)$$

Because an unknown number of mirroring operations may have been introduced during the data acquisition and image processing, the actual direction of curvature of the Ewald sphere is often unknown for a given data set. Therefore, the user may typically want to perform Ewald sphere correction with both directions. For this purpose, we implemented the `--reverse_curvature` option. If Ewald sphere curvature is a limiting factor for the resolution of a reconstruction, one direction of curvature should yield a higher resolution map than the other.

2.5 CPU vector acceleration

The large computational load of the cryo-EM pipeline requires RELION to be highly parallel. In RELION-2, a new code-path was introduced, which off-loaded the core computations onto graphical processing units (GPUs) [Kimanius et al., 2016]. This new path was by design less demanding of double precision arithmetics and memory, and provided a large improvement in data throughput by virtue of the greatly increased speed. However, the GPU acceleration is only available for cards from a single vendor, and it cannot use many of the large-scale computational resources available in existing centres, local clusters, or even researchers' laptops. In addition, the memory available on typical GPUs limits the box sizes that can be used, which could turn into a severe bottleneck for large particles. For RELION-3, we have developed a new general code path where CPU algorithms have been rewritten to mirror the GPU acceleration, which provides *dual* execution branches of the code that work very efficiently both on accelerators as well as the single-instruction, multiple-data vector units present on traditional CPUs. This has enabled lower-precision arithmetics, it

reduces memory requirements considerably, and it also makes it possible to exploit the very large amounts of memory that can be fitted on CPU servers. In addition, it provides a transparent source code where new algorithms may be accelerated immediately, and it enables a significant speedup compared to the *legacy* code path, which is nonetheless still present in RELION-3 for comparative purposes.

2.6 Python scripts for automated processing

As of release 2.0, RELION has used a pipelined approach, where the execution of individual command-line programs is organised into *jobs*. Multiple different types of jobs exist for the different tasks in the single-particle workflow, e.g. there are specific jobs for CTF estimation, particle picking, 2D classification and 3D refinement. Each job takes specific types of files, called *nodes* as input and produces other nodes as output. By using output nodes from one job as input for the next job, the entire process of structure determination is expressed as a directional graph. Parameters for jobs are typically provided through the GUI, where jobs can be executed or scheduled for later, possibly repeated, execution.

In RELION-3, we have expanded the functionality to schedule and execute jobs from the GUI to the command line by modifications in the `relion_pipeline` program. This has allowed the generation of (python) scripts to perform complicated, standardised tasks by execution of a series of jobs. All jobs executed by these scripts become part of the standard pipeline, and one can use the GUI to follow the process live, or to do other things in the same directory meanwhile. With RELION-3, we distribute two example scripts. We envision that these will serve as inspiration for scripts written by others to reflect their specific needs.

The first example script is called `bfactor_plot.py`. It automatically runs multiple 3D refinements on random subsets of an input data set, where each time the size of the subset is doubled from a user-specified minimum number of particles. The script then automatically generated a B-factor plot according to Rosenthal and Henderson [Rosenthal and Henderson, 2003], where $1/d^2$, with d being the resolution of each refinement, is plotted against the logarithm of the number of particles in the subset. The script also calculates the B-factor from the slope of a fitted line through the plotted points, and uses extrapolation of this line to make predictions about the number of particles needed to achieve higher resolutions.

The second example script is called `relion_it.py`. It is designed to provide on-the-fly feedback on the quality of micrographs, while they are being acquired at the microscope. The script repeatedly executes jobs for micrograph-based motion correction, CTF estimation, LoG-based auto-picking and particle extraction. In the mean time, the script keeps track of how many particles have been extracted, and can also execute 2D and/or 3D classification jobs in batches of a user-defined number of particles. If no 3D reference is available for 3D classification, an SGD initial model generation job will be executed to create one automatically. In addition, it is possible to replace the LoG-based auto-picking job by a template-based auto-picking job once a suitable 3D reference has been

obtained from 3D classification of the LoG-picked particles.

Although almost all parameters from the GUI are accessible through the `reliion_it.py` script, the procedure has been designed to rely on only a few dataset-specific parameters. Besides parameters describing the microscopy and computational setup, the most important parameters are d_{min} and d_{max} for the LoG-based auto-picking and box sizes for the extracted particles. The default thresholds for the LoG-based and template-based auto-picking jobs are set to relatively low values. This typically results in many false positives in the picking, which are then sorted out in the subsequent classifications. This way, completely non-interactive, on-the-fly feedback on data quality may be achieved.

2.7 Other developments

Besides the developments described above, RELION-3 comprises two major algorithmic advances that have already been described elsewhere.

Firstly, Bayesian particle polishing implements a Gaussian Process regression algorithm for estimating beam-induced motion tracks for individual particles, and an improved B-factor estimation algorithm for resolution-dependent weighting of individual movie frames. Although not strictly necessary, Bayesian polishing is typically performed after motion-correction at the micrograph (patch) level by MOTIONCOR2 [Zheng et al., 2017]. To expose the metadata of this algorithm better to the Bayesian polishing approach, we also implemented our own version of the MOTIONCOR2 algorithm. Unlike the UCSF implementation, our version does not use GPU-acceleration but runs on multi-core CPUs using OpenMP multi-threading. For $4k \times 4k$ movies, and using twelve cores, our program runs at comparable speeds to the UCSF implementation. More details about the Bayesian polishing and our implementation of the MOTIONCOR2 algorithm have been described by Zivanov et al [Zivanov et al., 2018].

Secondly, multi-body refinement implements an automated and iterative approach for focused refinement with partial signal subtraction on multiple independently moving parts of a complex [Nakane et al., 2018]. This approach is useful for improving densities of flexible multi-domain complexes, and yields insights into the type of inter-domain motions that are present in those.

In addition, many smaller programs and functionalities have been added throughout the single-particle processing workflow. We highlight the following. The motion correction, CTF estimation and auto-picking jobs now output PDF files with metadata plots for all micrographs. We have improved our 3D initial model generation by adopting a stochastic gradient descent (SGD) procedure that follows the implementation in cryoSPARC much more closely than our previous version. We achieved considerably speedups in 2D and 3D classification by using subsets of particles in the first few iterations, which was inspired by a similar implementation in cistem [Grant et al., 2018b]. We have added a program called `reliion_star_handler` that implements a range of useful operations for metadata files in the STAR format [Hall, 1991]. And we have implemented a program called `reliion_align_symmetry` that aligns the symmetry axes of a 3D map in arbitrary orientation according to the conventions in RELION.

3 Results and Discussion

3.1 CPU vector acceleration

The functions that perform the expectation step (particle alignment) in RELION-3 can now use the same code path and structure as the GPU-accelerated code that was introduced in RELION-2, even if executing on systems without GPUs. This has led to several key improvements. First, the algorithm strength reductions first introduced for the GPU code makes it possible to use single precision in most places, which cuts the memory footprint in half for large parts of the code. Second, data has been organised to better fit modern hardware. The legacy CPU code would for instance allocate space for all *possible* orientations under the current sampling level, while the new CPU code makes use of coalesced memory allocation limited to the orientations that will be investigated. These type of changes result in a higher compute intensity with more coalesced memory access in smaller allocations. It is evident that the generalised and more streamlined code path facilitates compiler optimisation compared to the previous CPU version. The standard 3D benchmark e.g. runs 1.5x and 2.5x faster on the previous (Broadwell) and current (Skylake) generations of x86 processors, respectively (Fig. 2A,B), without compromising the results (Fig. 2 - figure supplement 1). However, the new coalesced memory allocation also makes it possible to instruct the compiler to generate efficient code targeting the single-instruction, multiple-data (SIMD) vector units present in modern CPUs. These operate on many input data elements at once, somewhat similar to GPUs, and typically also provide better throughput for single-precision data. Examples of this includes the so-called AVX (advanced vector extensions) instructions present in all 64-bit x86 CPUs that operate on 8 single-precision floating-point numbers, or the AVX512 instructions present in the newest CPUs that extend the width to 16 single-precision numbers. By explicitly instructing the compiler to use AVX-instructions for the new generalised code path, we see a further performance increase of 5.4x on Skylake-generation CPUs (Fig. 2B). This enables x86 CPU nodes to approach the cost-efficiency of professional GPU hardware (Fig. 2C), with results that are qualitatively identical to the legacy code path (Fig. 2 - figure supplement 1).

Notably, use of AVX-instructions does not improve on performance for legacy CPU-execution at all since the code parallelism is not expressed clearly enough for the compiler to recognise it. The performance also depends on which compiler is used, as they are able to optimise the program to varying degrees. The GNU compiler (GCC, version 7.3) seems less able to convert RELIONs current formulation to efficient parallel instructions than the Intel compiler (ICC, version 2018.3.222), and presently the former provides no additional benefit on a Broadwell-generation CPU when the slightly newer AVX2 instructions are enabled. The better performance of ICC is not unreasonable given that these benchmarks target Intel CPUs, although the negligible benefit of enabling AVX2 support in GCC is a bit surprising - hopefully this will improve in future compiler versions.

A handful of things are worth emphasising:

- While the code has been streamlined to allow compilers to generate efficient x86 SIMD vector instructions, only preprocessor directives are used to aid the acceleration, and the new code path is therefore portable to any architecture.
- Since the vector acceleration executes on the CPU, it is straightforward to expand memory and use arbitrarily large box sizes (in contrast to the limitations present when offloading to (GPU) accelerators).
- For now, the legacy CPU execution is preserved for comparison. Legacy execution can be run with any `relion_refine` program by omitting the `--gpu` and/or `--cpu` flag.
- Due to complex compiler dependencies for GPU support, it might not be possible to enable simultaneous CPU and GPU acceleration support. Similarly, since the CPU acceleration targets the same routines as the GPU acceleration, it will not *per se* alleviate any CPU bottlenecks during GPU execution. The generalised path means RELION will branch into functions that launch highly parallel *either* on the CPU or GPU.
- Owing to the higher performance achieved through use of the non-free ICC compiler, and the fact that current GCC versions do not yet automatically generate AVX512 instructions, we will provide pre-build RELION-3 binaries optimised for the latest CPUs in addition to source code distribution.

