



## Image Processing in Cryo-Electron Microscopy of Single Particles: The Power of Combining Methods

**Carlos Oscar S. Sorzano, Amaya Jiménez-Moreno, David Maluenda, Erney Ramírez-Aportela, Marta Martínez, Ana Cuervo, Robert Melero, Jose Javier Conesa, Ruben Sánchez-García, David Strelak, Jiri Filipovic, Estrella Fernández-Giménez, Federico de Isidro-Gómez, David Herreros, Pablo Conesa, Laura del Caño, Yunior Fonseca, Jorge Jiménez de la Morena, Jose Ramon Macías, Patricia Losana, Roberto Marabini, and Jose-Maria Carazo**

### Abstract

Cryo-electron microscopy has established as a mature structural biology technique to elucidate the three-dimensional structure of biological macromolecules. The Coulomb potential of the sample is imaged by an electron beam, and fast semi-conductor detectors produce movies of the sample under study. These movies have to be further processed by a whole pipeline of image-processing algorithms that produce the final structure of the macromolecule. In this chapter, we illustrate this whole processing pipeline putting in value the strength of “meta algorithms,” which are the combination of several algorithms, each one with different mathematical rationale, in order to distinguish correctly from incorrectly estimated parameters. We show how this strategy leads to superior performance of the whole pipeline as well as more confident assessments about the reconstructed structures. The “meta algorithms” strategy is common to many fields and, in particular, it has provided excellent results in bioinformatics. We illustrate this combination using the workflow engine, Scipion.

**Key words** Single particle, Cryo-electron microscopy, Image processing, Scipion

---

### 1 Introduction

Cryo-electron microscopy (cryo-EM) is a quickly growing structural technique capable of yielding quasi-atomic models of biological macromolecules [1, 2]. Cryo-EM structures have already found applications in structure-based drug design [3–5]. Additionally, it has the advantage of potentially identifying different conformational states [6, 7]. The recent success of this technique is due to technological advances in sample preparation [8] and microscope,

image acquisition [9], and image-processing technologies [10]. In this chapter, we focus on this latter aspect. At its simplest, a pipeline for image processing enables the parameters for acquiring images from the electron microscope to be determined. These parameters include the gain of the camera, the beam-induced movement, the aberrations of the microscope (most importantly the defocus) for each micrograph/specimen, the orientation of each particle (or class of particles), and possible changes in magnification with respect to the nominal magnification. As in all identifications of parameters in a noisy environment, algorithms will always produce an estimate of those parameters, but they may be correctly or incorrectly identified and our task is to try to discern those parameters that have been incorrectly determined.

For this task, comparing and combining the output of several algorithms is an appropriate approach. The rationale is that the local minima of an algorithm will not be the local minima of another one. In this way, if two different algorithms, with different mathematics underneath, disagree about an estimate, one of the two has to be wrong. On the contrary, if both algorithms agree (within some tolerance), we cannot guarantee that both are right, but at least, it is the best estimate we can have with the tools available. Mathematically, we are interested in unbiased estimates of the parameters. A parameter is considered biased if its expected value (i.e., the average of many repetitions of the estimation process) does not converge to the true (although unknown) parameter. If our parameter estimate is unbiased, we can obtain a better estimate by averaging several estimates. By doing so, we are also reducing the variance of our estimation.

In this chapter, we follow this principle of combining different algorithms as a way of providing a more solid scientific support to structural claims. At present, the strategy of comparing different parameter estimates is not always possible at all the steps along the image-processing pipeline. One of the reasons is that the parametrization of the different processes is not always comparable (e.g., each frame alignment program encodes the beam-induced movement in a different way and the parameter estimates can neither be compared nor averaged). At those places where the comparison and/or averaging is possible, we do it. At those other places in which the comparison is not possible, we simply choose one algorithm that has proven to be robust and to produce good results in experimental cases. We concentrate on those methods currently accessible through Scipion [11], as the integration in a single platform makes comparisons easier.

As an example dataset we have chosen the Brome Mosaic Virus dataset [12] used in the Map Challenge [13] (EMPIAR Entry: 10010, EMDB entry: 6000). This dataset nicely illustrates the difficulties to get good estimates of the underlying parameters (especially those of 3D alignment). The dataset is formed by


424 movies of 37 frames of size  $4096 \times 3072$  pixels at a pixel size of  $0.99 \text{ \AA}$  taken at a JEOL 3200FSC with a DE-12 camera. The resolution reported at EMDB is  $3.8 \text{ \AA}$ , (the highest resolution reported for this dataset was  $3.5 \text{ \AA}$ , [13], the same as the one reported in this chapter). However, as discussed in Soranzo et al. [14], we should not base our full analysis of the results of a reconstruction on a single number associated to some way to measure resolution or reproducibility since the analysis process is much more complex than that (in fact, and just as an example, resolution is locally and directionally dependent, [15]). Indeed, we show in this chapter (and it is well known in the field) that significantly different maps may report the same resolution when assessed against one global number. What is more, the reported resolution is normally in the lower extreme of the resolution histogram [10, 16] so that the true resolution of the structure is typically lower than the one reported by a single number. Measures based on the ability of the map to accommodate an atomic model should be preferred, understanding that these are only possible for resolutions below  $4 \text{ \AA}$ , and in any case the reported resolution number should only be taken as a rough estimate of the quality of the map.

---

## 2 Methods

### 2.1 *From Frames to Valid Micrographs*

Frames are acquired at extremely short exposure times resulting in very low-contrast images, often due to the low count of electron hits. Along the acquisition process particles move under the electron beam. This was one of the reasons for the low-resolution maps of cryo-EM before the introduction of direct detector cameras. These devices have greatly improved the point spread function formally due to the scintillators (converters of electrons into photons) and have allowed quick scanning of many images, referred to as frames, with very little exposure. The set of frames corresponding to the same field of view is called a movie and the average of all the aligned frames is called a micrograph. These frames contain the structural information (Fig. 1 top), but before they can be used they have to be corrected for distortions introduced by the camera and the particle movement induced by the electron beam. The most important distortion introduced by the camera is the so-called camera gain. This refers to the fact that a uniform electron illumination is not transformed into a uniform readout of the camera (Fig. 1 middle). The reason is that each electron hit from the beam is amplified by the electronic circuitry of the camera to be translated into an electric potential that is finally read. This amplification depends on electronic currents that may change over time, and the gain correction image needs to be regularly measured. At this point, we can use the algorithm described in Sorzano et al. [17] in order to verify that the experimental images have been properly beam-corrected.



**Fig. 1** Example of data acquisition (from EMPIAR 10010). Many frames like the one at the top of the figure are acquired per image field. These frames contain the structural information, but they are distorted by the camera (the middle image shows the correction required for the camera) and by beam-induced movement that has to be corrected after correcting for the gain. Once these two distortions are corrected, and after averaging the aligned frames, particles can be identified on the micrographs (bottom)


There are several algorithms to correct for the beam-induced movement. One can think of the beam-induced movement as the finding of a function that for every pixel in the micrograph tells us where that pixel is coming from at every frame (frames are indexed by  $i$ ).

$$I_{\text{mic}}(x, y) = Ii(x + f_x^i(x, y), y + f_y^i(x, y))(1) \quad i$$

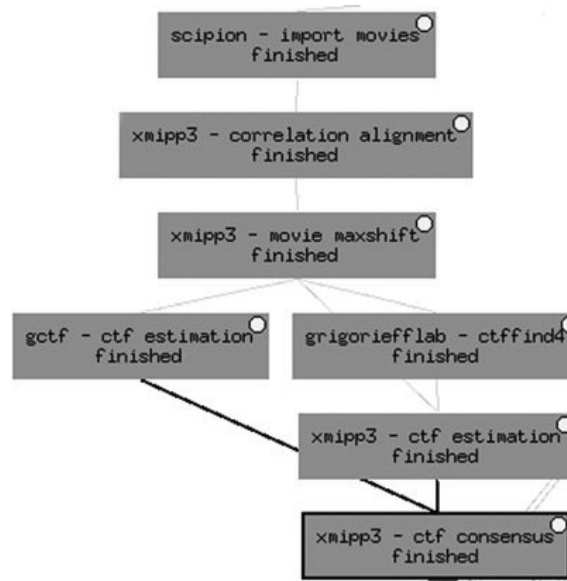
With no beam-induced movement, the functions  $f_x^i = f_y^i = 0$ . In general, we can expand the local beam shift in a Taylor series. The first MotionCorr algorithm [18] can be thought of as a Taylor series of order 0, in which the function is approximated by a constant. MotionCor2 [19] can be thought of as a second-order approximation of the series. Optical flow [20] would be a high order Taylor expansion, thanks to its free-form field. The problem of this latter approach is its computational cost. The solution for trading-off a high order expansion with a low computational cost is to use a deformation field expressed in terms of B-splines [21]. In this chapter, we have used this latter approach as implemented in Xmipp [22]. After estimating the deformation field, we can average the aligned frames to produce a micrograph as the one shown in Fig. 1 bottom.

The parameterization of the beam-induced movement varies among different software packages. Consequently, we cannot estimate this parameter by comparing the output of two different solutions. However, we can make “sanity checks” on the estimated motion field. For instance, the average shift between one frame and the next should not exceed a given threshold (5 Å in our example), and the overall trajectory of the movie should be below another value (15 Å in our example). These values are user defined, and they are meant to prevent incorrectly estimated fields or too quickly moving fields of view to progress along the image pipeline. In this example, only one movie out of the 424 movies failed to meet this requirement.

The next step is to estimate the aberrations of the microscope for each of the micrographs/specimens, the so-called Contrast Transfer Function (CTF) [23]. Again, the parameterization of the different packages is very diverse. However, they all compute the defocus, so that comparisons can be made at this point. In the example developed in this chapter, we calculated the defocus with GCTF [24] and CTFFind4 [25]. We then used the defocus estimated by CTFFind4 as the initialization for the CTF of Xmipp [23]. This latter algorithm has the advantage that it estimates the envelope of the CTF as well as many quality criteria of the CTF regarding its astigmatism, the visibility of the Thon rings, and the quality of the fit, such as the level of ice, correlation of the Thon rings with a 90° version of itself, visibility of the second ring, correlation between the model and the observed data between the first and third zero [26]. 77 out of the 423 micrographs were rejected for some reason related to the CTF (34 because the estimated maximum resolution was above 5 Å, 17 because the correlation between the observed and estimated CTF was below 0.03, 15 because the two CTF estimates did not coincide up to 4 Å [27], 11 due to astigmatism), Fig. 2 shows some of the rejected CTFs. 346 micrographs progressed to the next step. The workflow used for micrograph selection is shown in Fig. 3.



