Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Ultramicroscopy 102 (2004) 67–84
ultramicroscopy
www.elsevier.com/locate/ultramic

Noise bias in the refinement of structures derived from single particles

Alex Stewart, Nikolaus Grigorieff*

Howard Hughes Medical Institute and Department of Biochemistry, Rosenstiel Basic Medical Sciences Research Center, MS029, Brandeis University, 415 South Street, Waltham, MA 02454-9110, USA

Received 1 May 2004; received in revised form 14 August 2004; accepted 24 August 2004

Abstract

One of the main goals in the determination of three-dimensional macromolecular structures from electron microscope images of individual molecules and complexes (single particles) is a sufficiently high spatial resolution, about 4 Å, at which the interpretation with an atomic model becomes possible. To reach high resolution, an iterative refinement procedure using an expectation maximization algorithm is often used that leads to a more accurate alignment of the positional and orientational parameters for each particle. We show here the results of refinement algorithms that use a phase residual, a linear correlation coefficient, or a weighted correlation coefficient to align individual particles. The algorithms were applied to computer-generated data sets that contained projections from model structures, as well as noise. The algorithms show different degrees of over-fitting, especially at high resolution where the signal is weak. We demonstrate that the degree of over-fitting is reduced with a weighting scheme that depends on the signal-to-noise ratio in the data. The weighting also improves the accuracy of resolution measurement by the commonly used Fourier shell correlation. The performance of the refinement algorithms is compared to that using a maximum likelihood approach. The weighted correlation coefficient was implemented in the computer program FREALIGN.

PACS: 87.64.Bx; 87.64.Dz

Keywords: Image processing; Electron microscopy; Protein structure; FREALIGN

1. Introduction

Electron microscopic imaging of isolated (single) macromolecules and their assemblies has become one of the standard methods to determine their three-dimensional (3D) shape.

*Corresponding author. Tel.: +1 781 736 2444; fax: +1 781 736 2405.
E-mail address: niko@brandeis.edu (N. Grigorieff).

0304-3991/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.ultramic.2004.08.008

Such information has become essential for a more detailed understanding of the many molecular machines that perform complex functions inside living cells [1]. Single-particle electron microscopy, as it is often called, can be used to visualize the structure of macromolecular complexes that are not readily accessible by more traditional techniques, such as nuclear magnetic resonance (NMR) spectroscopy or X-ray crystallography: It does not require crystals, it does not impose an upper molecular mass limit on the structure under investigation, and it can be applied to minute amounts of only a few pmol of material (see, for example, Refs. [2,3]). The main shortcoming of single particle electron microscopy is its limited resolution which does not normally allow the interpretation of the structure by an atomic model. For example, highly symmetrical viruses can currently be resolved to about 7 Å resolution [4–6] while the asymmetrical ribosome has been resolved to about 10 Å resolution [7,8]. Furthermore, a lower molecular mass limit of about 50 kDa exists [9] because a minimum mass is required to generate sufficient scattering contrast for the alignment of images of the individual particles.

As the push to higher resolution has intensified over the last 10 years with the development of more powerful electron microscopes, faster computers used for image processing, and improved image-processing algorithms, an unbiased refinement of a 3D reconstruction, and a reliable assessment of its resolution have become more important. Both the refined structure and its measured resolution can be subject to significant errors, due to over-fitting of the 3D reconstruction to the often very noisy image data [10]. This paper will introduce a new strategy for the refinement of structures derived from images of single particles, with the goal of reducing noise-dependent over-fitting and accurate resolution measurement. We start with a brief review of a current typical refinement procedure and the source of over-fitting.

2. Expectation maximization algorithm

For the refinement of a structure determined from images of single particles, we assume that we have already a rough structure of the particle that was determined using, for example, the random conical tilt [11] or angular reconstitution method [12]. The goal of the refinement is then to obtain more accurate parameters that describe the x, y position of the particle in the image, and its 3 angles of orientation. A typical refinement procedure will use the rough 3D structure as a reference to re-align all particles and obtain improved particle parameters that can be used to calculate a better 3D structure. Such an approach is implemented, for example, in the projection matching procedure [13] or the minimization of a phase residual in reciprocal space between particle image and reference [14]. The iteration cycle of alignment and reconstruction is a typical implementation of the expectation maximization algorithm [15,16].