3.2 Per-particle defocus correction

Per-particle defocus variations within a micrograph arise from different particles being situated at different heights within the field of view. This may happen when the ice layer is relatively thick, and/or when the ice layer is not perpendicular to the electron beam. The latter may occur when the sample support film wrinkles upon flash freezing or when the sample stage is tilted in the microscope, for example to ameliorate a problem of preferred particle orientations [Tan et al., 2017, Naydenova and Russo, 2017]. As an extreme case of non-horizontal samples, we re-processed an influenza virus hemagglutinin dataset (EMPIAR-10097) that was collected at 40 degree stage tilt to overcome a strong preferred orientation of the particles [Tan et al., 2017]. At this tilt angle, the height differences within a single micrograph are more than 4000 Å. A reconstruction from this data set was originally published at 4.2 Å resolution. Recently, the same data were also re-processed in Warp and cryoSPARC, resulting in an improved resolution of 3.2 Å [Tegunov and Cramer, 2018].

We aligned and dose-weighted the 668 movies with our own implementation of the MOTIONCOR2 algorithm, and estimated per-micrograph defocus parameters using CTFFIND4 [Rohou and Grigorieff, 2015]. Because of the high tilt, Thon rings were very weak and the estimated resolution from CTFFIND4 was

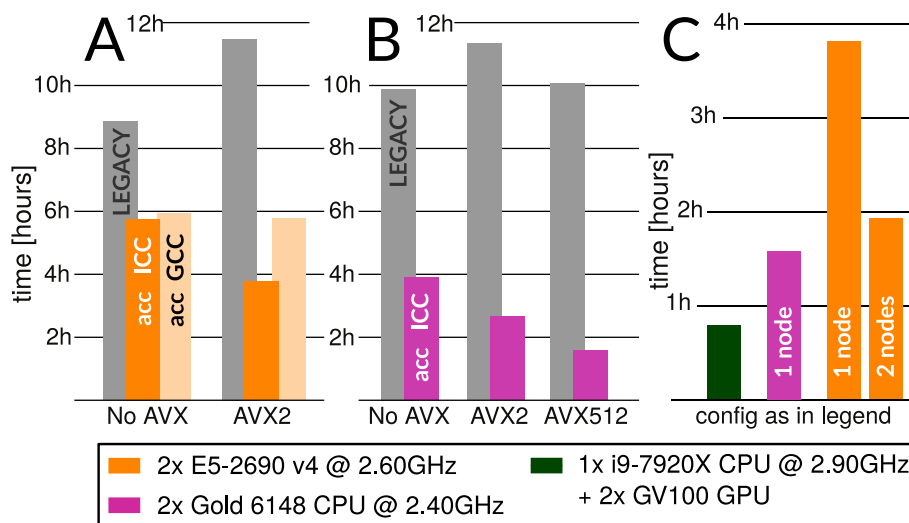


Figure 2: **Accelerated CPU performance** **A** Even when specific vector instructions are disabled, RELION-3 runs faster than RELION-2 even on the previous-generation Broadwell processors that are ubiquitous in many cryo-EM clusters worldwide. Enabling vectorisation during compilation with the Intel compiler benefits the new streamlined code path, but not the legacy code. **B** For latest-generation Skylake CPUs, the difference is much larger even with only AVX2 vectorisation enabled, and when enabling the new AVX512 instructions the performance is roughly 4.5x higher than the legacy code path. **C** The accelerated CPU code executing on dual-socket x86 nodes provides cost-efficiency that is at least approaching that of professional-class GPU hardware (but not consumer GPUs).

limited to about 10 Å. We applied LoG autopicking to an initial subset of 210 micrographs, and subjected the resulting 75,360 particles to 2D classification. We then used four good 2D class averages as templates for reference-based autopicking of all micrographs. The resulting 608,919 particles were then classified into 100 2D classes, from which we selected 361,507 particles that contributed to reasonable 2D class averages. We generated an initial 3D model using the SGD algorithm with C3 symmetry. Up to this point, we did not perform CTF-amplitude correction at resolutions lower than their first peak. A first 3D refinement with full CTF correction yielded only 6.58 Å resolution, and the FSC curve between the two half-maps showed strong oscillations, indicating inaccuracies in the CTF correction (Figure 4A,B). We then refined the per-particle defoci with a search range of ± 2500 Å, which improved the resolution to 4.35 Å. Repeating per-particle defocus refinement and 3D refinement two more times yielded 3.99 and 3.77 Å. Since a fourth iteration did not improve the resolution any further, we then proceeded with Bayesian polishing, which improved the resolution to 3.54 Å. We then used 3D classification into six classes to select

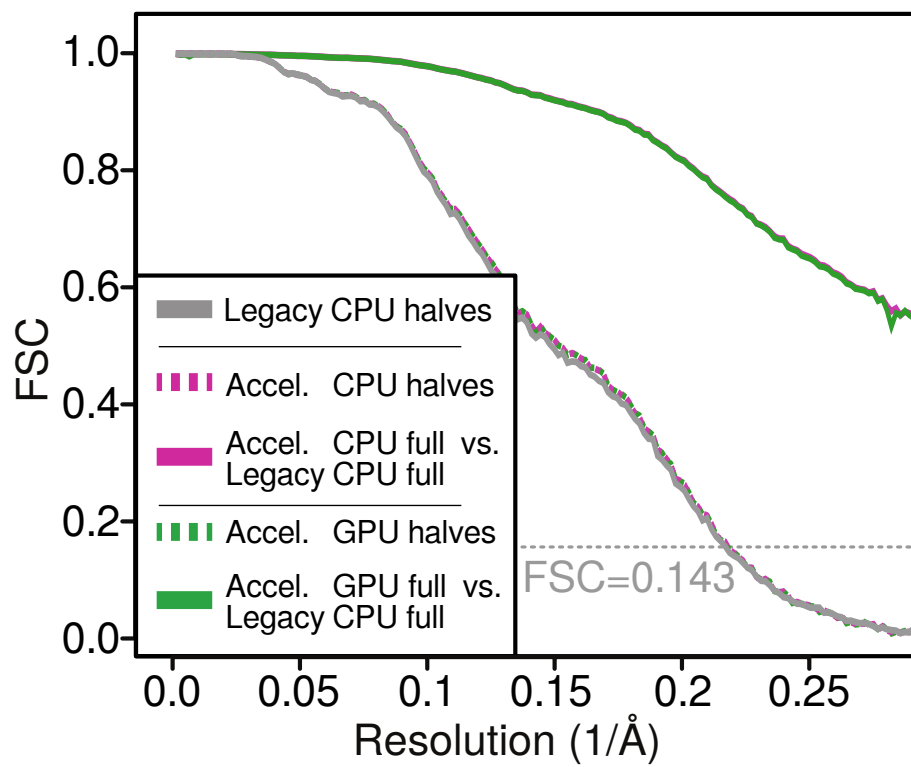


Figure 2 - figure supplement 1 FSCs comparing half-sets in each type of run, legacy-CPU, acc-CPU, and acc-GPU, all reach the same resolution and sampling accuracy. Additional FSCs compare the reconstruction of the acc-CPU and acc-GPU using the legacy-CPU as a base of comparison, showing numerical agreement beyond the reconstructed signal threshold. This validates that the quality of results has not been compromised from using lower precision or compiler optimisation.

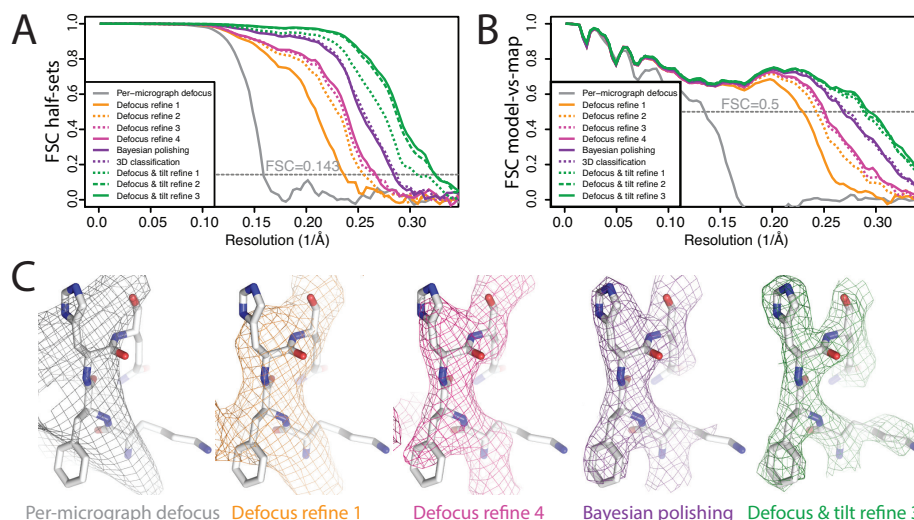


Figure 4: Per-particle defocus correction **A** FSC curves between independently refined half-maps for the different stages of processing as explained in the main text. **B** As in A, but FSC curves are between the cryo-EM maps and the corresponding atomic model (PDB-4FNK) [Ekiert et al., 2012]. **C** Representative density features for some of the maps for which FSC curves are shown in A and B.

a subset of 152,960 particles contributing to the best class, resulting in 3.49 Å resolution after 3D refinement. CTF refinement with beam-tilt correction and 3D refinement were then again iterated three times, giving 3.29 Å, 3.18 Å and finally 3.11 Å resolution. During these refinements, 1,364 particles were translated within 30 Å of their closest neighbour. Since duplicated particles may invalidate gold-standard separation and thereby inflate resolution, we removed these duplicate particles before the final 3D refinement using new functionality in the "Subset selection" jobtype on the GUI. Plots of the final defocus values for particles within a single micrograph clearly replicate the tilt geometry of the experiment (Figure 4 - figure supplement 1).

This example demonstrates the robustness of our per-particle defocus refinement against the initial values being far off. Using initial defoci that are closer to the correct ones will reduce the need for performing multiple cycles of per-particle defocus refinement. Therefore, in practice it would probably be faster to start from defoci provided by local CTF estimation programs, such as CTFTILT [Mindell and Grigorieff, 2003], GCTF [Zhang, 2016] or Warp [Tegunov and Cramer, 2018]. Currently, this requires some scripting by the user, as the wrappers to CTFFIND4 and Gctf on the RELION-3 GUI do not allow local CTF estimation.