**Fig. 2** Examples of rejected micrographs



**Fig. 3** Workflow used in this example to align the movie frames, estimate the CTF, and select micrographs according to their CTF quality

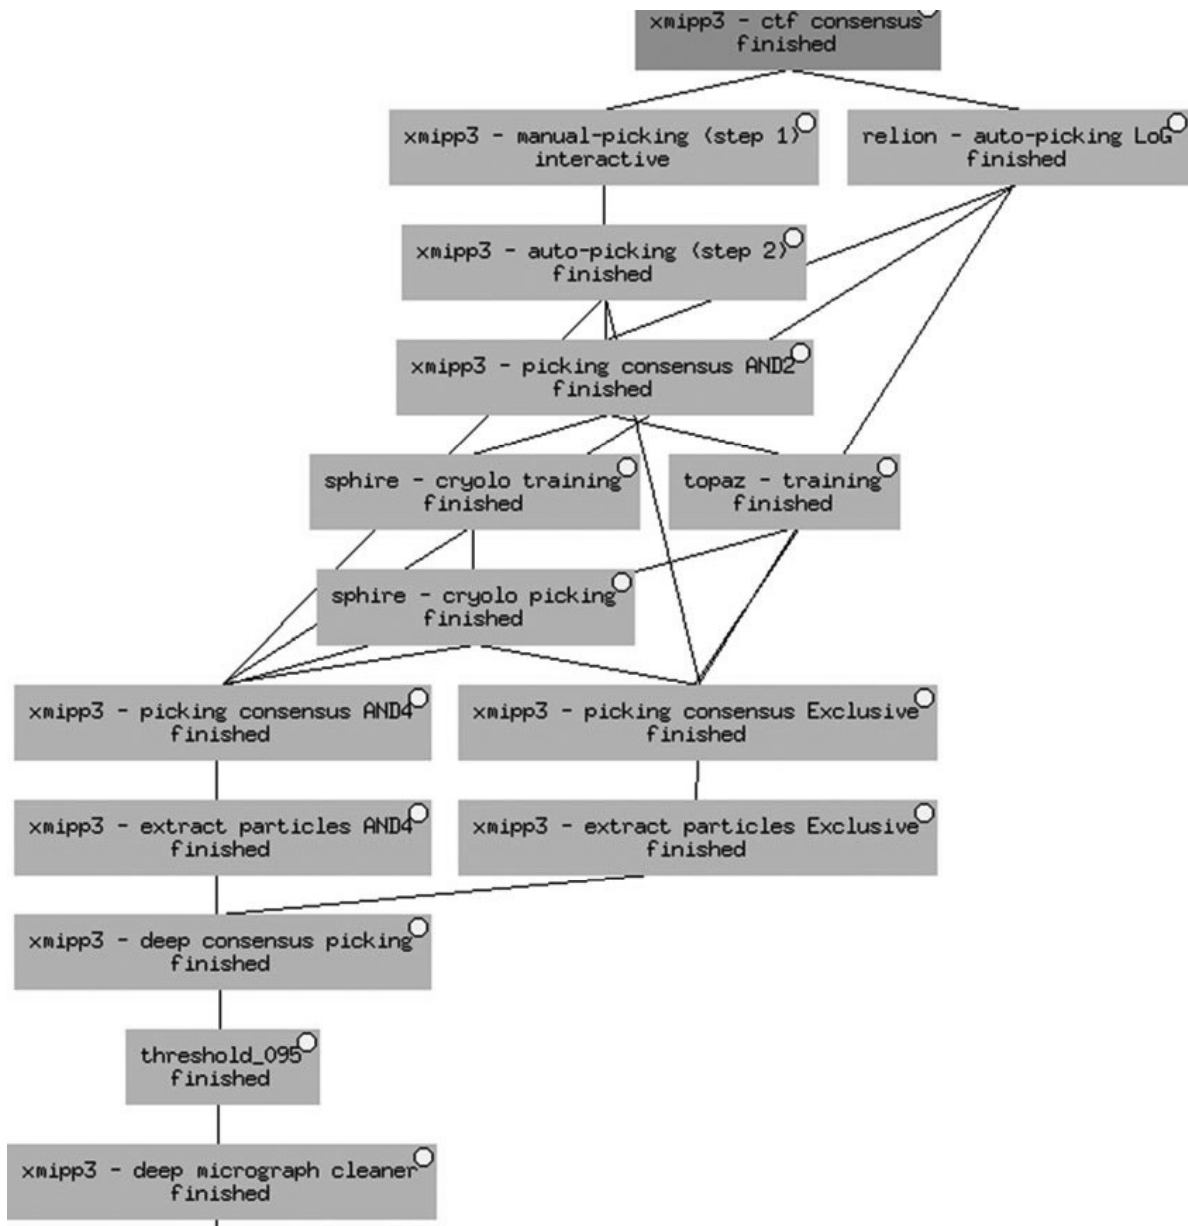
## 2.2 Finding Particles in Micrographs

The next step in the data analysis is to find particles (instances of the macromolecular complex we are working with) in the micrographs (see Fig. 1 bottom), where we refer to finding the coordinate of the center of the particle. There are three families of particle pickers, explained here in increasing order with respect to the total information they require. The first family consists of algorithms whose only input information is the particle size. Appion DoG picker [28], Relion LoG picker [29], and Sparx Gaussian picker belong to this family. The second family comprises algorithms that learn from the user or a pre-trained set how a particle looks like. These algorithms normally extract features from the images and employ a classifier to decide whether it seems to be a particle or it does not. Xmipp picker [30] and all the neural network pickers like Topaz [31] or Cryolo [32] belong to this second family. The third family

consists of algorithms that use image templates that are correlated with the micrographs at various orientations to find similar particles. Relion template picker and Gautomatch belong to this latter family. In very broad terms, the performance of the second and third families is similar, and both are better than the first family. As a rough estimate, although it obviously depends on the dataset, most algorithms have a false-positive rate between 10% and 30% (i.e., 10–30% of the found particles are not real particles) and a false-negative rate in a similar range (i.e., 10–30% of the true particles are missed). The current trend is to pick “almost everything” and then sort out the incorrectly selected particles during a 2D cleaning step (as well as other types of cleaning procedures).

In this example, we used the workflow depicted in Fig. 4. We used Relion LoG picker with a box size of 350 pixels (it found 41,695 coordinates). We also picked the particles using Xmipp picker that learns from the user the kind of particles he is interested in (it found 24,386 coordinates). We then identified those coordinates that were found by both pickers (19,173 of them) and used this subset as a positive set for training Cryolo (it found 28,596 coordinates) and Topaz (it found 28,180). Among all pickers they found 49,819 unique coordinates. 16,887 were found by all of them and 19,106 were found by only one of them. We used these two sets as the positive and negative training set, respectively, of a deep consensus approach [33]. The positive set is normally formed by well-centered particles, while the negative set is formed by off-centered particles or contaminants (see Fig. 5). The neural network of the deep consensus learns to distinguish between these two different kinds of particles and assigns a score between 0 (bad) and 1 (good) to each coordinate to assess how good a particle is. We used a threshold of 0.95 resulting in 32,880 particles. We see that this set is twice the size of the set of particles found by all of the algorithms (so that we have a lower false-negative rate). We then submit these coordinates to deep micrograph cleaner [34] that removes those coordinates that fall on for example, contaminations, aggregates, ice crystals, and carbon edges. Only 0.8% of the coordinates were in these kinds of bad regions (showing the power of deep consensus to eliminate these bad particles). The resulting coordinates are shown in Fig. 5 in which we can see that there is a good compromise between having a high positive rate, with low false positives and low false negatives.


Despite the careful particle selection performed by the very sophisticated algorithms above (machine and deep learning algorithms), there could still be images in the dataset that do not really correspond to particles. The screening method described in [35] aimed at identifying large deviations from the dataset. Its application to the particles coming out from the workflow in Fig. 4 identified 230 particles (0.7% of the dataset) that did not follow the general trends of the dataset. Some of them can be seen in



**Fig. 4** Workflow used in this example to find particles in the micrographs

Fig. 6. Note that this algorithm is only capable of identifying gross deviations from the main population and that the small amount of particles (0.7%) is an indicator of two things: first, the particle selection described above is pretty accurate; second, all algorithms, no matter how sophisticated they are, have false positives and false negatives and the false positives of one algorithm do not need to be the false positives of another algorithm. In this regard, this is a good example of the need to use several methods to identify possible “pathologies” in the dataset.