The expectation maximization family of algorithms is a prescription for incrementally improving a model based on the current state of knowledge about the model and data. The process proposes finding a "complete" data set with which one could produce the best possible model without any additional knowledge. In our case this "complete" data set includes both the individual particle images, and the particle orientations which we do not know. The expectation maximization algorithm estimates the hidden data (the orientation parameters for the particles) using the current model and the image data. The algorithm cycles back and forth, iterating between predicting the model structure given the complete data, and estimating the particle orientation parameters. In order for this process to lead to an improved model (as measured by an increased likelihood of the data), we must be sure that the alignment function correctly weights the information contained in the images to produce a strictly better set of particle orientations. Previous implementations have used a variety of correlation style functions to align the particle images. We propose to use a correlation function that is tuned such that given a "true" model of the structure, the error in the estimated orientations will be minimal in a least squares sense. This requires weighting down those signal components with low signal-to-noise ratio (SNR).

It has been shown that the convergence of a correlation-based refinement algorithm onto the correct structure is inferior to a maximum likelihood approach at very small SNR [17]. In the direct maximum likelihood approach, the assignment of parameters to each particle in an image is avoided, and the refined structure contains each image at all possible alignment parameters, weighted by the probability of the structure given any particular set of alignment parameters. The calculation of the refined structure is computationally expensive and an implementation of the algorithm has so far only been demonstrated in two dimensions (2D, [17]). Our goal in this paper is, therefore, to formulate a weighting scheme that improves the performance of the correlation-based refinement. It will be shown below that the use of a carefully weighted correlation coefficient performs significantly better than a simple linear correlation coefficient, approaching the performance of a maximum likelihood refinement.

3. Resolution measurement

At any point in the refinement, the resolution of a structure can be estimated from features in the structure if these features correspond to elements of known structure. For example, the visualization of α-helical secondary structure indicates that a resolution of about 9 Å has been achieved [18]. At a resolution that does not correspond to a threshold at which known features become visible, it is more difficult to assess the resolution accurately. The most commonly used resolution measure is the Fourier shell correlation (FSC, [19]). It evaluates the correlation in resolution zones between two 3D reconstructions of a particle that are calculated from two halves of the data. The assumption in the analysis is that the noise in these two test reconstructions that remains after alignment of the particles and averaging over each half data set is not correlated. However, in practice, this assumption is often not valid because, as described for the expectation maximization algorithm above, the particles from both half data sets are usually aligned to a single, overall reference 3D reconstruction that is derived from a previous alignment cycle. This leads to a bias in the two reconstructions that correlates the noise in the test reconstructions, and the FSC becomes unreliable, usually indicating significantly higher correlation than would be expected for truly uncorrelated noise [10].

We evaluate here the degree of bias in 3D reconstructions induced by the reference structure, for different similarity measures used in the alignment. We illustrate that a reduced bias can be achieved by careful weighting of the data according to its SNR, and that the reduced bias renders the resolution measurement by the FSC method more reliable. We show that the reduced bias significantly improves the overall particle alignment and resolution of the reconstruction. The study uses computer-simulated data sets for which the underlying true structure is known a priori. As similarity measures, we consider a weighted phase residual, the linear (unweighted) correlation coefficient, as well as a newly devised weighted correlation coefficient. Finally, we compare the performance of the new weighted correlation coefficient with that of a refinement using a direct maximum likelihood algorithm.