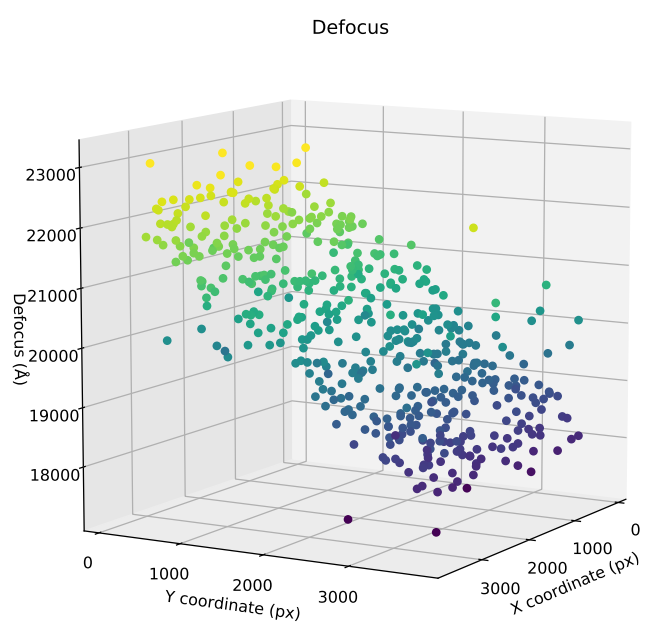


Figure 4 - figure supplement 1 Per-particle defocus estimates (along Z) are plotted against the X,Y-coordinates of the particles in a representative micrograph.

3.3 Beam tilt correction

Until to date, beam tilt has been largely ignored in cryo-EM single-particle analysis. This is partially because of a lack of tools to estimate and correct for beam tilt, and partially because in many cases the amount of beam tilt did not affect the achievable resolution to a noticeable extent. Beam tilt results from sub-optimal coma-free alignment of the microscope. Based on our experience, we estimate that currently available microscope alignment procedures may still leave up to a quarter mrad of beam tilt in the data. This will not limit the resolution of 3 Å structures, but as shown for β -galactosidase in section 3.6, such amounts of beam tilt are limiting for 2 Å structures. Therefore, as resolutions keep improving, beam tilt correction will gain in importance.

There are data acquisition scenarios where beam tilt becomes a limiting factor at much lower resolutions. With the elevated costs of high-end electron microscopes, faster data acquisition methods are highly sought after. A rate limiting step in many data acquisition schemes is to physically move the sample stage from one hole in the carbon to the next, as one needs to wait for mechanical stage drift to settle. Because it is much faster to shift the electron beam instead, one could achieve higher data acquisition rates by using beam-shift to image multiple holes before shifting the stage. Beam-shift is already used extensively to take multiple images within a single hole, but because it introduces beam tilt, its use for imaging multiple holes has only recently attracted attention [Cheng et al., 2018]. Although it is possible to actively compensate for the introduced beam tilt using the microscope's deflection coils [Eades, 2006, Glaeser et al., 2011], this procedure puts more stringent requirements on microscope alignment.

We evaluated our methods for beam tilt estimation and correction using two data sets on horse spleen apo-ferritin (purchased from Sigma), which were acquired with or without active beam tilt compensation. The data were recorded using serialEM [Mastronarde, 2005] in combination with a patch in the Thermo Fischer microscope software to perform active beam tilt compensation. (At the time of our experiment, active beam-tilt compensation was not yet implemented in serialEM.) For both data sets, we used the same R2/2 Quantifoil grid; the same Titan Krios microscope with a Gatan energy filter (with a slit width of 20 eV); the same K2-summit detector; and we used beam-shift to take five images in nine holes each before shifting the stage. This procedure was repeated nine times for the first data set and ten times for the second, because the area imaged in the latter contained fewer particles. We only used active beam tilt compensation for the first data set. Movies with a total dose of 41 electrons per Å² were fractionated into 40 frames with a pixel size of 1.04 Å. Template-based auto-picking yielded 104,992 particles for the first and 109,646 particles for the second data set. After selection of suitable 2D classes, the first and second data sets comprised 79,985 and 65,131 particles, respectively.

The first data set, with active beam tilt compensation, yielded a resolution of 2.27 Å after standard 3D auto-refinement, and 2.20 Å after Bayesian polishing and a second round of refinement (Figure 6A,B). The second data set, without active beam tilt compensation, yielded 2.92 Å after the first refinement,

and 2.83 Å after Bayesian polishing. Beam tilt estimation was then performed independently for nine subsets of the data: one subset for each of the nine holes that were imaged before shifting the stage, and each subset thus comprising 10×5 movies. Plots of the phase differences measured for each of the nine holes clearly show the effects of beam tilt at higher resolution (Figure 6 - Figure supplement 1). Smaller phase shifts in the opposite direction at intermediate resolution indicate that during refinement, the images have shifted off-center in an attempt to compensate for the beam tilt-induced phase shifts. A third round of refinement with correction for the estimated beam tilts improved the resolution to 2.18 Å. Plots of the estimated beam tilt against the beam-shift recorded by serialEM (Figure 6 - Figure supplement 2) revealed a linear relationship of 0.19 mrad of beam tilt per μm of beam-shift in both the X and Y directions. We then also ran a fourth refinement, where we predicted the beam tilt for each movie separately based on the recorded beam-shift values by serialEM. This resulted in a resolution of 2.23 Å. Features in the reconstructed density confirmed the improvement in resolution caused by beam tilt correction (Figure 6C). These results illustrate that beam tilt estimation and correction in RELION-3 may be a feasible alternative to active beam tilt compensation at the microscope. Future studies will need to explore which of these methods is preferable under different experimental conditions.

3.4 Ewald sphere correction

Ewald sphere correction has also been largely ignored, although a few contributions have recently described high-resolution reconstructions where Ewald sphere correction did make a difference [Tan et al., 2018, Zhu et al., 2018]. To test our Ewald-sphere correction implementation, we used particle images of the P22 virion capsid protein [Hryc et al., 2017], which we downloaded from EMPIAR-10083 [Iudin et al., 2016]. To reduce memory consumption and to accelerate the processing, we down-sampled the particles to 1.54 Å/pixel, and skipped Fourier padding during refinement. As we did not have access to the original micrographs, we used the defocus parameters for 45,150 particles as they were refined by JSPR in the original publication, which were kindly provided by Wah Chiu. We ignored variations in magnification as estimated by JSPR, because magnification correction is not implemented in RELION-3. Standard 3D auto-refinement, with correction for an estimated beam tilt of 0.6 mrad in X and -0.13 mrad in Y, gave a reconstruction with an estimated resolution of 3.54 Å (Figure 9A,B). Subsequent Ewald sphere correction improved this to 3.32 Å. Features in the reconstructed density confirmed the improvement in resolution caused by beam tilt correction (Figure 9C). Ewald sphere correction with the opposite curvature gave only 3.92 Å. In theory, if one knows the number of mirroring operations during data acquisition and image processing, Ewald sphere correction could therefore also be used to determine the absolute hand of a structure. However, at the resolutions where Ewald sphere correction matters, this is often no longer a relevant question.

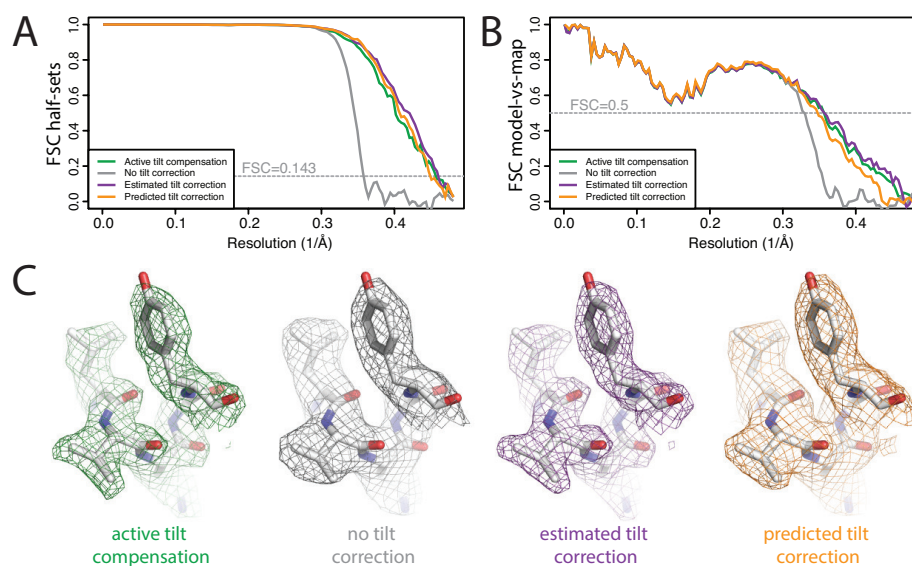


Figure 6: **Beam tilt correction** **A** FSC curves between independently refined half-maps for the different stages of processing as explained in the main text. **B** As in A, but FSC curves are between the cryo-EM maps and the corresponding atomic model (PDB-2W00) [de Val et al., 2012]. **C** Representative density features from the 2.20 Å map from the data set with active beam tilt compensation (left); the 2.92 Å map from the data set without active beam tilt compensation and without beam tilt correction (middle); and the 2.18 Å map from the data set without active beam tilt compensation but with beam tilt correction (right).