**Fig. 5** Example of coordinates found for a particular micrograph by: at least one of the algorithms (top left), all of the algorithms (top right), only one algorithm (bottom left), deep consensus and micrograph cleaner (bottom right)


### 2.3 2D Classification

The principle of using multiple algorithms to perform the same task is well illustrated by the 2D classification step that we perform next. The goal of 2D classification is to group particles with similar shape, understanding that the shape is determined by the projection direction of the particle and the macromolecule that is imaged (its composition and conformational state). We submitted all the particles surviving the previous steps (a total of 32,299) to 2D classification. We did it using CryoSparc [36] and CL2D ([37]). The 2D classes of each one of the programs are shown in Fig. 7. At the sight of these results, we make two important notes. The first note is that CryoSparc2D classes tend to be either much larger or much smaller than the ones of CL2D. The mathematical reason for this was described in Sorenzo et al. [37] and has to do with the fact that classes with many images assigned (no matter if the assigned images really look like the class average or not) tend to have lower noise and attract even more images. This is a characteristic observed in Relion and CryoSparc, and less so in CL2D. Despite this attraction problem, the less populated classes of CryoSparc still represent spherical particles, which indicate the good quality of the particle selection performed up to this point (the typical Relion or CryoSparc 2D classification result has a few classes with many particles,




**Fig. 6** Example of particles that do not follow the general trends of the dataset (Vargas et al., 2013)

and many empty classes with mostly noise). The second note is that CL2D classes are sharper, and they were capable of identifying sets of particles with a varying illumination background (marked in red in Fig. 7 bottom). This intensity gradient largely drives the alignment of these particles, and they should be excluded from further analysis as they are not well aligned. 3852 particles belonged to these classes with illumination gradients. Additionally, CL2D has the possibility of looking for outliers within the classes [38] Outliers are defined as those particles whose Mahalanobis distance to the centroid of the cluster is larger than a given threshold (typically, 3). Particles with a distance smaller than the threshold are referred to as the core of the class. CL2D is a hierarchical algorithm in which particles are first classified into a small number of classes (typically, 4), and then these classes are split into 2 (resulting in 8, 16, 32, ... classes). Particles are allowed to choose any of the existing classes at any moment. Two particles belong to the stable core of a class if they were always together along the classification process (i.e., when there were 4, 8, 16, ... classes, they were always together). We have observed that empty particles and contaminants tend to be randomly spread over the classes and, consequently, they tend to jump between them. 6.6% of the particles were identified as outliers within its class, resulting in a total of 26,793 particles. At this point, we have a set of particles assigned to 2D classes that are only



**Fig. 7** 2D classes calculated by CryoSparc2D (top) and CL2D (bottom)

roughly centered. Therefore, we can refine the coordinates within the micrograph so that they correspond better to the particle center. By doing so, we will be able to identify that some of the allegedly independent coordinates were actually pointing to the same particle (see Fig. 5 bottom right). The protocols within Scipion for this analysis are Xmipp center particles and Xmipp remove duplicates. This step reduced the set of particles by 9%, leaving 24,199 particles available for 3D reconstruction. Another important lesson from this part of the analysis is that among all




**Fig. 8** Example of unique coordinates after 2D classification, centering particles and removal of duplicated particles

identified “particles” at the beginning, selected by at least one sensible and sophisticated algorithm (49,819), less than half of them (24,199) were really independent and supposedly to be good particles (see Fig. 8). This is in agreement with standard observations in the field in which many of the “picked particles” are discarded along the 3D analysis. We should see it as an indication of the fact that most of the picked particles were not really particles, but only suggested coordinates (some of them even pointing to the same particles).

All the 2D classification analysis was performed with particles whose pixel size was 3.46 Å (the box size at this sampling was  $100 \times 100$ ). There are two reasons for this. The first one is that they occupy less space and all calculations are faster. For this 2D screening task, we do not need “atomic” resolution. The second one, and maybe more important, is that making the images smaller removes a lot of noise that may cause image misalignment. In this way, we are increasing the signal-to-noise ratio and having a better analysis with a resolution adapted to the needs of the task at hand.

#### **2.4 Constructing an Initial Map**

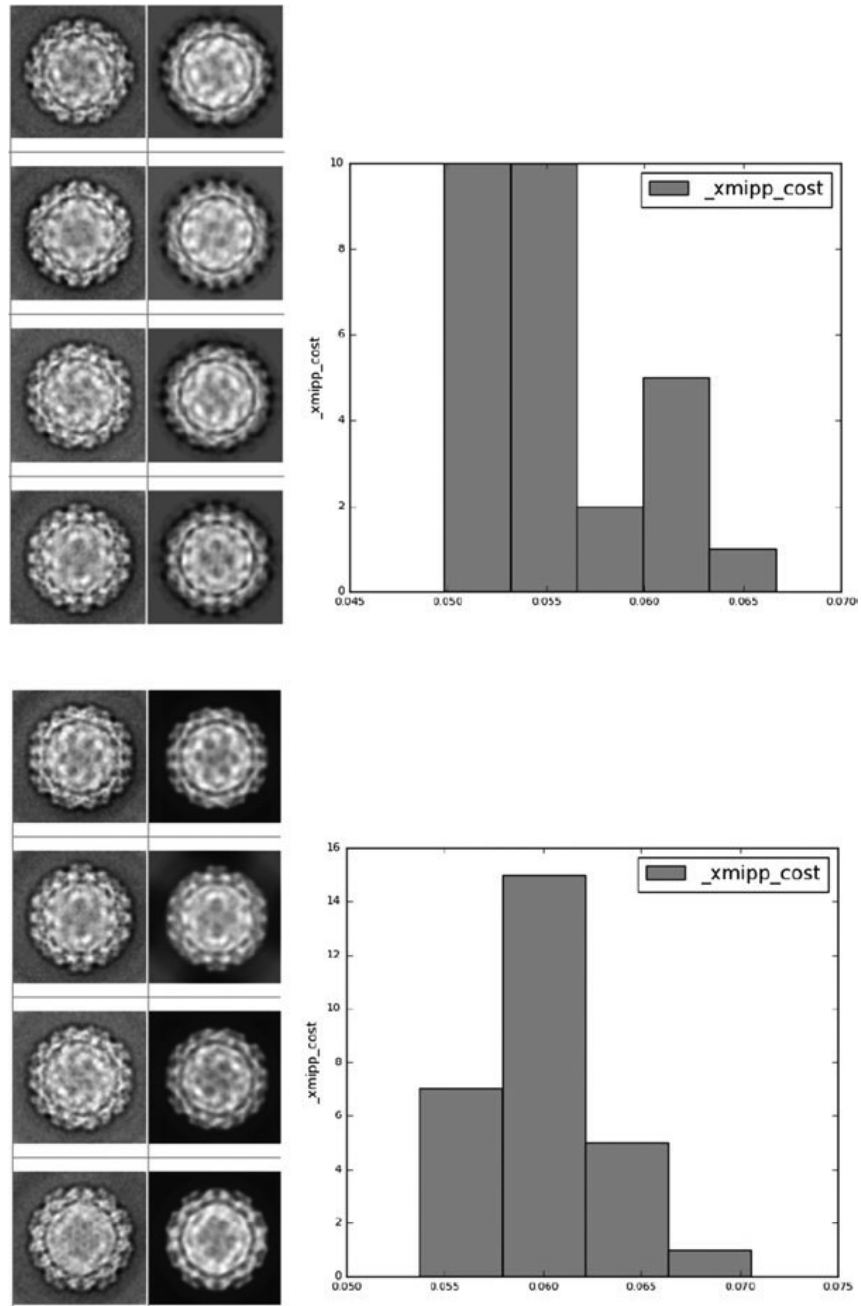
Once we have selected a set of particles, we may proceed to construct an initial volume with them. This is one of the critical steps of the whole procedure because if a bad volume is constructed, it is very likely that the subsequent algorithms cannot escape from this local minimum. Scipion integrates many different algorithms for this task, including Xmipp Ransac [39] and Reconstruct Significant [40], CryoSparc [36], Relion Stochastic Gradient Descent [41], Eman [42], and Simple [43]. All these algorithms employ different flavors of stochastic optimization (an optimization that is not too greedy and allows steps in which the goal function is worse than it was in the previous step), and a simplified optimization landscape (normally by filtering the images, reducing their size, or working with class averages rather than raw images). The goal of these two



**Fig. 9** Example of incorrect (left) and correct (right) initial volume for the Brome Mosaic Virus


strategies is to avoid local minima and try to find the global minimum. However, this is not always achieved for all the proposed volumes. Figure 9 shows an incorrect (local minimum) and correct (close to the global minimum) initial volume proposed by two different algorithms (in this case RANSAC, incorrect, and reconstruct significant, correct; RANSAC proposes a set of candidates and only one of them, shown in the figure, was incorrect). This example was relatively easy, and most initial volume algorithms returned a correct structure. However, this is not always the case and, depending on the dataset, most of the candidates may correspond to local minima. The suggested procedure in this chapter is to run multiple algorithms proposing a diverse population of candidates to initial volume. Traditionally, it was the user who had to choose one of them to continue the processing. At present, there are algorithms that are capable of considering all these candidates and, either by letting them evolve (swarm consensus, [44]) or by comparing them [45] automatically decide (normally correctly) a suitable initial volume.

All this analysis can be performed at relatively low resolution (pixel size 3.6 Å in our example) as a way to speed-up calculations, and more importantly, to smooth the goal function landscape by removing noise. One of the most useful analyses before going further is to compare the re-projections of the initial volume with the 2D class averages. Even the least matching classes should be in good agreement. Figure 10 shows these comparisons for the two initial volumes proposed in Fig. 9. Differences are subtle (sometimes they are much more obvious), showing the difficulty encountered by the initial volume algorithms, but an analysis of the similarity between these re-projections (see Fig. 10) indicates which one of the two is better (the bottom one, since the histogram of correlations is shifted towards higher values). If experimental SAXS data is available, then the SAXS curve of the different



**Fig. 10** Comparison of the re-projections of the initial volume and the 2D classes for an incorrect initial volume (top) and a correct one (bottom). For each of the volumes, the least correlating 2D class averages are shown (left) along with the corresponding re-projection (right) and the histogram of all correlations within the comparison

proposals to initial candidate can be simulated and compared to the experimental SAXS curve. Jimenez et al. (2019) [46] shows that this strategy can indeed distinguish between incorrect and correct initial volumes.



**Fig. 11** The 3D Classification of 24,199 particles of the Brome Mosaic Virus into two classes by Relion. The first class (left) received 68.4% of the particles, while the second class (right) received 31.6%

### **2.5 Achieving an Homogeneous Population**

The next step in the pipeline is to separate the selected particles into structurally homogeneous datasets. This is done through 3D classification steps. Using the same rationale as in the previous steps, we may perform this analysis at an intermediate pixel size (1.73 Å in our example, resulting in a box size of 200 pixels). The two algorithms for 3D classification in Scipion are Relion and CryoSparc. We used Relion to separate into two classes the 24,199 particles selected in previous steps. 68.4% of the particles went to a class similar to the initial volume and 31.6% to another class (see Fig. 11).

In most publications, the classification step is executed only once. However, we were interested in the reproducibility of this class assignment, and we ran this process three times. The number of particles assigned to the first class ranged from 68.4% to 82.8%, but the two classes were invariably the same. However, and very interestingly, only 55.8% of the images were always assigned to the first class in all three classifications (this set can be calculated using the consensus classes 3D protocol of Scipion). This is an indication of the instability of the classification process. The belonging or not to a class is another parameter that needs to be estimated in a very noisy environment, and consequently its estimation is prone to errors. This statistical principle is not generally well rooted in the 3DEM community, and most papers perform a single classification step taking its result as the “ground true” classification of the dataset at hand. The take-home message should be to trust only those subsets of images that consistently have been assigned to the same class because the class assignment is another parameter to estimate (and one of the most important ones), and there could be errors in its estimation. Naturally, the key point is to know how this clear instability associated to the angular estimation translates into the map itself.