4. Weighted correlation coefficient

A widely used similarity measure between images that will be considered here is the linear correlation coefficient, CC, defined as

$$\mathrm{CC}(X, Y) = \frac{\sum_{i=1}^{N\times N} x_i y_i}{\sqrt{\sum_{i=1}^{N\times N} x_i^2 \sum_{i=1}^{N\times N} y_i^2}}. \quad (1)$$

X and Y are a particle image and a reference image of dimensions $N \times N$, and with pixel values x_i and y_i, respectively. In Eq. (1) we assume that the pixel averages of x_i and y_i are zero, according to a normalization procedure where a constant is added to each pixel value, followed by a multiplication of a second constant, to yield images with a zero average and constant variance across the data set. An alignment procedure will attempt to maximize CC by applying rotations and translations either to the particle image, or to the reference image. If the reference is a 3D

reconstruction, the reference image is a projection of this reconstruction, and the alignment procedure will vary the three angles defining the projection, as well as the two positional coordinates, to maximize CC. Eq. (1) can also be written in Fourier space:

$$\mathrm{CC}(X,Y)=\frac{\sum_{k\in[0,0.5]}F_X(\mathbf{k})F_Y^{\dagger}(\mathbf{k})}{\sqrt{\sum_{k\in[0,0.5]}\left|F_X(\mathbf{k})\right|^2\sum_{k\in[0,0.5]}\left|F_Y(\mathbf{k})\right|^2}}. \tag{2}$$

Here, F_X and F_Y are the Fourier transformations of images X and Y, respectively, and $\mathbf{k}$ is the reciprocal space coordinate. The magnitude of $\mathbf{k}$, k, indicates resolution and assumes values between 0 and 0.5, the Nyquist frequency. In the alignment of images of single particles, the range of k is often limited to exclude very low frequencies that could contain strong components due to a density gradient across the image that are unrelated to the particle, as well as high frequencies that are very noisy. The linear correlation coefficient contains an implicit linear weighting by the amplitudes of both images: Fourier terms with large amplitudes will determine the outcome of the sum in the numerator in Eq. (2) more strongly than terms with small amplitudes.

We will further consider a weighted phase residual that has been used as a similarity measure in the computer program FREALIGN [14], used for the refinement of single particle 3D reconstructions:

$$\mathrm{PRES}(X,Y)=\frac{\sum_{k\in[0,0.5]}\Delta\Phi_{X,Y}(\mathbf{k})\left|F_X(\mathbf{k})\right|}{\sum_{k\in[0,0.5]}\left|F_X(\mathbf{k})\right|}. \tag{3}$$

$\Delta\Phi_{X,Y}$ is the phase difference between Fourier terms of the Fourier transformations of images X and Y. In Eq. (3), the weight for the phase residual contains only the amplitudes of the particle image, and not also the amplitudes of the reference image, to avoid over weighting of strong frequency components in the data [14,20].

The amplitude weighting in the linear correlation coefficient and the weighted phase residual is based on the assumption that a large amplitude of a Fourier term is indicative of a strong signal. Such an assumptions holds usually true, for example, in X-ray or electron diffraction experiments where the amplitude indicates how many photons or electrons have been measured, respectively. When images are considered, however, the overall SNR is very small, for example, $\frac{1}{8}$ for a particle of 80 kDa molecular weight [21], due to the low electron dose used to minimize damage to the specimen. For larger macromolecular complexes, the overall SNR is higher but still becomes very small towards high resolution. Under conditions of weak signal, the above assumption that the amplitude of a Fourier term indicates the strength of the signal is invalid because the amplitudes are dominated by random noise.

Ideally, the weights used in a similarity measure would be based on the true signal contained in the particle image and the reference. A fundamental problem at small SNR values is therefore the accurate determination of the true SNR. A very sensitive SNR measurement has been proposed that derives from the correlation coefficient between two images [22]. We propose here a new similarity measure that uses weights derived from correlation coefficients.