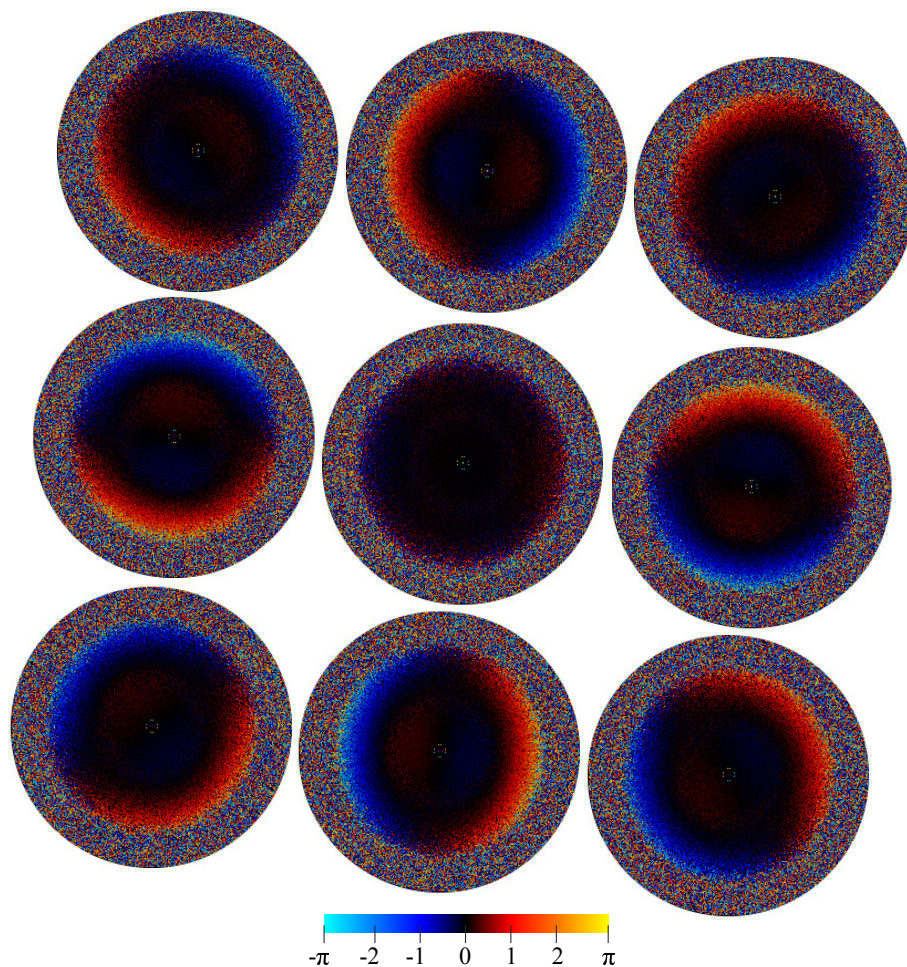


Figure 6 - figure supplement 1 The average per-pixel phase differences (in radians) between reference projections and the individual experimental particle images (i.e. the phase angle of q_k in Eq. 8). The nine plots correspond to subsets of images collected at the nine different holes without active beam-tilt compensation. Note that the applied beam shift values have been recorded in stage coordinates giving a 175 degree X-axis rotation for the magnification used, which is reflected in the relative position of the nine plots. The inside regions showing the opposite slope to the global trend result from a systematic error in particle alignment caused by beam tilt.

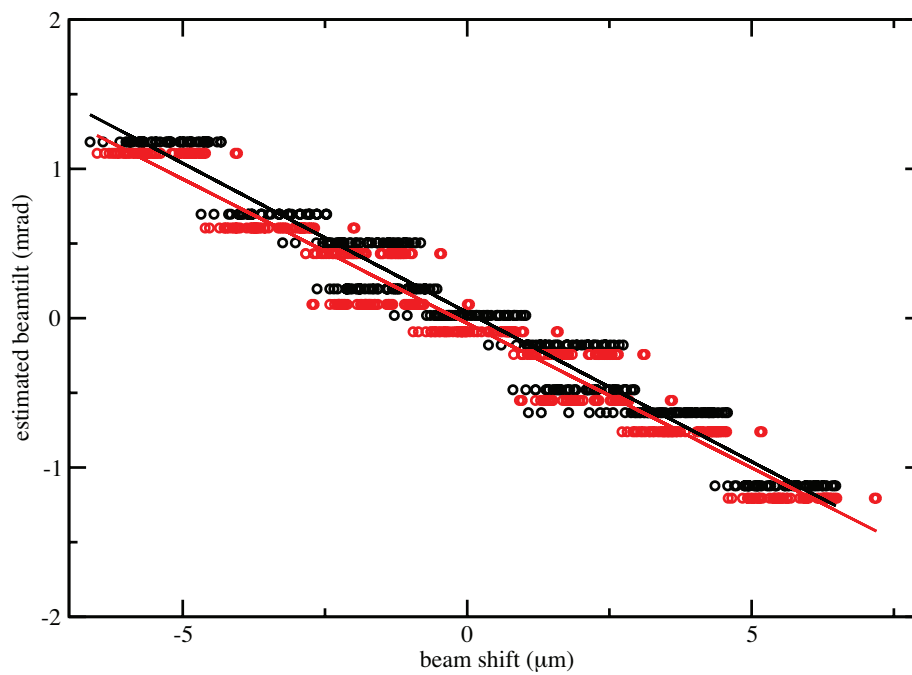


Figure 6 - figure supplement 2 The estimated beam tilt is plotted against the beam-shift values from serialEM in the X-direction (black circles) and in the Y-direction (red circles). Linear fits through these points indicate there is 0.19 mrad of beam tilt for each μm of beam shift.

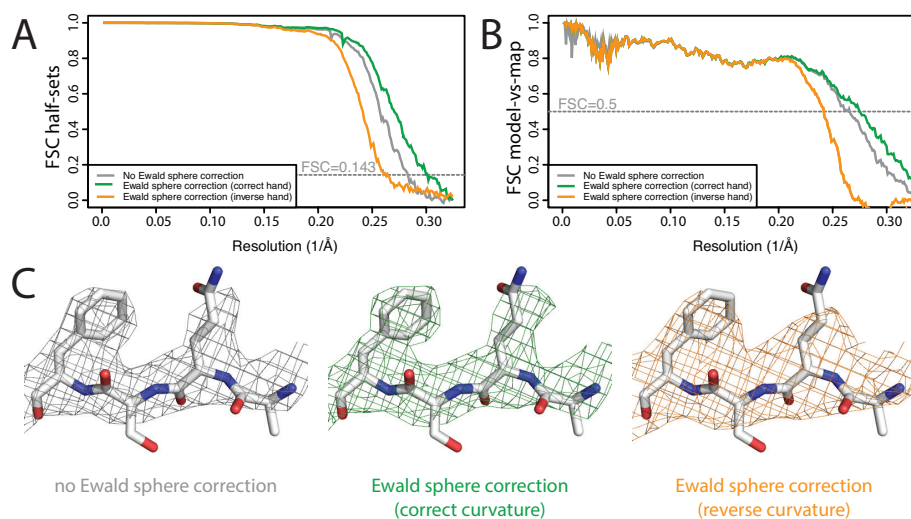


Figure 9: **Ewald sphere correction** **A** FSC curves between independently refined half-maps without Ewald sphere correction (grey); with Ewald sphere correction with the correct curvature (green) and with Ewald sphere correction with the inverse curvature (orange). **B** As in A, but FSC curves are between the cryo-EM and the corresponding atomic model (PDB-5UU5) [Hryc et al., 2017]. **C** Representative density features without Ewald sphere correction (grey) with Ewald sphere correction and the correct curvature (green) and with Ewald sphere correction and the reverse curvature (orange).

3.5 Non-interactive data exploration

To illustrate the general applicability of the `relion_it.py` script, we re-processed our previously published γ -secretase data set [Bai et al., 2015b,a], which we deposited at EMPIAR-10194. We chose this data set as a representative of a non-straightforward data set, as γ -secretase is a small (130 kDa ordered mass), asymmetrical, membrane-embedded protein complex.

Using default parameters, we performed motion correction of the 2,925 16-frame movies in our own implementation of the MotionCor2 algorithm, and estimated CTF parameters in gCTF [Zhang, 2016]. LoG-based picking (with $d_{min} = 120$ and $d_{max} = 180$) at the default threshold led to 2.6 million particles (Figure 10A), which were divided in batches of 10,000 particles. Only the first batch was used for 2D classification, which revealed the presence of different views of the complex (Figure 10B). The script could have run similar 2D classification on all batches, for example to monitor data quality during data acquisition.

Instead, we chose to let the script proceed with SGD initial model generation and 3D classification on the same batch of 10,000 particles (Figure 10C). After 3D classification, the script automatically selects the 3D class with the highest resolution reconstruction. As all four classes had an equal resolution of 23 Å, the script selected the largest class (with an accumulated weight of 34.5%). However, the second largest class (with an accumulated weight of 34.4%) was actually the correct structure, which was obvious from a visual inspection of the density (Figure 10C). At this stage we intervened. We stopped the script and edited the file `RELION_IT_2NDPASS_3DREF` to point to the correct map. This intervention probably reflects the weakest point in our otherwise non-interactive procedure. Robust, non-interactive selection of 3D class reconstructions (or 2D class averages) is a subject of ongoing research in our groups.

Upon re-starting the script, template-based auto-picking with the correct 3D structure and the default threshold (0.4) led to 1.8 million particles, which were then used for 2D classification into 200 classes. At this point, the non-interactive script finished and the rest of the data processing was performed manually. Selection of the best 2D class averages led to a subset comprising 665,102 particles, which were subjected to Bayesian particle polishing and two rounds of 3D classification into eight classes. Two classes, comprising 333,377 particles, were selected for refinement of the per-particle defocus values and separate estimations of the beam tilt in the six microscopy sessions that this data set comprises. A second 3D classification yielded a majority class with 240,777 particles, which led to a final resolution of 3.29 Å (Figure 10E-G).

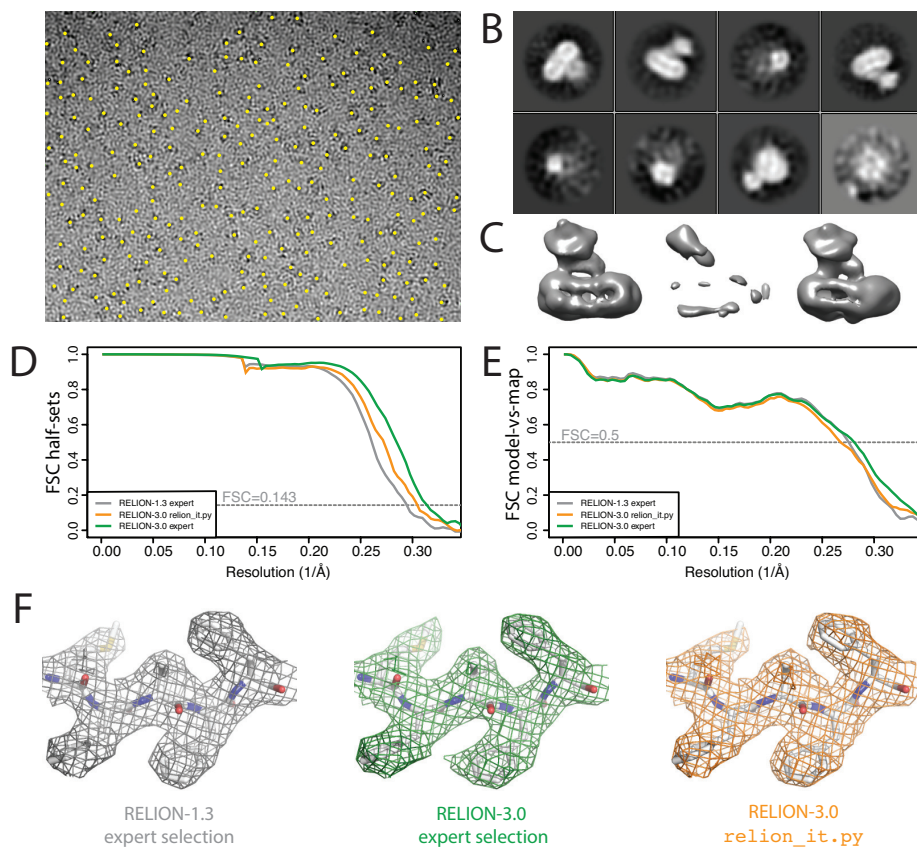
When we originally processed this data set in RELION-1.3, we obtained a map at 3.40 Å resolution from 159,549 particles, which were selected after 3D classification of 412,272 polished particles [Bai et al., 2015b]. This set of particles was obtained through careful tuning of template-based auto-picking parameters for each of the six microscopy sessions, and careful selection of 2D and 3D classes. To assess whether the non-interactive pre-processing procedure can compete with this expert-driven approach, we used the same set of 412,272