In very general terms, let us assume that among all the particles assigned to class 1, there is a fraction  $\alpha$  (between 0 and 1) of particles that truly belong to class 1 and  $1 - \alpha$  that truly belong to class 2. Then, by miss-estimating this  $1 - \alpha$  fraction, the estimated class 1 volume would be (in a simplified manner, because the true dependence also depends on the angular assignment).

$$\widehat{V}_1 = \alpha V_1 + (1 - \alpha) V_2$$

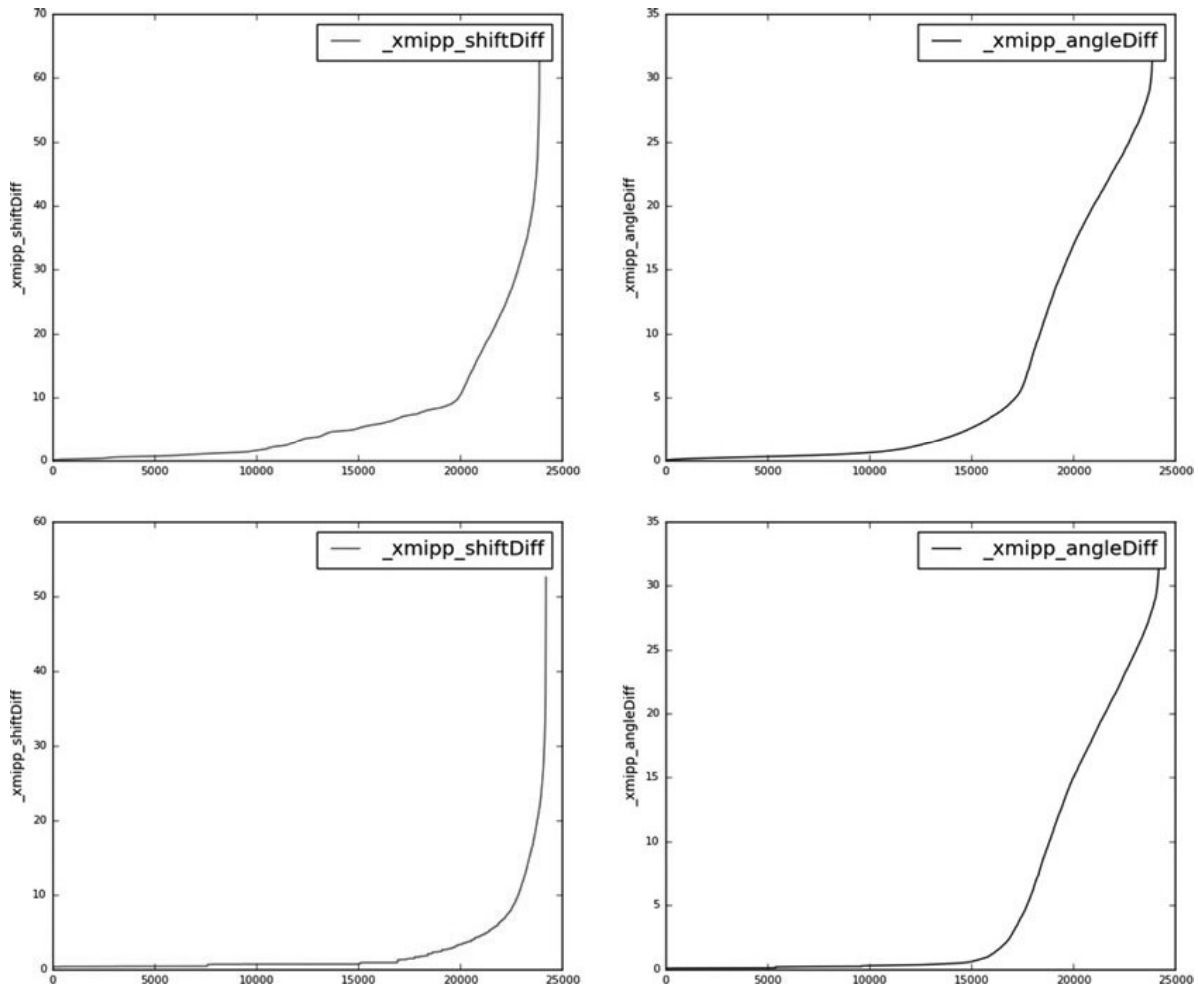
Still, the way the additional term  $(1 - \alpha) V_2$  translates into the map is virtually impossible to know. If errors in the alignment follow a more or less random pattern, then the result will be the introduction of a blur-like in the map, which may not affect its interpretation after sharpening and applying a threshold. Indeed, this is the case for main class 1 since the map does not virtually change. However, this is not the case for class 2, where the lack of analysis of alignment instability would have led to totally wrong conclusions, as we present in the following.

The fact that between 31.6% and 17.2% of the particles are assigned to class 2 would lead us to think that class 2 is a truly existing class in the dataset (we may even try to find its biological role). However, we suspected that this was an image-processing artifact. For testing this hypothesis, we selected the particles assigned to class 2 and ran an initial volume protocol (Relion stochastic gradient descent, in this example) obtaining the volume in class 1! Actually, in this dataset there is only one distinguishable conformation. In this example, class 2 comprises a mixture of particles with a correct angular assignment and particles with a mirrored angular assignment. This exercise brings two very important take home messages: (1) class separation could be partially or totally artifactual due to the image processing; (2) even if the parameter of class belonging is correctly estimated, the angular orientation of a particle with respect to the class may still be incorrect. The only way of verifying if this is our case is by estimating the class parameter and angular assignment multiple times, preferably with different algorithms based on different mathematics and only trust those classes and angular assignments that are consistent between algorithms and runs. Still, being consistent is not a guarantee of being correct. But, being inconsistent is a guarantee of being incorrect.

## **2.6 Refining an Homogeneous Population**

Once we have divided the set of particles into structurally homogeneous datasets (in our example, the 24,199 particles belonged to a single structural class) the final step is to align the particles with respect to an initial volume that must be refined. Xmipp highres [38], Relion [47], and CryoSparc [36] can be used for this task within Scipion. From the previous steps, we now understand that we are estimating the pose parameters of the particles and that these





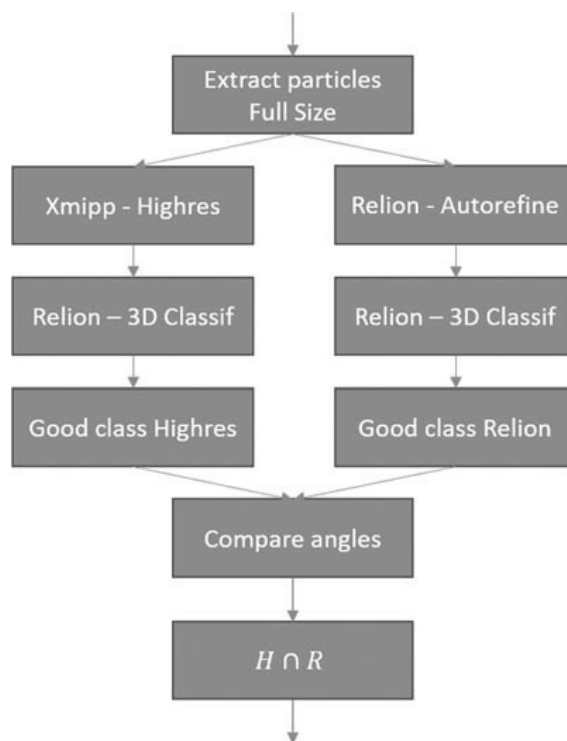
**Fig. 12** Comparison of the shifts (left, in pixels) and angles (right, in degrees) of Relion and Highres (top), and two independent runs of Relion (bottom)

parameters may be estimated with noise. For this reason, we ran Relion and Xmipp Highres. Both algorithms agreed in the angular assignment of the image for about 50% of the particles (see Fig. 12, top). For illustrative purposes, we show in the same figure (bottom) the shift and angular comparison of two independent runs of Relion (they agree for about 75% of the particles).

This disagreement between two reasonable angular assignment methods (or even a method with itself) speaks about the difficulty of estimating parameters in noisy environments and the need to confirm the validity of those parameters. As with the incorrect 3D classification, if a fraction  $\alpha$  of the particles are correctly aligned, while  $1 - \alpha$  are incorrectly aligned, our reconstruction will be

$$\hat{V} = \alpha V_{\text{correct}} + (1 - \alpha) V_{\text{incorrect}}$$

Identifying the correctly assigned particles is a difficult task with experimental data as the ground truth is never known. At this point, we have found useful performing a 3D classification of the aligned



**Fig. 13** Starting from a common set of particles at full size, we perform two independent angular assignments with Xmipp Highres and Relion Autorefine. The outputs of both programs are classified into 3D without re-estimating the angular orientation, and the images assigned to a good class are identified. We then compare the angular assignment of these good images between Highres and Relion. We label as  $H \cap R$  the set of images with similar angular assignment in Highres (H) and Relion (R)