To introduce a more accurate weighting scheme, we consider first an alignment procedure that performs a least-squares fit of the reference to the particle image. We seek to minimize

$$L=\sum_{k\in[0,0.5]}\left|F_X(\mathbf{k})-F_Y(\mathbf{k})\right|^2. \tag{4}$$

In the following, we will assume that the reference contains no noise. This assumption is not valid, in practice, since the reference is usually derived from the noisy particle images and, therefore, also contains noise. The effect of noise in the reference will be considered later. The particle image can be written as

$$F_X(\mathbf{k})=F_Y(\mathbf{k})+N_X(\mathbf{k}). \tag{5}$$

N_X is the difference between the particle image and the reference. If these are not aligned, the difference will arise from the misalignment, as well as the noise in the particle image. After approximate alignment, however, the main difference will be due to the noise.

Eq. (4) does not contain any weighting that is based on the SNR of the particle image, SNR_X. Such a weighting can be introduced through a Wiener filter W:

$$L_W = \sum_{k\in[0,0.5]} \left|F_X(\mathbf{k})W_X(\mathbf{k}) - F_Y(\mathbf{k})\right|^2 \tag{6}$$

with

$$W_X(\mathbf{k}) = \frac{1}{\left(1 + 1/\mathrm{SNR}_X(\mathbf{k})\right)} \tag{7}$$

and

$$\mathrm{SNR}_X(\mathbf{k}) = \frac{\left|F_Y(\mathbf{k})\right|^2}{\left|N_X(\mathbf{k})\right|^2}. \tag{8}$$

The SNR of individual Fourier terms is difficult to measure. Frank and Al-Ali [22] showed, however, that the correlation coefficient between two images is a sensitive measure of the SNR:

$$\mathrm{SNR} = \frac{\mathrm{CC}}{1-\mathrm{CC}}. \tag{9}$$

Here, CC is the correlation coefficient between two images that have the same SNR. Since we assume for the moment that our reference image does not contain noise, this formula becomes (see Appendix A)

$$\mathrm{SNR} = \frac{\mathrm{CC}^2}{1-\mathrm{CC}^2}. \tag{10}$$

The average SNR of a group of Fourier terms can, therefore, be estimated from the correlation coefficient between these terms and the reference image. The SNR can vary significantly across the Fourier spectrum and depends on the structure we are to determine. It has its strongest dependence on the resolution, however, and it is therefore reasonable to evaluate average SNR values for resolution zones.

To relate Eq. (6) to a correlation coefficient, we note that, if both particle image and reference image are normalized as described above, the minimization of L_W is equivalent to the maximization of

$$\begin{aligned} R_W &= \sum_{k\in[0,0.5]} F_X(\mathbf{k})W_X(\mathbf{k})F_Y(\mathbf{k}) \\ &\approx \sum_i W_i \sum_{k\in[k_i,k_i+\Delta k]} F_X(\mathbf{k})F_Y(\mathbf{k}) \\ &= \sum_i W_i \mathrm{CC}_i \sqrt{\sum_{k\in[k_i,k_i+\Delta k]} \left|F_X(\mathbf{k})\right|^2 \sum_{k\in[k_i,k_i+\Delta k]} \left|F_Y(\mathbf{k})\right|^2} \\ &= \sum_i \mathrm{CC}_i^3 \sqrt{\sum_{k\in[k_i,k_i+\Delta k]} \left|F_X(\mathbf{k})\right|^2 \sum_{k\in[k_i,k_i+\Delta k]} \left|F_Y(\mathbf{k})\right|^2}. \end{aligned} \tag{11}$$

CC_i is the average correlation coefficient in the resolution zone $[k_i, k_i+\Delta k]$. W_i is the Wiener filter constant calculated using the average SNR in this zone. The index i runs over all resolution zones considered in the alignment. For the last equality in Eq. (11), we used the relation

$$W_i = \mathrm{CC}_i^2 \tag{12}$$

that follows from Eqs. (7) and (10). In Eq. (11), the square root depends on the amplitudes of the Fourier terms, as well as the number of terms in the sums. As discussed before, an amplitude weighting does not properly reflect the signal contained in the data. We therefore ignore the amplitude contribution to the weighting and write for the function that needs to be maximized in an alignment

$$R = \frac{1}{\sum_i N_i} \sum_i N_i \mathrm{CC}_i^3. \tag{13}$$

N_i is the number of Fourier terms in the resolution zone $[k_i, k_i+\Delta k]$.