particles also for Bayesian polishing, refinement of per-particle defocus values, beam tilt estimation and 3D classification. The resulting selection of 246,602 particles gave a final resolution of 3.18 Å (Figure 10E-G). The improvement from 3.40 Å to 3.18 Å represents the contribution of the new functionality in RELION-3, whereas the gain from 3.29 Å to 3.18 Å represents the added value of the expert user. However, the map from the automatically pre-processed data does not correlate with the atomic model as well as one would expect from the half-map FSC curves. This suggests a higher fraction of false positives in the data resulting from the `relion_it.py` script compared to the expert selection, which may be caused by the relatively low default for the auto-picking threshold in the script. These results illustrate that non-interactive pre-processing gets close to the performance of a human expert, but for challenging data sets there is still a benefit in careful tuning of parameters. Future investigations will aim to reduce this gap.

3.6 High-resolution refinement I: β -galactosidase

To illustrate the combined effects of the RELION-3 implementation on high-resolution structure determination, we re-processed the β -galactosidase data set at EMPIAR-10061 [Bartesaghi et al., 2015] that we also used to illustrate RELION-2.0 [Kimanius et al., 2016]. This data set was made available to the community by the Subramaniam group. Using RELION-2.0, we previously obtained a 2.2 Å resolution reconstruction (EMD-4116).

We aligned the original, super-resolution movies using the UCSF implementation of MOTIONCOR2 (version 1.0.5) [Zheng et al., 2017] in 4 by 4 patches with dose weighting and 2-fold binning. CTF parameters were estimated using CTFFIND 4.1.10 [Rohou and Grigorieff, 2015]. We initially picked 4,084 particles using the LoG-based algorithm from a subset of 462 movies and classified these particles into 20 2D classes. Six of the resulting 2D class averages were used for template-based auto-picking of 229,553 particles from the entire data set. After an initial 2D classification into 200 classes, we chose to reject only the most obvious outliers, resulting in a data set of 217,752 particles. Although many of the selected particles were probably sub-optimal, we chose to keep them until the Bayesian polishing step, where relatively high particle numbers per micrograph may improve the description of the motion tracks. An initial 3D auto-refinement with these particles yielded a map with an estimated resolution of 2.31 Å (Figure 11A,B). Re-refining the same data set after only per-particle defocus refinement or beam tilt estimation improved the resolution to 2.21 or 2.22 Å, respectively. Application of both gave 2.08 Å resolution. We then performed Bayesian polishing, which gave 1.94 Å, and classified the resulting shiny particles into six classes. The 134,900 particles that contributed to the three best 3D classes were then re-refined to 1.92 Å resolution. A second round of per-particle defocus refinement and beam tilt estimation, followed by another 3D auto-refinement yielded a final resolution of 1.90 Å. Ewald sphere correction for this data set did not improve the resolution by at least one Fourier shell, although a small increase in FSC values near the estimated resolution was visible.



The reconstructed density showed features that one would expect at this resolution (Figure 11C). FSC curves between the maps and PDB-5A1A confirmed that RELION-3 improved resolution by 0.3 Å compared to RELION-2.0 (Figure 11 - figure supplement 1). Application of the `bfactor_plot.py` script produced an estimate for the overall B-factor of 56 Å² (Figure 11 - figure supplement 2). By extrapolation, the script predicts a resolution of 1.75 Å when using four times more particles, and 1.69 Å when using eight times more particles.

Recently, the Subramaniam group also re-processed this data set using new methods for defocus estimation, motion correction and radiation damage weighting [Bartesaghi et al., 2018]. The resulting map (EMD-7770) was reported as an atomic resolution structure. The resolution according to the 0.143 FSC criterion was 1.9 Å. Despite the same nominal resolution of 1.9 Å, the RELION-3 map correlated with PDB-5A1A to 0.15 Å higher resolution than EMD-7770 (Figure 11 - figure supplement 2). In Bartesaghi et al. [2018] it is argued that the FSC curves may under-estimate the true resolution of cryo-EM maps, and a point is made for showing figures of high-resolution density features to claim atomic resolution instead. We disagree. The high-resolution figures in that paper were made using a map that was sharpened to resolutions well beyond the true resolution, resulting in high levels of noise. This noise was subsequently removed from regions outside the protein density with a high-resolution mask. As this noise still exists within the protein region, the resulting map is prone to incorrect interpretations. FSC curves between half-maps and between model and map are not without their own problems. They do not account for modelling errors or local variations in resolution, and they can be inflated by overfitting or the use of suboptimal masks. Nevertheless, provided suitably soft and low-resolution masks are used and care is taken to avoid overfitting [Scheres and Chen, 2012, Brown et al., 2015], FSC curves provide an objective overall assessment of map quality. Discarding this information in favour of a few selected density figures would be a mistake.

3.7 High-resolution refinement II: human apo-ferritin

For a final illustration of the combined functionality in RELION-3, we processed a data set on a human apo-ferritin sample, which was a gift from Fei Sun. These data were collected on the same K2 detector and Titan Krios microscope as the horse spleen apo-ferritin data set in section 3.3. We used serialEM to record seven images per hole in the R2/2 Quantifoil grid, before moving the stage to the next hole. Movies with a total dose of 47 electrons per Å² were fractionated into 40 frames with a pixel size of 0.814 Å. We have deposited these data at EMPIAR-10200.

We again used the `relion_it.py` script for non-interactive pre-processing of the data. Using default parameters, the script launched motion correction of the 1,255 40-frame movies in our own implementation of the MOTIONCOR2 algorithm, and estimated CTF parameters in gCTF [Zhang, 2016]. LoG-based picking (with $d_{min} = 110$ and $d_{max} = 150$) at the default threshold led to 723,524 particles, which were divided in batches of 10,000 particles. Only the

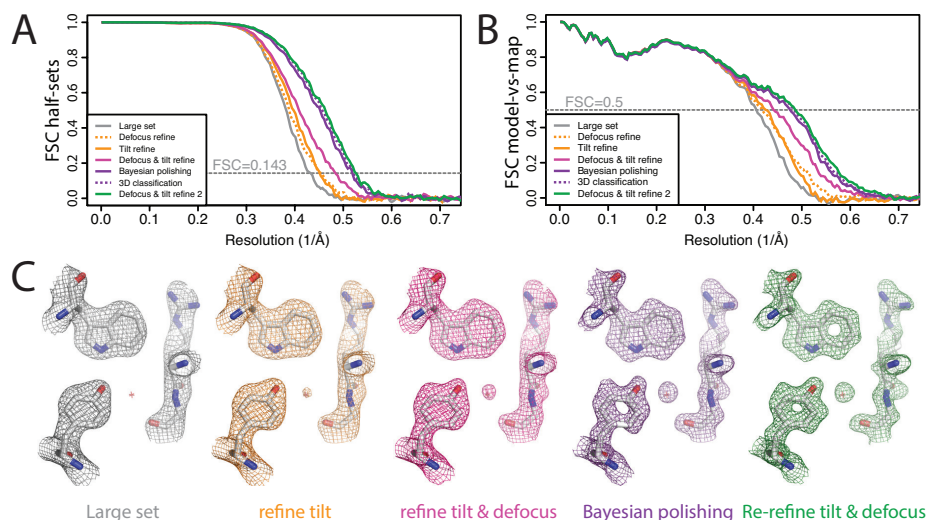


Figure 11: **High-resolution refinement: β -galactosidase** **A** FSC curves between independently refined half-maps for the different stages of processing as explained in the main text. **B** FSC curves between PDB-5A1A and the maps at the same stages of processing as in A. **C** Representative density features for some of the maps for which FSC curves are shown in A and B.

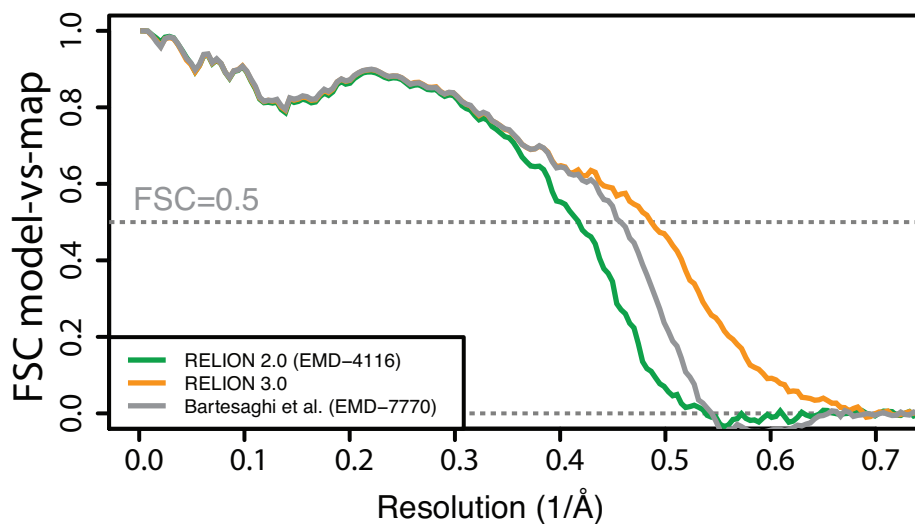


Figure 11 - figure supplement 1 FSC curves between PDB-5A1A and EMD-4116, EMD-7770 and the new map reconstructed in RELION-3.