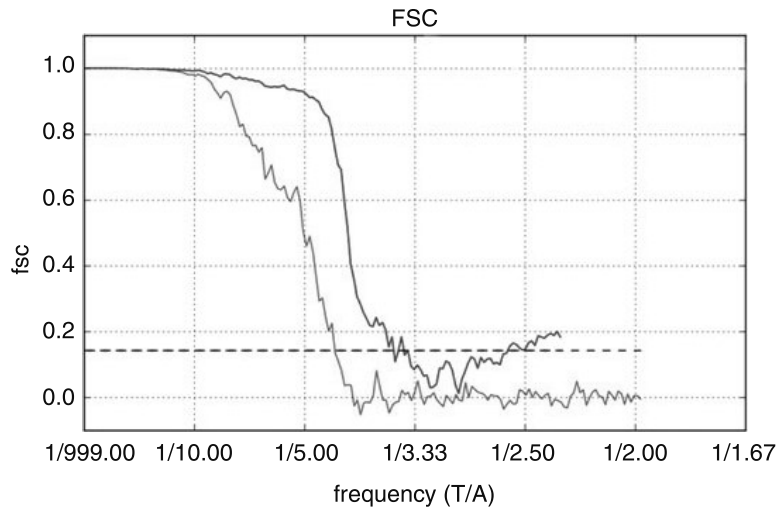
particles in two classes without re-estimating the angular parameters, but just the class parameters. We have observed that normally the set of particles is separated into a class that gives rise to a high-quality reconstruction (and we assume that most of the particles in that class are reasonably aligned) and another class with lower resolution (and we assume that most of the particles in that class are not well aligned). We may then compute the intersection of the two good classes (one from Highres and one from Relion). 10,471 particles (47.2% of the original subset) were classified into both classes as good, and the corresponding angular assignments differed in less than  $2.5^\circ$  and 4 pixels (1% of the image size). See Fig. 13 for a graphical summary of this strategy. Although we can never guarantee that the angular assignment is correct with experimental data, we know that this set of 10,471 particles were assigned very similar 3D pose by two independent 3D reconstruction algorithms and that they were assigned to a high-quality 3D class by two independent executions of a 3D classification algorithm. At this point, and to the best of our ability to estimate parameter errors, we would say that these images correspond to well-assigned particles

within a relatively homogeneous class. A drawback of this approach is that about half the particles have not made their way to the final reconstruction, and indeed this is a problem for an example in which the resolution is known to be limited by the number of particles (see Fig. 21 and related comments). Perhaps better strategies can be derived in the future to reduce this loss, possibly following ways conceptually comparable with some of the methods for “rescuing” particles while picking that were presented in past section, but they will always be rooted in the key realization that some way to estimate parameter instability has to be incorporated into the whole approach.

### **2.7 Final 3D Reconstruction with Consensus Geometry**

Scipion has the possibility of merging different sets of particles while keeping their angular assignment. Additionally, the angular assignments by any of the two methods (Xmipp Highres or Relion Autorefine) are both expressed in a common geometrical framework [48] so that we may continue the 3D reconstruction process benefiting from our best estimates disregarding their origin (although CryoSparc has not been used in this experiment, the internal angular consistency within Scipion would be the same). We used Xmipp Highres local iterations for this task. This algorithm allows local refinement of the 3D pose parameters (Euler angles and in-plane shifts), as well as possible anisotropic magnification errors, local defocus values, and gray normalization parameters. In Fig. 14, we show the FSC of Highres and Relion Autorefine as well as representative slices of the two reconstructions. The calculated resolution by Highres is 3.5 Å. However, this number should not be used as the sole measure of map goodness. More interesting than the low end of the FSC curve is its behavior before it starts to fall down (before the 0.5 threshold). We observe that the Highres reconstruction is much more consistent (closer to 1) in a wider range of frequencies. This behavior of the Highres reconstruction is also observed in the slices, as shown in Fig. 14, where the Relion reconstruction seems to be a low-pass filtered version of the reconstruction of Highres. It should be noted, however, that Highres introduces some nonlinear constraints in the reconstruction process for noise suppression; from this point of view, it can be considered that it incorporates some form of masking, while in Relion this is a separated process. The introduction of these nonlinear constraints also explains the fact that the FSC does not always fall to zero.

Local resolution can be measured in Scipion by any of the standard tools in the field: Blocres [49], Resmap [50], MonoRes [10], or DeepRes [51]. We note that this latter method is mask-invariant, which will be an important feature later on in this work. We have found that the resolution reported by the FSC is typically in the lower extreme of the resolution histogram reported by the local resolution tools [51]. In this case, most voxels are in the range

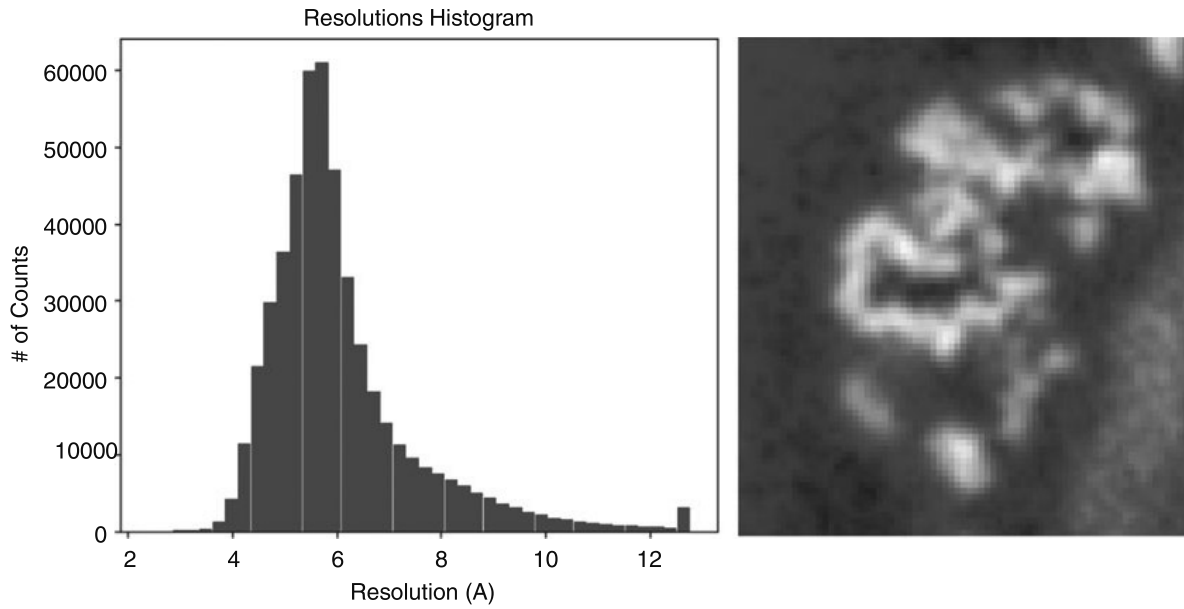


**Fig. 14** Top: FSC of Relion Autorefine (green) and Xmipp Highres local refinement (blue) after the particle selection and angular assignment strategy described along the chapter. Bottom: Representative slice of the reconstruction with Highres (left) and Relion Autorefine (right)

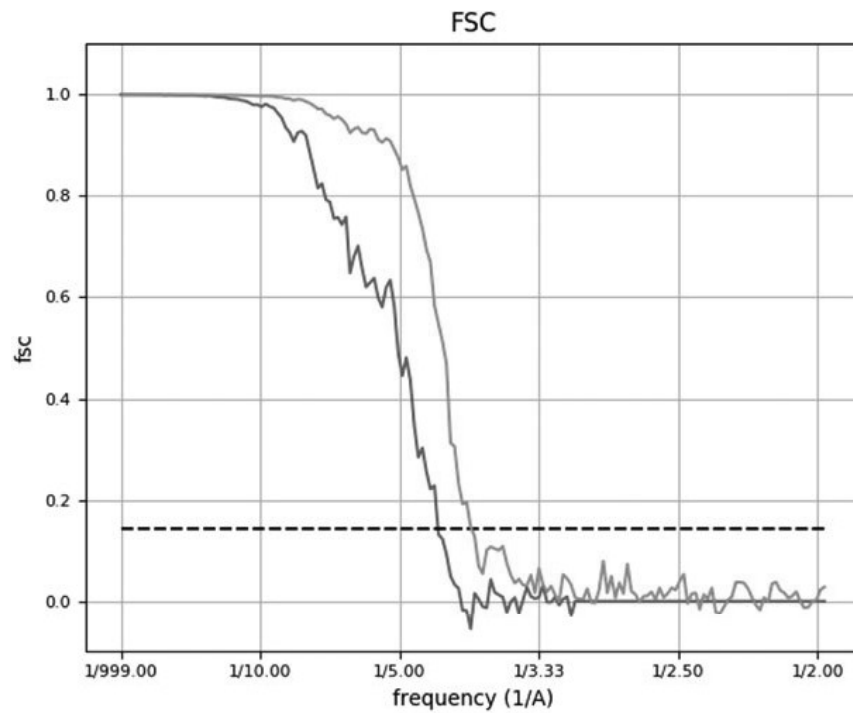
between 4 and 6 Å as reported by DeepRes, and this estimation is consistent with the visual appearance of a close-up of the capsid structure (see Fig. 15).

## 2.8 Post-Processing

At present, the two most widely post-processing tools are masking and B-factor correction [52]. The FSC is insensitive to the B-factor correction because it is an isotropic filter, but it is affected by the mask [14]. Actually, the improvement in FSC that we see in Relion post-processing (see Fig. 16) is purely due to the change of mask between the FSC measured during the reconstruction and one resulting from the mask used in the post-processing. The dependence on the mask poses a real problem to the FSC as an objective evaluator of the quality of the reconstructed map, as different reasonable masks (or even over-smooth masks) result in different resolution measurements (see Fig. 17).