The correlation coefficients CC_i in Eq. (13) are themselves subject to significant error. An image with $N \times N$ pixels that contains solely noise will have an expected correlation coefficient with a reference that is approximately Gaussian-distributed around zero with a standard deviation of

$$w = \frac{1}{N}. \tag{14}$$

To reduce the dependence of the alignment algorithm on the small and noisy correlation coefficients, we use the standard deviation in Eq. (14) as a filter constant and replace the correlation coefficient in Eq. (13) by a filtered

correlation coefficient

$$R_{\mathrm{f}} = \frac{1}{\sum_i N_i} \sum_i N_i \mathrm{CC}_{\mathrm{f},i}^3 \tag{15}$$

with

$$\mathrm{CC}_{\mathrm{f},i} = \frac{\mathrm{CC}_i^2}{\mathrm{CC}_i + w}. \tag{16}$$

Although for many biological single-particle imaging projects the resolution is clearly the major factor in determining the SNR of the Fourier data, there are important classes of structures for which this is not the case. Helical structures, for example, produce concentrated energy in layer lines which dominate the transforms of images of Actin, Myosin, Tubulin and many other biopolymers. Icosahedral viruses have strong spot-like reflections from their projected structure. In these cases, our resolution zone processing of the data is non-ideal, but estimating the SNR of strong localized signals can be done by scanning through the data in other ways. Strong signal generally can look out for itself. It is most essential to correctly weight the weak data, particularly at high resolution.

It is worth pointing out that the correlation coefficient in Eq. (10) between particle image and reference is only related to the SNR if particle image and reference are perfectly aligned. For a misaligned particle, the correlation coefficient will be smaller than expected for a particular SNR, and therefore, Eq. (10) will give an underestimate of the true SNR in the image. For the purpose of finding the best alignment by maximizing R in Eq. (13), however, the correlation coefficient still provides a useful measurement of the goodness of fit, and the SNR calculated in Eq. (10) can be taken as an "effective" SNR that measures only the signal that is aligned with the reference, while the remaining mismatched signal, together with the noise, contributes to the difference between image and reference.

In the derivations above, we assumed that the reference contains no noise. In the more realistic scenario of a noisy reference, the correlation coefficient in Eq. (10) will be affected by the noise in the particle image, as well as the reference. Similarly to the case of a misaligned particle, the result will be an underestimated SNR in the image. The error in the SNR estimate will be greater at high resolution than at lower resolution, and it will lead to smaller weights W_i. Since at small SNR values, the estimate of the SNR will be relatively noisy, this error is unlikely to have a significant effect on the overall alignment process.

5. Uniform convergence

The use of the similarity measure given by Eq. (15) assigns each resolution zone in the data an appropriate weight to achieve an alignment with reduced bias. Since the refinement of a structure determined from images of individual particles usually requires several iterations of alignment and 3D reconstruction, we also have to consider the convergence behavior of our algorithm. Many existing refinement algorithms use the reconstruction of a particle from an earlier refinement cycle as the reference for the next iteration. Such algorithms often converge onto local optima, preventing further alignment that would equally satisfy terms in all resolution zones. In our refinement algorithm, we consider the difference between the current model structure and the new reconstruction given the updated parameters. This difference should be interpreted as a direction of a step or gradient to be applied to the earlier model to reach an improved structure (as measured by the function discussed above).

One important observation concerning the Fourier representation of the gradient is the following: A small change in particle orientation or position may lead to only a tiny change in the low frequency components of the gradient, but will cause a much larger change in the high-frequency components. Because of this, high frequencies will tend to converge more rapidly than low frequencies and can lock the system into a local optimum based on the high-frequency components of the images, which in general have little signal and plenty of noise. This causes instability upon iteration, and is a consequence of the following consideration. A gradient should not generally be directly applied to a model, since it is a covariant entity. The components of different spatial frequency need to be scaled so that they have uniform

influence on the model. To apply the gradient to the model in such a way that the convergence is uniform in all components, independent of spatial frequency, we must filter the gradient by the inverse spatial frequency. Such a filter does not guarantee a convergence onto the global optimum but onto a point that lies closer to it compared to a refinement without a filter.