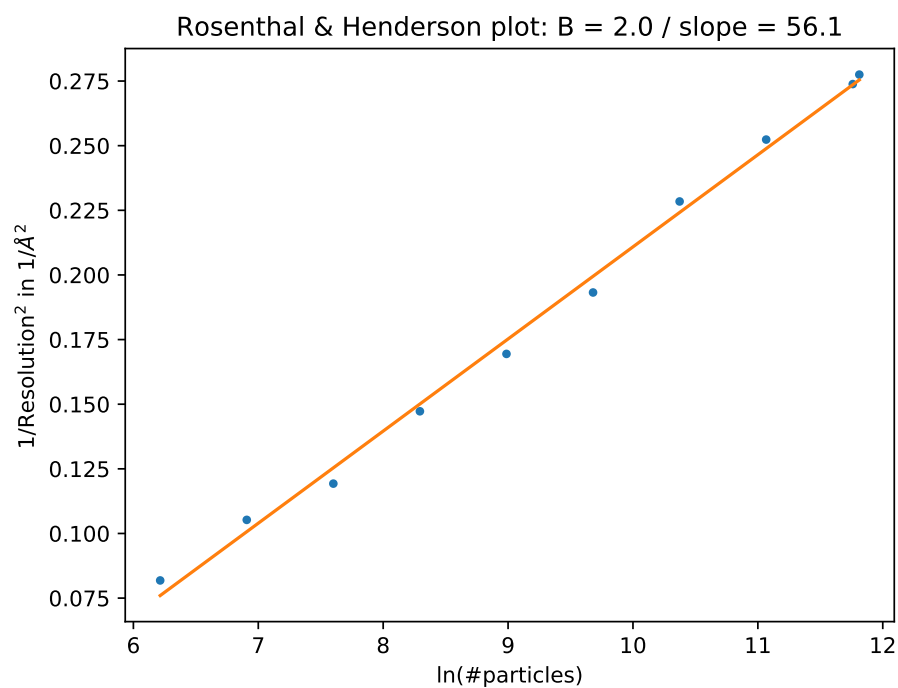


Figure 11 - figure supplement 2 The B-factor plot that was generated automatically by the `bfactor_plot.py` script. The script performed nine 3D refinement and post-processing jobs with subsets of increasing sizes for the beta-galactosidase data set. Fitting a straight line through the inverse of the resolution squared of each refinement against the natural logarithm of the number of particles in the corresponding subset yielded an estimated B-factor of 56 \AA^2 .

first batch was used for initial model generation and 3D classification. Automated selection of the highest resolution 3D class (at 6.5 Å) yielded a good apo-ferritin model, which was used for template-based auto-picking using the default (0.4) threshold. The resulting 427,694 particles were used for 2D classification into 200 classes.

We then proceeded processing in an interactive manner. Selection of suitable 2D class averages led to a subset of 426,450 particles, which were subjected to Bayesian polishing and per-particle defocus refinement. A final 3D refinement yielded a resolution of 1.65 Å, which is two resolution shells below the Nyquist frequency (Figure 14A,B). Estimation of the beam tilt yielded values below 0.1 mrad, and refinement with these values did not improve the map. As expected at this resolution, the map shows spherical-like features around some of the atoms (Figure 14C). However, as these are not yet separated, it would be a stretch to call this an atomic resolution structure. Application of the `bfactor_plot.py` script produced an estimate for the overall B-factor of 66 Å² (Figure 14 - figure supplement 1).

Besides the selection of suitable 2D and 3D classes, which as mentioned is a topic of ongoing research in our groups, all of the calculations performed to yield this map could be incorporated into an extended version of the `reliion_it.py` script. Thereby, completely automated, high-resolution cryo-EM structure determination of relatively straightforward samples like apo-ferritin is within close reach. This will be useful for the standardised processing of data sets on test samples to assess microscope performance. In the future, improved versions of procedures like this may facilitate high-throughput cryo-EM structure determination, for example for the screening of small-molecule compounds bound to drug targets.

4 Conclusions

The new functionality in RELION-3 allows faster, more automated and higher resolution cryo-EM structure determination compared to previous versions. CPU acceleration makes it possible to use a much wider range of hardware efficiently for RELION, including laptops and desktops, and it avoids the box size limits still present when offloading to GPU accelerators. The functionality to schedule and execute jobs from the command-line, coupled with python scripting, provides a versatile and extendable platform to standardise and automate complicated tasks. Corrections for per-particle defocus variations, beam tilt and the Ewald sphere already improve resolutions for currently available data sets, and will become even more important as the resolutions of typical cryo-EM structures improve in the future. RELION is distributed under the GNU General Public License (GPL, version 2), meaning it is completely free to use or modify by anyone, which we hope will maximise its usefulness to the community.

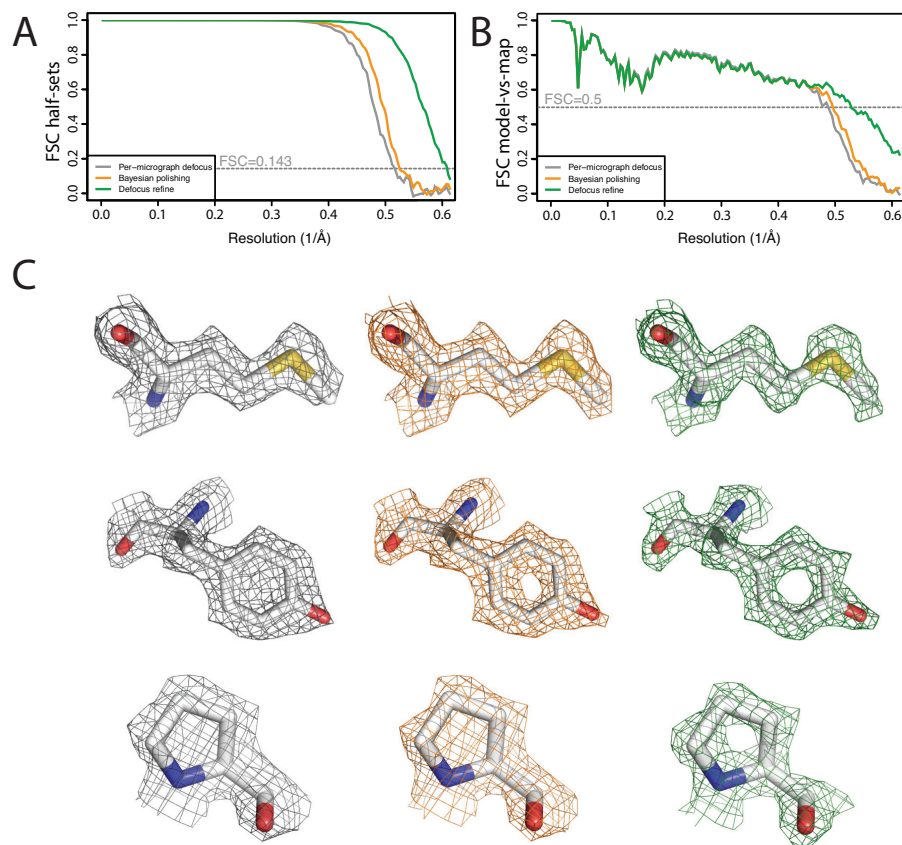


Figure 14: **High-resolution refinement: apo-ferritin** **A** FSC curves between independently refined half-maps at different stages of the processing. **B** FSC curves between PDB-5N27 [Ferraro et al., 2017] and the cryo-EM map at different stages of the processing. **C** Representative density features at different stages of the processing.

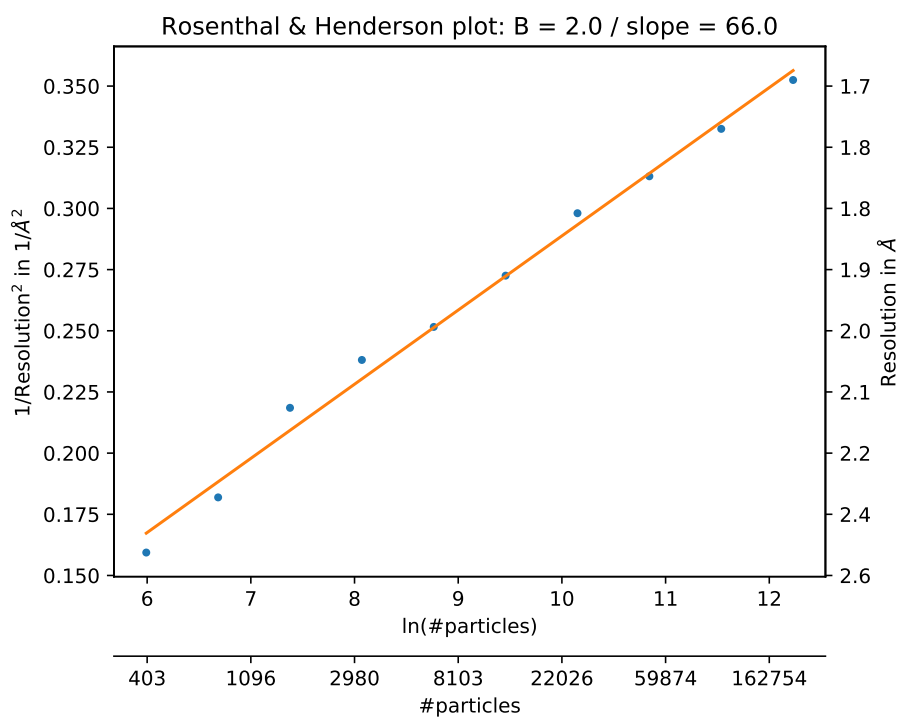


Figure 14 - figure supplement 1 The B-factor plot that was generated automatically by the `bfactor_plot.py` script. The script performed ten 3D refinement and post-processing jobs with subsets of increasing sizes for the high-resolution apo-ferritin data set. Fitting a straight line through the inverse of the resolution squared of each refinement against the natural logarithm of the number of particles in the corresponding subset yielded an estimated B-factor of 66 Å².

Acknowledgements

We thank Charles W. Congdon from Intel for extensive assistance with implementing the accelerated CPU code; Chris Russo and Richard Henderson for helpful discussions; and Jake Grimmett, Toby Darling and Stefan Fleischmann for help with high-performance computing. This work was funded by the UK Medical Research Council (MC_UP_A025_1013 to SHWS); the Swiss National Science Foundation (SNF: P2BSP2 168735 to JZ); the Swedish Research Council (2017-04641 to EL) and the Knut and Alice Wallenberg Foundation (to EL).