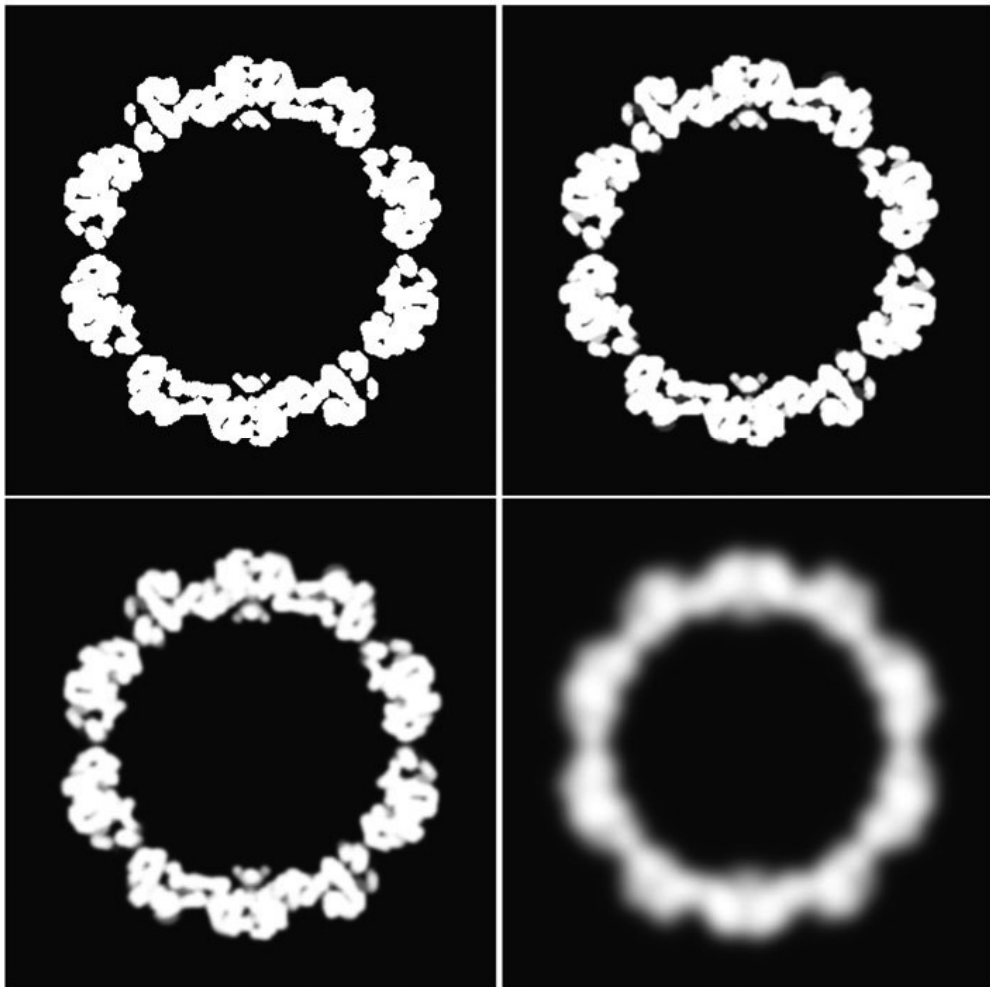
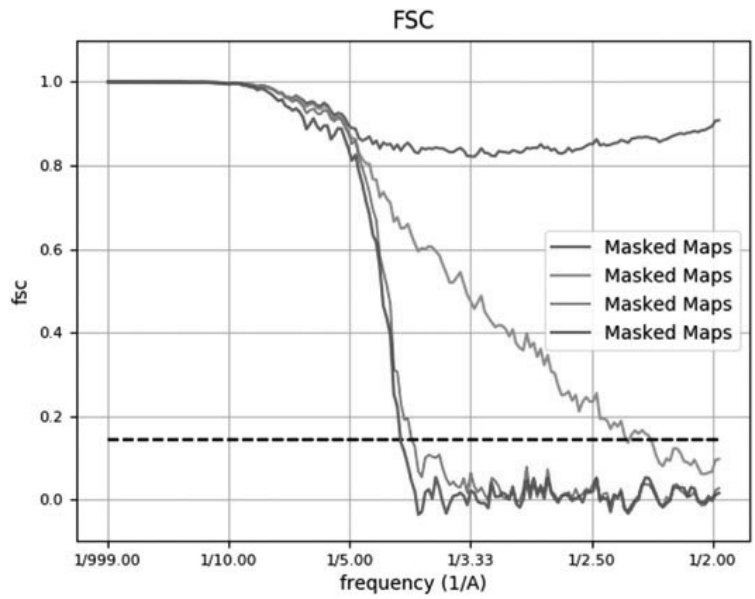


**Fig. 15** Left: Local resolution histogram calculated by DeepRes. Right: Zoomed version of one of the capsomers of the Highres slice shown in Fig. 14




**Fig. 16** Difference observed in Relion post-processing between the unmasked (blue) and masked reconstructions (orange)

Beside the dependence on the mask, the current post-processing practice of boosting high frequencies by applying a B-factor has the problem that it does not take into account local



**Fig. 17** Top: Different FSCs resulting from different masks. A tight, sharp mask (top left) results in an FSC that does not drop from around 1. If this mask is smoothed with a Gaussian of  $\sigma = 1$  (almost undistinguishable by eye with respect to the sharp mask), 2 or 10 (clearly oversmoothed), the FSC drops from almost 1 to a different curve. Interestingly, the FSC does not distinguish the maps masked with masks as different as the ones convolved with  $\sigma = 2$  and  $\sigma = 10$




**Fig. 18** Top: Representative slice of the B-factor corrected map by Relion post-processing (left) and Local Deblur (right). Bottom: Zoomed regions of the maps above

differences in cryo-EM maps (usually studied by the local resolution) and produces maps whose Fourier spectrum does not match the expected decay predicted by the diffraction theory [53]. Alternative methods have been proposed based on the matching of the spectrum falloff of the reconstructed map to the falloff of a fitted atomic model [54] or the use of the local resolution to locally deblur the reconstructed map [51]. The drawback of the first approach is that it requires fitting an atomic model to the reconstructed map, which is a whole job in itself, but it works very well otherwise. In Fig. 18, we show the maps produced by B-factor correction of the Relion autorefine map and by local deblur of the Xmipp Highres. The appearance of both reconstructions is totally different at the level of noise (B-factor is expected to boost noise, while local deblur to suppress it) and, consequently, map interpretability is directly affected by the appearance of the map we are looking at. Moreover, in Ramírez-Aportela et al. [51] we showed that the decay of the deblurred maps in Fourier space corresponds

to the expected behavior of biological macromolecules as predicted by the diffraction theory. The local resolution histogram of the deblurred map, as measured by DeepRes (which is not affected by masks), shifted from the 4–6 Å region to the 3.5–4.5 Å.

## 2.9 Validation

An important part of the image-processing pipeline is the checking of some necessary conditions that cannot guarantee the correctness of the map, but failing to meet them guarantees its failure. First, we must visually inspect the slices of the map: stripes, especially radial, or artifacts outside the map should not be present. This visual inspection is most useful in the raw result of the 3D reconstruction process (Fig. 14). The angular distribution of the map should also be inspected. For symmetric maps (especially highly symmetrical like the one used in this work), it is difficult to evaluate the homogeneity of the angular distribution as angles are only estimated within the asymmetric unit. We may break the symmetry by randomly assigning to each particle one of the equivalent positions in the projection sphere and, then, evaluate whether or not the angular distribution is uniformly covered. In Fig. 19, we show the angular distribution (with broken symmetry) of the example developed along the chapter. In spite of not being perfectly uniform, it is sufficiently varied as to prevent large regions of the Fourier space



**Fig. 19** Projection sphere in which the projection direction of all particles is represented



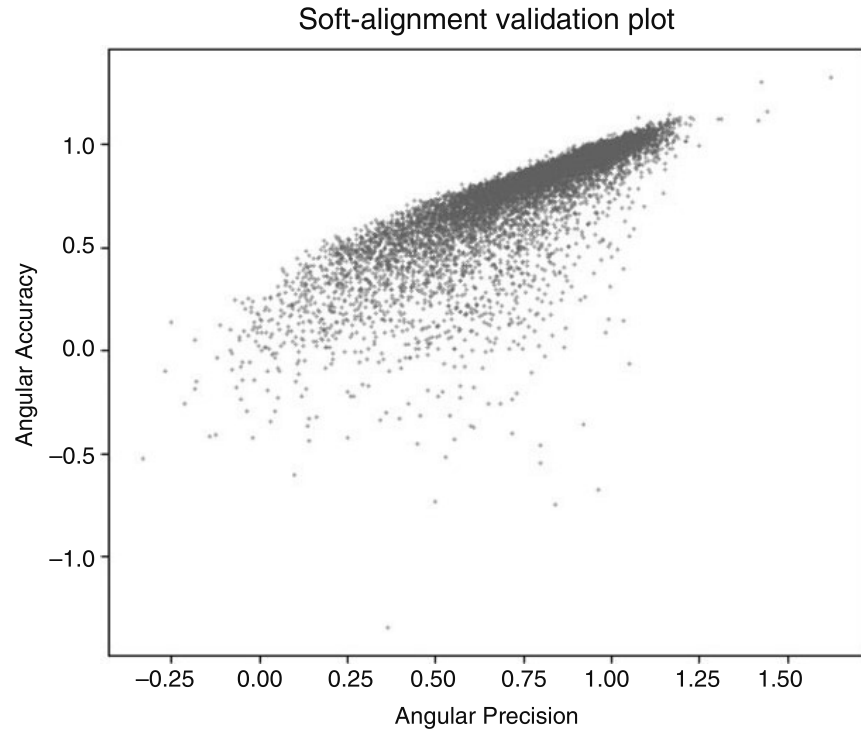
not being measured. Although there are algorithms to quantify the evenness of the angular distribution [55], these are not generally applied and thresholds to decide whether the angular distribution is sufficiently good or not are not available.

Vargas et al. [56, 57] proposed a couple of metrics to measure the “alignability” of the particle set (they referred to them as soft validation indicators). The first measure (angular precision) analyzed the distance between the top  $N$  best angular assignments (particles whose best  $N$  angular assignments are scattered over the projection sphere are less reliable than particles whose best  $N$  angular assignments are all clustered around the same projection direction; by default  $N = 7$ ). The second measure (angular accuracy) analyzes the consistency between the final angular assignment, which followed a global to local optimization with the risk of getting trapped into a local minimum, and a global angular assignment performed de novo when the final structure is known (those particles for which the two angular assignments coincide are more reliable than those for which they do not). These measures are translated into a score per particle between  $-1$  and  $1$ . Ideally, good particles should have a score above  $0.5$  in both metrics. These two metrics are accessible in Scipion under the protocol multi-reference alignability and the result for this example is shown in Fig. 20. The average precision and accuracy in the dataset are both  $0.89$  that is well above the  $0.5$  threshold that is required to have a reliable angular assignment.