Note that obtaining the model step by filtering the gradient does not change the stationary (convergence) point of the iteration, only the path taken. Without the filter the algorithm loses the ability to reject spurious high-resolution detail. Other techniques for managing the convergence might be employed such as gradually increasing the resolution between iterations or regularization with smoothness constraints. The advantage of the proposed method is that it is data-driven and does not require manual tuning or interference. The use of a step to modify the model allows us to consider many possible additional constraints that are not discussed here.

As shown below, the step filtering enables the refinement algorithm to deal with very low SNR. However, the convergence of the algorithm, especially at the high-resolution end of the structure, is slow (see the comparison with maximum likelihood refinement). We investigated, therefore, also a single-step refinement algorithm that does not use step-filtering but uses an additional weight in the refinement function that is proportional to the inverse spatial frequency to reduce non-uniform convergence. Thus, the weighted correlation coefficient in Eq. (15) is further modified to

$$R'_{\mathrm{f}} = \sum_i \mathrm{CC}^3_{\mathrm{f},i} \tag{17}$$

Although the single-step refinement using Eq. (17) is not optimal for data with very low SNR (e.g. $\frac{1}{200}$), it performs remarkably well at SNR values commonly encountered in experimental data (e.g. $\frac{1}{50}$).

6. Constraints

Over-fitting of noise is a common problem in algorithms that attempt to fit a model to a set of measurements. The degree of over-fitting can be controlled with constraints that introduce a so-called regularization to the fitting problem. In the reconstruction of 3D volumes from projections, non-linear constraints such as maximum entropy [23] or discrete density [24] are often used. The maximum entropy constraint attempts to balance the information content of the model against goodness-of-fit, while discrete density constraints force the range of densities within the model to agree with prior knowledge of sample composition (such as the simple condition that the density of the model be positive). Non-linear constraints can help minimize over-fitting and can be used to restrict the space of possible models in physically reasonable ways. However, in this paper we concentrate on the form and behavior of the function evaluating the fit between model and data. The behavior of these functions is best observed in the absence of non-linear constraints.

We introduce here a new type of constraint that is based on the principle difference in the character of noise and signal. Besides being random, noise in an image also lacks "coherence": The density of any point in an image is not correlated with the density of any other point in the image. The same is true in Fourier space. Signal, on the other hand, has coherence. This means that alignment of one part of the signal image will also align all other parts. In Fourier space, the alignment of signal at low resolution will automatically align the signal at high resolution. The signal in a perfectly aligned image will correlate positively across the entire resolution range, whereas the alignment of noise in one resolution zone will not lead to any bias of correlation towards positive or negative values in another resolution zone. Alignment algorithms commonly optimize an overall correlation coefficient, thus biasing the correlation towards positive values across the entire resolution range.

As explained above, this leads to a bias in the 3D reconstruction which is amplified by further iterations. To reduce this bias and amplification, we modify the alignment goal by optimizing only the absolute value of the correlation coefficient in each resolution zone. The target function given by

Eq. (15) is replaced by

$$\tilde{R}_{\mathrm{f}} = \frac{1}{\sum_i N_i} \sum_i N_i \left| \mathrm{CC}^3_{\mathrm{f},i} \right|, \tag{18}$$

whereas Eq. (17) is replaced by

$$\tilde{R}'_{\mathrm{f}} = \sum_i \left| \mathrm{CC}^3_{\mathrm{f},i} \right|. \tag{19}$$

For the alignment of noise-free images, the value of $\tilde{R}_{\mathrm{f}}$ ($\tilde{R}'_{\mathrm{f}}$) assumes 1 when the image is aligned perfectly, and will be smaller for misaligned images. The absolute value in Eqs. (18) and (19) introduces an ambiguity in the alignment because, for a certain alignment, the same value of $\tilde{R}_{\mathrm{f}}$ ($\tilde{R}'_{\mathrm{f}}$) is obtained after inverting either the image or the reference. A particle of certain shape and symmetry that would produce the negative image of itself upon rotation would, therefore, have at least two possible alignments. Proteins with such properties do not exist, however, and this ambiguity will not be considered in the following discussion.