References

- Vahid Abrishami, Javier Vargas, Xueming Li, Yifan Cheng, Roberto Marabini, Carlos Óscar Sánchez Sorzano, and José María Carazo. Alignment of direct detection device micrographs using a robust Optical Flow approach. *Journal of Structural Biology*, 189(3):163–176, March 2015. ISSN 1095-8657. doi: 10.1016/j.jsb.2015.02.001.
- Xiao-chen Bai, Eeson Rajendra, Guanghui Yang, Yigong Shi, and Sjors HW Scheres. Sampling the conformational space of the catalytic subunit of human gamma-secretase. *eLife*, 4:e11182, December 2015a. ISSN 2050-084X. doi: 10.7554/eLife.11182. URL <https://elifesciences.org/articles/11182>.
- Xiao-Chen Bai, Chuangye Yan, Guanghui Yang, Peilong Lu, Dan Ma, Linfeng Sun, Rui Zhou, Sjors H. W. Scheres, and Yigong Shi. An atomic structure of human gamma-secretase. *Nature*, 525(7568):212–217, September 2015b. ISSN 1476-4687. doi: 10.1038/nature14892.
- Alberto Bartesaghi, Alan Merk, Soojay Banerjee, Doreen Matthies, Xiongwu Wu, Jacqueline L. S. Milne, and Sriram Subramaniam. 2.2 Å resolution cryo-EM structure of beta-galactosidase in complex with a cell-permeant inhibitor. *Science (New York, N.Y.)*, 348(6239):1147–1151, June 2015. ISSN 1095-9203. doi: 10.1126/science.aab1576.
- Alberto Bartesaghi, Cecilia Aguerrebere, Veronica Falconieri, Soojay Banerjee, Lesley A. Earl, Xing Zhu, Nikolaus Grigorieff, Jacqueline L. S. Milne, Guillermo Sapiro, Xiongwu Wu, and Sriram Subramaniam. Atomic Resolution Cryo-EM Structure of beta-Galactosidase. *Structure (London, England: 1993)*, 26(6):848–856.e3, June 2018. ISSN 1878-4186. doi: 10.1016/j.str.2018.04.004.
- Tanmay A. M. Bharat, Christopher J. Russo, Jan Löwe, Lori A. Passmore, and Sjors H. W. Scheres. Advances in Single-Particle Electron Cryomicroscopy Structure Determination applied to Sub-tomogram Averaging. *Structure (London, England: 1993)*, 23(9):1743–1753, September 2015. ISSN 1878-4186. doi: 10.1016/j.str.2015.06.026.

- Nikhil Biyani, Ricardo D. Righetto, Robert McLeod, Daniel Caujolle-Bert, Daniel Castano-Diez, Kenneth N. Goldie, and Henning Stahlberg. Focus: The interface between data collection and data processing in cryo-EM. *Journal of Structural Biology*, 198(2):124–133, 2017. ISSN 1095-8657. doi: 10.1016/j.jsb.2017.03.007.
- Alan Brown, Fei Long, Robert A. Nicholls, Jaan Toots, Paul Emsley, and Garib Murshudov. Tools for macromolecular model building and refinement into electron cryo-microscopy reconstructions. *Acta Crystallographica. Section D, Biological Crystallography*, 71(Pt 1):136–153, January 2015. ISSN 1399-0047. doi: 10.1107/S1399004714021683.
- Anchi Cheng, Edward T. Eng, Lambertus Alink, William J. Rice, Kelsey D. Jordan, Laura Y. Kim, Clinton S. Potter, and Bridget Carragher. High Resolution Single Particle Cryo-Electron Microscopy using Beam-Image Shift. *bioRxiv*, page 306241, April 2018. doi: 10.1101/306241. URL <https://www.biorxiv.org/content/early/2018/04/26/306241>.
- J. M. de la Rosa-Trevín, A. Quintana, L. Del Cano, A. Zaldívar, I. Foche, J. Gutiérrez, J. Gómez-Blanco, J. Burguet-Castell, J. Cuenca-Alba, V. Abrishami, J. Vargas, J. Otón, G. Sharov, J. L. Vilas, J. Navas, P. Conesa, M. Kazemi, R. Marabini, C. O. S. Sorzano, and J. M. Carazo. Scipion: A software framework toward integration, reproducibility and validation in 3d electron microscopy. *Journal of Structural Biology*, 195(1):93–99, July 2016. ISSN 1095-8657. doi: 10.1016/j.jsb.2016.04.010.
- Natalia de Val, Jean-Paul Declercq, Chang Kee Lim, and Robert R. Crichton. Structural analysis of haemin demetallation by L-chain apoferritins. *Journal of Inorganic Biochemistry*, 112:77–84, July 2012. ISSN 1873-3344. doi: 10.1016/j.jinorgbio.2012.02.031.
- Alwyn Eades. Obtaining TEM images with a uniform deviation parameter. *Ultramicroscopy*, 106(4-5):432–438, March 2006. ISSN 0304-3991. doi: 10.1016/j.ultramic.2005.12.002.
- Damian C. Ekiert, Arun K. Kashyap, John Steel, Adam Rubrum, Gira Bhabha, Reza Khayat, Jeong Hyun Lee, Michael A. Dillon, Ryann E. O’Neil, Aleksandr M. Faynboym, Michael Horowitz, Lawrence Horowitz, Andrew B. Ward, Peter Palese, Richard Webby, Richard A. Lerner, Ramesh R. Bhatt, and Ian A. Wilson. Cross-neutralization of influenza A viruses mediated by a single antibody loop. *Nature*, 489(7417):526–532, September 2012. ISSN 1476-4687. doi: 10.1038/nature11414.
- Rafael Fernandez-Leiro and Sjors H. W. Scheres. Unravelling biological macromolecules with cryo-electron microscopy. *Nature*, 537(7620):339–346, 2016. ISSN 1476-4687. doi: 10.1038/nature19948.
- Giarita Ferraro, Silvia Ciambellotti, Luigi Messori, and Antonello Merlino. Cisplatin Binding Sites in Human H-Chain Ferritin. *Inorganic Chemistry*, 56

- (15):9064–9070, August 2017. ISSN 0020-1669. doi: 10.1021/acs.inorgchem.7b01072. URL <https://doi.org/10.1021/acs.inorgchem.7b01072>.
- Robert M. Glaeser, Dieter Typke, Peter C. Tiemeijer, James Pulokas, and Anchi Cheng. Precise beam-tilt alignment and collimation are required to minimize the phase error associated with coma in high-resolution cryo-EM. *Journal of Structural Biology*, 174(1):1–10, April 2011. ISSN 1095-8657. doi: 10.1016/j.jsb.2010.12.005.
- Timothy Grant, Alexis Rohou, and Nikolaus Grigorieff. cis TEM, User-friendly software for single-particle image processing. *eLife*, 7, March 2018a. ISSN 2050-084X. doi: 10.7554/eLife.35383.
- Timothy Grant, Alexis Rohou, and Nikolaus Grigorieff. cisTEM, user-friendly software for single-particle image processing. *eLife*, 7, March 2018b. ISSN 2050-084X. doi: 10.7554/eLife.35383.
- Nikolaus Grigorieff. FREALIGN: high-resolution refinement of single particle structures. *Journal of Structural Biology*, 157(1):117–125, January 2007. ISSN 1047-8477. doi: 10.1016/j.jsb.2006.05.004. URL <http://www.ncbi.nlm.nih.gov/pubmed/16828314>.
- Sydney R. Hall. The STAR file: a new format for electronic data transfer and archiving. *J. Chem. Inf. Comput. Sci.*, 31(2):326–333, 1991. ISSN 0095-2338. doi: 10.1021/ci00002a020. URL <http://dx.doi.org/10.1021/ci00002a020>.
- Shaoda He and Sjors H. W. Scheres. Helical reconstruction in RELION. *Journal of Structural Biology*, in press, 2017. ISSN 1047-8477. doi: 10.1016/j.jsb.2017.02.003. URL <http://www.sciencedirect.com/science/article/pii/S1047847717300199>.
- Corey F. Hryc, Dong-Hua Chen, Pavel V. Afonine, Joanita Jakana, Zhao Wang, Cameron Haase-Pettingell, Wen Jiang, Paul D. Adams, Jonathan A. King, Michael F. Schmid, and Wah Chiu. Accurate model annotation of a near-atomic resolution cryo-EM map. *Proceedings of the National Academy of Sciences*, 114(12):3103–3108, March 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1621152114. URL <http://www.pnas.org/content/114/12/3103>.
- Andrii Iudin, Paul K. Korir, José Salavert-Torres, Gerard J. Kleywegt, and Ardan Patwardhan. EMPIAR: a public archive for raw electron microscopy image data. *Nature Methods*, 13(5):387–388, May 2016. ISSN 1548-7105. doi: 10.1038/nmeth.3806.
- Wen Jiang, Matthew L. Baker, Joanita Jakana, Peter R. Weigele, Jonathan King, and Wah Chiu. Backbone structure of the infectious epsilon15 virus capsid revealed by electron cryomicroscopy. *Nature*, 451(7182):1130–1134, February 2008. ISSN 1476-4687. doi: 10.1038/nature06665.

- Dari Kimanius, Björn O Forsberg, Sjors HW Scheres, and Erik Lindahl. Accelerated cryo-EM structure determination with parallelisation using GPUs in RELION-2. *eLife*, 5, November 2016. ISSN 2050-084X. doi: 10.7554/eLife.18722. URL <http://elifesciences.org/lookup/doi/10.7554/eLife.18722>.
- O. L Krivanek and P. A Stadelmann. Effect of three-fold astigmatism on high resolution electron micrographs. *Ultramicroscopy*, 60(1):103–113, August 1995. ISSN 0304-3991. doi: 10.1016/0304-3991(95)00050-B. URL <http://www.sciencedirect.com/science/article/pii/030439919500050B>.
- Alp Kucukelbir, Fred J Sigworth, and Hemant D Tagare. Quantifying the local resolution of cryo-EM density maps. *Nature methods*, 11(1):63–65, January 2014. ISSN 1548-7105. doi: 10.1038/nmeth.2727.
- Werner Kühlbrandt. The resolution revolution. *Science (New York, N.Y.)*, 343(6178):1443–1444, March 2014. ISSN 1095-9203. doi: 10.1126/science.1251652.
- Xueming Li, Paul Mooney, Shawn Zheng, Christopher R Booth, Michael B Braunfeld, Sander Gubbens, David A Agard, and Yifan Cheng. Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-EM. *Nature methods*, 10(6):584–590, June 2013. ISSN 1548-7105. doi: 10.1038/nmeth.2472.
- David N. Mastrorarde. Automated electron microscope tomography using robust prediction of specimen movements. *Journal of Structural Biology*, 152(1):36–51, October 2005. ISSN 1047-8477. doi: 10.1016/j.jsb.2005.07.007.
- Robert A. McLeod, Julia Kowal, Philippe Ringler, and Henning Stahlberg. Robust image alignment for cryogenic transmission electron microscopy. *Journal of Structural Biology*, 197(3):279–293, 2017. ISSN 1095-8657. doi: 10.1016/j.jsb.2016.12.006.
- G. McMullan, S. Chen, R. Henderson, and A.R. Faruqi. Detective quantum efficiency of electron area detectors in electron microscopy. *Ultramicroscopy*, 109(9):1126–1143, August 2009. ISSN 0304-3991. doi: 10.1016/j.ultramic.2009.04.002. URL <http://www.sciencedirect.com/science/article/pii/S0304399109001120>.
- Alan Merk, Alberto Bartesaghi, Soojay Banerjee, Veronica Falconieri, Prashant Rao, Mindy I. Davis, Rajan Pragani, Matthew B. Boxer, Lesley A. Earl, Jacqueline L. S. Milne, and Sriram Subramaniam. Breaking Cryo-EM Resolution Barriers to Facilitate Drug Discovery. *Cell*, 165(7):1698–1707, June 2016. ISSN 1097-4172. doi: 10.1016/j.cell.2016.05.040.
- Joseph A. Mindell and Nikolaus Grigorieff. Accurate determination of local defocus and specimen tilt in electron microscopy. *Journal of Structural Biology*, 142(3):334–347, June 2003. ISSN 1047-8477. doi: 10.1016/