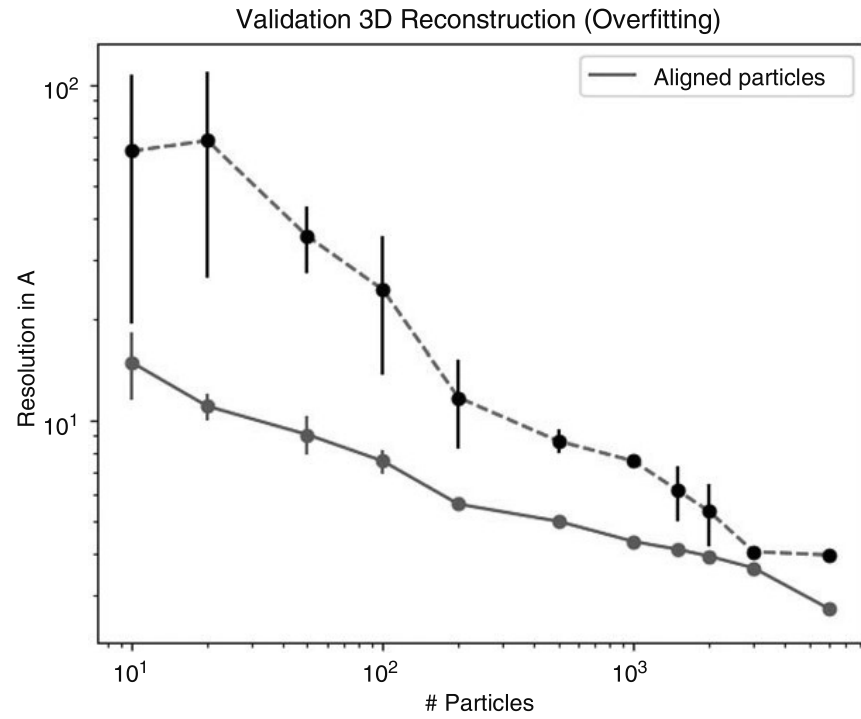
Finally, we may wonder whether by adding more particles we would gain more resolution or, otherwise, the resolution is limited by other factors (heterogeneity, misalignment, low-resolution images, microscope aberrations, etc.). This can be determined by a ResLog plot [58] that shows the increase in resolution as new particles are added. Heymann (2015) [59] showed that an increasing number of pure noise particles also increased the resolution of the reconstructed volume. The idea is that the resolution obtained at a given number of noise images (i.e., no particles inside that can also be referred to as noise particles) should always be smaller than the resolution obtained with the same number of true particles (see Fig. 21). However, if the map is overfitted, the resolution obtained with noise particles reaches a similar level to the one of supposedly true particles. This plot also answers the question of whether the resolution is limited by the number of particles or by other factors. If the plot saturates (reaches a plateau), then the resolution is not limited by the number of particles since adding more particles does not seem to increase the resolution. In our example, the resolution is limited by the number of particles.

## **2.10 Interpretation and Model Building**

Once we have validated the map, we may proceed to the final step: visualizing the map and interpreting it. Isosurfaces at a given threshold or any other more sophisticated segmentation tool (like




**Fig. 20** The angular precision and accuracy for the set of particles used in the 3D reconstruction of the virus shown in this chapter




**Fig. 21** Plot of the resolution vs the number of particles used for the reconstruction. The solid line at the bottom part refers to true particles, while the one above corresponds to noise particles

the one based on the False Discovery Rate, [60]) play an important role in this regard as they help to provide a continuous surface that contains the macromolecule (see Fig. 22).

Biological conclusions can be derived at this point. However, the most common next step is to build an atomic model with the help of the constraints given by the EM measurement. The map and atomic model together are usually referred to as a hybrid model (see Fig. 23). This model building is rather time consuming and can be regarded in itself as full of decisions and workflow branches as the process leading to the map that has been described in this chapter. The parameters sought in this case are the spatial location of each one of the atoms in the macromolecule. Scipion is expanding in



**Fig. 22** Isosurface representation of the virus capsid (top) and a zoomed version of it (bottom). The whole capsid has been represented by a gray mesh, while the unit cell of the virus (and a small region around it) has been highlighted in colors according to its distance to the virus center



**Fig. 23** Detail of the fitting process in which the EM map is used as a spatial constraint to construct an atomic model of the macromolecule

that domain, and it has incorporated some of the most common algorithmic tools used for this task [61]. The interested reader is referred to this latter publication and to the extensive tutorial about model building at <https://scipion-em.github.io/docs/docs/user/user-documentation.html#tutorials> for expanding this information.

---

### 3 Conclusions

There is an old saying in Electron Microscopy that goes like “the good thing about 3D Electron Microscopy is that you always get a volume; the bad thing about 3D Electron Microscopy is that you always get a volume.” It refers to the fact that by acquiring some images at the microscope and by processing them we will always get a map that pretends to be a faithful model compatible with our measurements. Our success in effectively achieving this goal depends on our ability to choose good micrographs preserving structural information to high resolution, particles that really correspond to a single structure, and being able to find the relative orientations of these particles in space. Making incorrect decisions in any of these steps will necessarily degrade the quality of the reconstructed map. In the limit, bad data quality, population mixtures, or incorrect alignments may yield the reconstructed map useless. Validation tools can be regarded as sanity checks that good maps satisfy. However, meeting the requirements of the

validation tools do not automatically make our reconstruction good. But, on the contrary, failing to meet these requirements, almost surely raises a warning on the quality of our reconstruction.

The EM old saying can be actually generalized to a wider problem: parameter estimation (“the good thing about parameter estimation is that you always get a value; the bad thing about parameter estimation is that you always get a value”). The whole process of image processing in EM is about finding parameters (whether a micrograph preserves information at high resolution or not, whether this small image is a centered particle or not, whether this particle belongs to this population or not, which is the 3D pose of the particle with respect to this volume. The old EM saying is just a particularization of this more general principle to the parameters describing the Coulomb potential in space of the macromolecule under study. All these parameters are estimated by computational algorithms that, by definition, will have false positives, false negatives, biases, and variances associated to the parameter estimation. This task is further complicated by the fact that the Signal-to-Noise Ratio (SNR) of the measurements is between  $1/10$  and  $1/100$  (i.e., there is 10–100 times more noise than signal) if we talk about micrographs, or several hundredths of this if we talk about movie frames. In an experimental study, we can never be sure that we have correctly identified all the parameters as the true values are unknown. The most we can do is to estimate the same parameter multiple times (preferably with different algorithms based on completely “orthogonal” mathematical principles), and only trust those estimates that are consistently estimated to similar values. This principle does not preclude bias (the different estimates may consistently point to a wrong answer), but the probability of this event is smaller if multiple estimates point to the same value, especially in the case of the use of several algorithms rather than multiple runs of the same algorithm. In this task, having a platform like Scipion that integrates over 30 packages with more than 300 protocols solving high-level image-processing tasks is not only convenient, but an imperative requirement if we are to combine the results from multiple runs and algorithms without having to deal with all the internal details and conventions of the different software packages.

At present, there seems to be an obsession in reporting structures with a resolution number as low as possible. However, it has been already shown [10, 51] and illustrated in this chapter that the resolution reported by the FSC (even if it is coming from two independent halves of the data) is at the low end of the local resolution range, and consequently, it is not the resolution of the map, but the resolution of the best resolved voxels in the map. Moreover, the FSC can be inadvertently distorted by the choice of the mask used for its calculation. Map-model comparisons [62] are alternative, objective measurements of the quality of the map.

Although they have the drawback that they also depend on the quality of the fitting, so that a good map can be poorly evaluated due to a poor quality fitting. Probably, there does not exist a single measure that can unambiguously determine the quality of a map, and all figures of merit can be fooled by pathological features. Our opinion is that reporting measures like the ones shown along the chapter, specifically targeting the identification of incorrectly estimated parameters, should be promoted as a good practice within the structural biology community. Relying on the algorithm to correctly deal with the miss-estimates is a dangerous practice that should be avoided in a solid scientific work.

Single Particle Analysis by EM has witnessed in the last years an incredible boost in its throughput and in the range of interesting biological problems addressed. In a way, this boost has been promoted by hardware improvements as well as faster, more robust, and easy-to-use algorithms. However, no algorithm is free from errors, particularly so because at this low level of SNR currently the correctness of many estimated parameters is very difficult to assess. As a field, we have given a huge step forward in reproducibility and open science by making public first the reconstructed maps (EMDB, [63]) and more recently the raw data (EMPIAR, [64]). We foresee a near future in which the intermediate steps connecting the raw data and the final map are also made public in such a way that the correctness of the estimated parameters can be verified, and even the data could be reprocessed at those points where the first data analysis is suspected to produce suboptimal parameter estimates.

---

## Acknowledgments

The authors would like to acknowledge economical support from: The Spanish Ministry of Economy and Competitiveness through Grants BIO2016-76400-R(AEI/FEDER, UE), the “Comunidad Autónoma de Madrid” through Grant: S2017/BMD-3817. Instituto de Salud Carlos III through Grant: PT17/0009/0010 (ISCIII-SGEFI / ERDF). European Union (EU) and Horizon 2020 through grants: CORBEL (INFRADEV-1-2014-1, Proposal: 654248) Instruct ULTRA (Proposal: 731005), EOSC Life (Proposal: 824087), HighResCells (Proposal: 810057), IMpaCT (Proposal: 857203), EOSC—Synergy (Proposal: 857647), iNEXT-Discovery (Proposal: 871037), and European Regional Development Fund-Project “CERIT Scientific Cloud” (No. CZ.02.1.01/0.0/0.0/16\_013/0001802). The authors acknowledge the support and the use of resources of Instruct, a Landmark ESFRI project.

## References

1. Benjin X, Ling L (2020) Developments, applications, and prospects of cryo-electron microscopy. *Protein Sci* 29:872–882
2. Lyumkis D (2019) Challenges and opportunities in cryo-em single-particle analysis. *J Biol Chem* 294:5181–5197
3. Eisenstein M (2018) Drug designers embrace cryo-EM. *Nat Biotechnol* 36:557–558
4. Scapin G, Potter CS, Carragher B (2018) Cryo-em for small molecules discovery, design, understanding, and application. *Cell Chem Biol* 25:1318–1325
5. Saur M, Hartshorn MJ, Dong J, Reeks J et al (2019) Fragment-based drug discovery using cryo-em Drug discovery today. doi: <https://doi.org/10.1016/j.drudis.2019.12.006>
6. Jonic S (2017) Computational methods for analyzing conformational variability of macromolecular complexes from cryo-electron microscopy images. *Curr Opin Struct Biol* 43:114–121
7. Sorzano COS, Jiménez A, Mota J, Vilas JL et al (2019) Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. *Acta Crystallogr Sect F, Struct Biol Commun* 75:19–32
8. Arnold SA, Müller SA, Schmidli C et al (2018) Miniaturizing EM sample preparation: opportunities, challenges, and “visual proteomics”. *Proteomics* 18:e1700176
9. Faruqi AR, McMullan G (2018) Direct imaging detectors for electron microscopy. *Nucl Instrum Methods Phys Res, Sect A* 878:180–190
10. Vilas JL, Gómez-Blanco J, Conesa P et al (2018) MonoRes: automatic and unbiased estimation of local resolution for electron microscopy maps. *Structure* 26:337–344
11. de la Rosa-Trevín JM, Quintana A, Del Cano L et al (2016) Scipion: a software framework toward integration, reproducibility and validation in 3D electron microscopy. *J Struct Biol* 195:93–99
12. Wang Z, Hryc CF, Bammer B et al (2014) An atomic model of bromo mosaic virus using direct electron detection and real-space optimization. *Nat Commun* 5:4808
13. Heymann JB, Marabini R, Kazemi M et al (2018) The first single particle analysis map challenge: a summary of the assessments. *J Struct Biol* 204:291–300
14. Sorzano COS, Vargas J, Oton J et al (2017) A review of resolution measures and related aspects in 3D electron microscopy. *Prog Biophys Mol Biol* 124:1–30
15. Vilas JL, Tagare HD, Vargas J et al (2020) Measuring local-directional resolution and local anisotropy in cryo-EM maps. *Nat Commun* 11:55
16. Ramírez-Aportela E, Mota J, Conesa P et al (2019) Deep-res: a new deep-learning- and aspect-based local resolution method for electron-microscopy maps. *IUCRj* 6:1054–1063
17. Sorzano COS, Fernández-Giménez E, Peredo-Robinson V et al (2018) Blind estimation of DED camera gain in electron microscopy. *J Struct Biol* 203:90–93
18. Li X, Mooney P, Zheng S, Booth CR et al (2013) Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-EM. *Nat Methods* 10:584–590
19. Zheng SQ, Palovcak E, Armache JP et al (2017) Motion-cor2: anisotropic correction of beam-induced motion for improved cryo-electron microscopy. *Nat Methods* 14:331–332
20. Abrishami V, Vargas J, Li X, Cheng Y et al (2015) Alignment of direct detection device micrographs using a robust optical flow approach. *J Struct Biol* 189:163–176
21. Tegunov D, Cramer P (2019) Real-time cryo-electron microscopy data preprocessing with warp. *Nat Methods* 16:1146–1152
22. de la Rosa-Trevín JM, Otón J, Marabini R et al (2013) Xmipp 3.0: an improved software suite for image processing in electron microscopy. *J Struct Biol* 184(2):321–328
23. Sorzano COS, Jonic S, Núñez Ramírez R et al (2007) Fast, robust and accurate determination of transmission electron microscopy contrast transfer function. *J Struct Biol* 160:249–262
24. Zhang K (2016) Gctf: real-time ctf determination and correction. *J Struct Biol* 193:1–12
25. Rohou A, Grigorieff N (2015) Ctfind4: fast and accurate defocus estimation from electron micrographs. *J Struct Biol* 192:216–221
26. Maluenda D, Majtner T, Horvath P et al (2019) Flexible workflows for on-the-fly electron-microscopy single-particle image processing using scipion. *Acta Crystallogr Sect D, Struct Biol* 75:882–894
27. Marabini R, Carragher B, Chen S et al (2015) Ctf challenge: result summary. *J Struct Biol* 190:348–359
28. Voss NR, Yoshioka CK, Radermacher M et al (2009) Dog picker and tiltpicker: software tools to facilitate particle selection in single

- particle electron microscopy. *J Struct Biol* 166 (2):205–213
29. Scheres SHW (2015) Semi-automated selection of cryo-em particles in relion-1.3. *J Struct Biol* 189:114–122
  30. Abrishami V, Zaldívar-Peraza A, de la Rosa-Trevín JM et al (2013) A pattern matching approach to the automatic selection of particles from low-contrast electron micrographs. *Bioinformatics* 29:2460–2468
  31. Bepler T, Morin A, Rapp M, Brasch J et al (2019) Topaz: a positive-unlabeled convolutional neural network cryoem particle picker that can pick any size and shape particle. *Microsc Microanal* 25:986–987
  32. Wagner T, Merino F, Stabrin M et al (2019) Sphire-cryolo is a fast and accurate fully automated particle picker for cryo-EM. *Commun Biol* 2:218
  33. Sanchez-Garcia R, Segura J, Maluenda D et al (2018) Deep consensus, a deep learning-based approach for particle pruning in cryo-electron microscopy. *IUCrJ* 5:854–865
  34. Sánchez-García R, Segura J, Maluenda D et al (2020) Micrograph cleaner: a python package for cryo-EM micrograph cleaning using deep learning. *bioRxiv*. <https://doi.org/10.1101/677542>
  35. Vargas J, Abrishami V, Marabini R et al (2013) Particle quality assessment and sorting for automatic and semiautomatic particle-picking techniques. *J Struct Biol* 183:342–353
  36. Punjani A, Brubaker MA, Fleet DJ (2017) Building proteins in a day: efficient 3D molecular structure estimation with electron cryomicroscopy. *IEEE Trans Pattern Anal Mach Intell* 39:706–718
  37. Sorzano COS, Bilbao-Castro JR, Shkolnisky Y et al (2010) A clustering approach to multi-reference alignment of single-particle projections in electron microscopy. *J Struct Biol* 171:197–206
  38. Sorzano COS, Vargas J, de la Rosa-Trevín JM et al (2014) Outlier detection for single particle analysis in electron microscopy. In: *Proc. Intl. Work-Conference on Bioinformatics and Biomedical Engineering, IWBBIO*, p 950
  39. Vargas J, Álvarez-Cabrera AL, Marabini R et al (2014) Efficient initial volume determination from electron microscopy images of single particles. *Bioinformatics* 30:2891–2898
  40. Sorzano COS, Vargas J, de la Rosa-Trevín JM et al (2015) A statistical approach to the initial volume problem in single particle analysis by electron microscopy. *J Struct Biol* 189:213–219
  41. Scheres SHW (2012) Relion: implementation of a Bayesian approach to cryo-EM structure determination. *J Struct Biol* 180:519–530
  42. Tang G, Peng L, Baldwin PR, Mann DS et al (2007) Eman2: an extensible image processing suite for electron microscopy. *J Struct Biol* 157:38–46
  43. Reboul CF, Eager M, Elmlund D, Elmlund H (2018) Single-particle cryo-EM- improved *ab initio* 3D reconstruction with simple/prime. *Protein Sci* 27:51–61
  44. Sorzano COS, Vargas J, Vilas JL et al (2018) Swarm optimization as a consensus technique for electron microscopy initial volume. *Appl Anal Optim* 2:299–313
  45. Gomez-Blanco J, Kaur S, Ortega J, Vargas J (2019) A robust approach to ab initio cryo-electron microscopy initial volume determination. *J Struct Biol* 208:107397
  46. Jimenez A, Jonic S, Majtner T et al (2019) Validation of electron microscopy initial models via small angle x-ray scattering curves. *Bioinformatics* 35:2427–2433
  47. Kimanius D, Forsberg BO, Scheres SH, Lindahl E (2016) Accelerated cryo-EM structure determination with parallelisation using GPUs in RELION-2. *elife* 5:e18722
  48. Sorzano COS, Marabini R, Vargas J et al (2014) *Computational Methods for Three-Dimensional Microscopy Reconstruction*, Springer, chap Interchanging geometry information in electron microscopy single particle analysis: mathematical context for the development of a standard, pp 7–42
  49. Cardone G, Heymann JB, Steven AC (2013) One number does not fit all: mapping local variations in resolution in cryo-em reconstructions. *J Struct Biol* 184:226–236
  50. Kucukelbir A, Sigworth FJ, Tagare HD (2014) Quantifying the local resolution of cryo-EM density maps. *Nat Methods* 11:63–65
  51. Ramírez-Aportela E, Vilas JL, Glukhova A et al (2019) Automatic local resolution-based sharpening of cryo-EM maps. *Bioinformatics* 36:765–772
  52. Fernández JJ, Luque D, Castón JR, Carrascosa JL (2008) Sharpening high resolution information in single particle electron cryomicroscopy. *J Struct Biol* 164(1):170–175
  53. Vilas JL, Vargas J, Martínez M et al (2020b) Re-examining the spectra of macromolecules: current practice of spectral quasi b-factor flattening. *J Struct Biol* 209:107447
  54. Jakobi AJ, Wilmanns M, Sachse C (2017) Model-based local density sharpening of cryo-EM maps. *elife* 6:e27131



55. Naydenova K, Russo CJ (2017) Measuring the effects of particle orientation to improve the efficiency of electron cryomicroscopy. *Nat Commun* 8:629
56. Vargas J, Otón J, Marabini R et al (2016) Particle alignment reliability in single particle electron cryomicroscopy: a general approach. *Sci Rep* 6:21626
57. Vargas J, Melero R, Gómez-Blanco J et al (2017) Quantitative analysis of 3D alignment quality: its impact on soft-validation, particle pruning and homogeneity analysis. *Sci Rep* 7:6307
58. Stagg SM, Noble AJ, Spilman M, Chapman MS (2014) Reslog plots as an empirical metric of the quality of cryo-em reconstructions. *J Struct Biol* 185:418–426
59. Heymann B (2015) Validation of 3dem reconstructions: the phantom in the noise. *AIMS Biophys* 2:21–35
60. Beckers M, Jakobi AJ, Sachse C (2019) Thresholding of cryo-em density maps by false discovery rate control. *IUCrJ* 6(1):18–33
61. Martínez M, Jiménez-Moreno A, Maluenda D et al (2020) Integration of cryo-EM model building software in Scipion. *J Chem Inf Model* 26:2533–2540
62. Afonine PV, Klaholz BP, Moriarty NW et al (2018) New tools for the analysis and validation of cryo-em maps and atomic models. *Acta Crystallogr Sect D, Struct Biol* 74:814–840
63. Patwardhan A (2017) Trends in the electron microscopy data bank (emdb). *Acta Crystallogr Sect D: Struct Biol* 73:503–508
64. Iudin A, Korir PK, Salavert-Torres J et al (2016) Empiar: a public archive for raw electron microscopy image data. *Nat Methods* 13:387–388