- S1047-8477(03)00069-8. URL <http://www.sciencedirect.com/science/article/pii/S1047847703000698>.
- Toshio Moriya, Michael Saur, Markus Stabrin, Felipe Merino, Horatiu Voicu, Zhong Huang, Pawel A. Penczek, Stefan Raunser, and Christos Gatsogiannis. High-resolution Single Particle Analysis from Electron Cryo-microscopy Images Using SPHIRE. *Journal of Visualized Experiments: JoVE*, (123), 2017. ISSN 1940-087X. doi: 10.3791/55448.
- Takanori Nakane, Dari Kimanius, Erik Lindahl, and Sjors Hw Scheres. Characterisation of molecular motions in cryo-EM single-particle data by multi-body refinement in RELION. *eLife*, 7, June 2018. ISSN 2050-084X. doi: 10.7554/eLife.36861.
- Katerina Naydenova and Christopher J. Russo. Measuring the effects of particle orientation to improve the efficiency of electron cryomicroscopy. *Nature Communications*, 8(1):629, 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-00782-3.
- J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, January 1965. ISSN 0010-4620. doi: 10.1093/comjnl/7.4.308. URL <https://academic.oup.com/comjnl/article/7/4/308/354237>.
- Ardan Patwardhan. Trends in the Electron Microscopy Data Bank (EMDB). *Acta Crystallographica. Section D, Structural Biology*, 73(Pt 6):503–508, June 2017. ISSN 2059-7983. doi: 10.1107/S2059798317004181.
- Ali Punjani, John L. Rubinstein, David J. Fleet, and Marcus A. Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nature Methods*, 14(3):290–296, March 2017. ISSN 1548-7105. doi: 10.1038/nmeth.4169.
- Alexis Rohou and Nikolaus Grigorieff. CTFFIND4: Fast and accurate defocus estimation from electron micrographs. *Journal of Structural Biology*, 192(2): 216–221, November 2015. ISSN 1095-8657. doi: 10.1016/j.jsb.2015.08.008.
- Peter B Rosenthal and Richard Henderson. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *Journal of Molecular Biology*, 333(4):721–745, October 2003. ISSN 0022-2836. URL <http://www.ncbi.nlm.nih.gov/pubmed/14568533>.
- John L. Rubinstein and Marcus A. Brubaker. Alignment of cryo-EM movies of individual particles by optimization of image translations. *Journal of Structural Biology*, 192(2):188–195, November 2015. ISSN 1095-8657. doi: 10.1016/j.jsb.2015.08.007.
- Christopher J. Russo and Richard Henderson. Ewald sphere correction using a single side-band image processing algorithm. *Ultramicroscopy*, 187:26–33, April 2018. ISSN 1879-2723. doi: 10.1016/j.ultramic.2017.11.001.

- W. O. Saxton. Observation of lens aberrations for very high-resolution electron microscopy. I. Theory. *Journal of Microscopy*, 179(2):201–213, August 1995. ISSN 1365-2818. doi: 10.1111/j.1365-2818.1995.tb03633.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2818.1995.tb03633.x>.
- Sjors H W Scheres. A Bayesian view on cryo-EM structure determination. *Journal of Molecular Biology*, 415(2):406–418, January 2012a. ISSN 1089-8638. doi: 10.1016/j.jmb.2011.11.010. URL <http://www.ncbi.nlm.nih.gov/pubmed/22100448>.
- Sjors H W Scheres. RELION: Implementation of a Bayesian approach to cryo-EM structure determination. *Journal of Structural Biology*, 180(3):519–530, December 2012b. ISSN 1047-8477. doi: 10.1016/j.jsb.2012.09.006. URL <http://www.sciencedirect.com/science/article/pii/S1047847712002481>.
- Sjors H W Scheres. Semi-automated selection of cryo-EM particles in RELION-1.3. *Journal of Structural Biology*, 189(2):114–122, February 2015. ISSN 1095-8657. doi: 10.1016/j.jsb.2014.11.010.
- Sjors H W Scheres and Shaoxia Chen. Prevention of overfitting in cryo-EM structure determination. *Nature methods*, 9(9):853–854, September 2012. ISSN 1548-7105. doi: 10.1038/nmeth.2115.
- Yong Zi Tan, Philip R. Baldwin, Joseph H. Davis, James R. Williamson, Clinton S. Potter, Bridget Carragher, and Dmitry Lyumkis. Addressing preferred specimen orientation in single-particle cryo-EM through tilting. *Nature Methods*, 14(8):793–796, August 2017. ISSN 1548-7105. doi: 10.1038/nmeth.4347.
- Yong Zi Tan, Sriram Aiyer, Mario Mietzsch, Joshua A. Hull, Robert McKenna, Joshua Grieger, R. Jude Samulski, Timothy S. Baker, Mavis Agbandje-McKenna, and Dmitry Lyumkis. Sub-2 Å Ewald Curvature Corrected Single-Particle Cryo-EM. *bioRxiv*, page 305599, April 2018. doi: 10.1101/305599. URL <https://www.biorxiv.org/content/early/2018/04/21/305599>.
- Guang Tang, Liwei Peng, Philip R Baldwin, Deepinder S Mann, Wen Jiang, Ian Rees, and Steven J Ludtke. EMAN2: an extensible image processing suite for electron microscopy. *Journal of Structural Biology*, 157(1):38–46, January 2007. ISSN 1047-8477. doi: 10.1016/j.jsb.2006.05.009. URL <http://www.ncbi.nlm.nih.gov/pubmed/16859925>.
- Dimitry Tegunov and Patrick Cramer. Real-time cryo-EM data pre-processing with Warp. *bioRxiv*, page 338558, June 2018. doi: 10.1101/338558. URL <https://www.biorxiv.org/content/early/2018/06/14/338558>.
- M. van Heel, G. Harauz, E. V. Orlova, R. Schmidt, and M. Schatz. A new generation of the IMAGIC image processing system. *Journal of Structural Biology*, 116(1):17–24, February 1996. ISSN 1047-8477. doi: 10.1006/jsbi.1996.0004.

- N. R. Voss, C. K. Yoshioka, M. Radermacher, C. S. Potter, and B. Carragher. DoG Picker and TiltPicker: Software tools to facilitate particle selection in single particle electron microscopy. *Journal of Structural Biology*, 166(2):205–213, May 2009. ISSN 1047-8477. doi: 10.1016/j.jsb.2009.01.004. URL <http://www.sciencedirect.com/science/article/pii/S1047847709000197>.
- David Woolford, Ben Hankamer, and Geoffery Ericksson. The Laplacian of Gaussian and arbitrary z-crossings approach applied to automated single particle reconstruction. *Journal of Structural Biology*, 159(1):122–134, July 2007. ISSN 1047-8477. doi: 10.1016/j.jsb.2007.03.003.
- Xuekui Yu, Lei Jin, and Z. Hong Zhou. 3.88 Å structure of cytoplasmic polyhedrosis virus by cryo-electron microscopy. *Nature*, 453(7193):415–419, May 2008. ISSN 1476-4687. doi: 10.1038/nature06893.
- F. Zemlin, K. Weiss, P. Schiske, W. Kunath, and K. H. Herrmann. Coma-free alignment of high resolution electron microscopes with the aid of optical diffractograms. *Ultramicroscopy*, 3:49–60, January 1978. ISSN 0304-3991. doi: 10.1016/S0304-3991(78)80006-0. URL <http://www.sciencedirect.com/science/article/pii/S0304399178800060>.
- Kai Zhang. Gctf: Real-time CTF determination and correction. *Journal of Structural Biology*, 193(1):1–12, January 2016. ISSN 1047-8477. doi: 10.1016/j.jsb.2015.11.003. URL <http://www.sciencedirect.com/science/article/pii/S1047847715301003>.
- King Zhang, Ethan Settembre, Chen Xu, Philip R. Dormitzer, Richard Bellamy, Stephen C. Harrison, and Nikolaus Grigorieff. Near-atomic resolution using electron cryomicroscopy and single-particle reconstruction. *Proceedings of the National Academy of Sciences of the United States of America*, 105(6):1867–1872, February 2008. ISSN 1091-6490. doi: 10.1073/pnas.0711623105.
- Shawn Q. Zheng, Eugene Palovcak, Jean-Paul Armache, Kliment A. Verba, Yifan Cheng, and David A. Agard. MotionCor2: anisotropic correction of beam-induced motion for improved cryo-electron microscopy. *Nature Methods*, 14(4):331–332, April 2017. ISSN 1548-7105. doi: 10.1038/nmeth.4193.
- Dongjie Zhu, Xiangxi Wang, Qianglin Fang, James L. Van Etten, Michael G. Rossmann, Zihe Rao, and Xinzheng Zhang. Pushing the resolution limit by correcting the Ewald sphere effect in single-particle Cryo-EM reconstructions. *Nature Communications*, 9(1):1552, April 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-04051-9.
- Jasenko Zivanov, Takanori Nakane, and Sjors H. W. Scheres. A Bayesian approach to beam-induced motion correction in cryo-EM single-particle analysis | bioRxiv. *bioRxiv*, 384537, 2018. URL <https://www.biorxiv.org/content/early/2018/08/03/384537>.