The use of Eq. (18) or (19) in the alignment of noise images that do not contain signal leads to a rather interesting result: For a particular noise image, there is an equal chance for every resolution zone to be aligned "in phase" (positive CC) or "out-of-phase" (negative CC). This means there is no overall bias of the alignment towards a positive correlation between noise image and reference, and therefore, the average (3D reconstruction) calculated from the aligned noise images will not carry any bias either. This is a remarkable result because, for pure noise images, Eqs. (18) and (19) predict that there will be no bias in an iterative refinement, and the FSC used for resolution measurement, as described above, will be close to zero, indicating the absence of signal.

When both signal and noise are present, the noise can introduce a small bias in the alignment because, due to the signal, the correlation between image and reference is now biased towards positive values. This positive bias depends on the strength of the signal. At low resolution where the SNR is usually highest, the alignment using Eq. (18) or (19) will produce mostly positive correlation values, and the bias will be strongest. However, the stronger signal at low resolution is less prone to over-fitting, and the noise bias will not have a strong effect on the overall alignment of the particle. As the signal becomes weaker towards higher resolution, the chance of a negative correlation due to noise (or misaligned signal) will increase. The alignment of the low-resolution signal will help in the alignment of the high-resolution signal because of its "coherence", as described above. At a resolution where the signal is negligible, no positive bias exists, as discussed in the previous paragraph. The use of Eq. (18) or (19) in the alignment ensures, therefore, that only parts of the data that contain signal will add up in phase when averaged, and that the parts that contain predominantly noise will remain uncorrelated with the average. This property will be demonstrated in the following section using computer simulations.

Fig. 1. Refinement runs using the phase residual (PRES), linear correlation coefficient (CC) and weighted correlation coefficient (weighted). The test structures used to generate the simulated data sets consisted of two cubes of 15 pixels length on the side, and overlapping at one corner (panel A), and eight randomly distributed spheres with a diameter of five pixels (panel B). Panel C shows FSC curves between refined structures and the noise-free reference. The refinement was carried out with 30,000 random projections of the test structure with the two cubes (panel A), using the three different refinement functions. Panel D shows FSC curves between two 3D reconstructions from two halves of the data set derived from the two cubes after refinement using the three different refinement functions. These curves are normally used to estimate the resolution of the refined structure. Panel E compares the FSC curve between the refined structure and the noise-free reference (labeled model weighted, same as in panel C) with the predicted FSC curve (labeled $\mathrm{FSC}_{\mathrm{C}}$ Weighted, Eq. (20)) based on the FSC curve between the two 3D reconstructions from two halves of the data set (see panel D). Panels F–H show histograms of the absolute errors in the angles after refinement using the phase residual, linear correlation coefficient and weighted correlation coefficient, respectively. PHI, THETA and PSI are the three Euler angles (rotation about the z-axis, tilt about the y-axis and rotation about the new z-axis). As expected, the distribution is widest for the phase residual and narrowest for the weighted correlation coefficient.

7. Computer simulations

In the following, the results from a number of computer simulations are shown to demonstrate differences in the refinement algorithms using the similarity measures given by Eqs. (2), (3) and (19). Each function was implemented in the program FREALIGN [14] used to carry out the refinement. Two test structures were constructed on the computer and are shown in Figs. 1A and B. Units of dimension and frequency are given in pixel and 1/pixel, respectively. The first structure consists of two cubes that overlap at one corner. Although this structure contains a two-fold symmetry, the

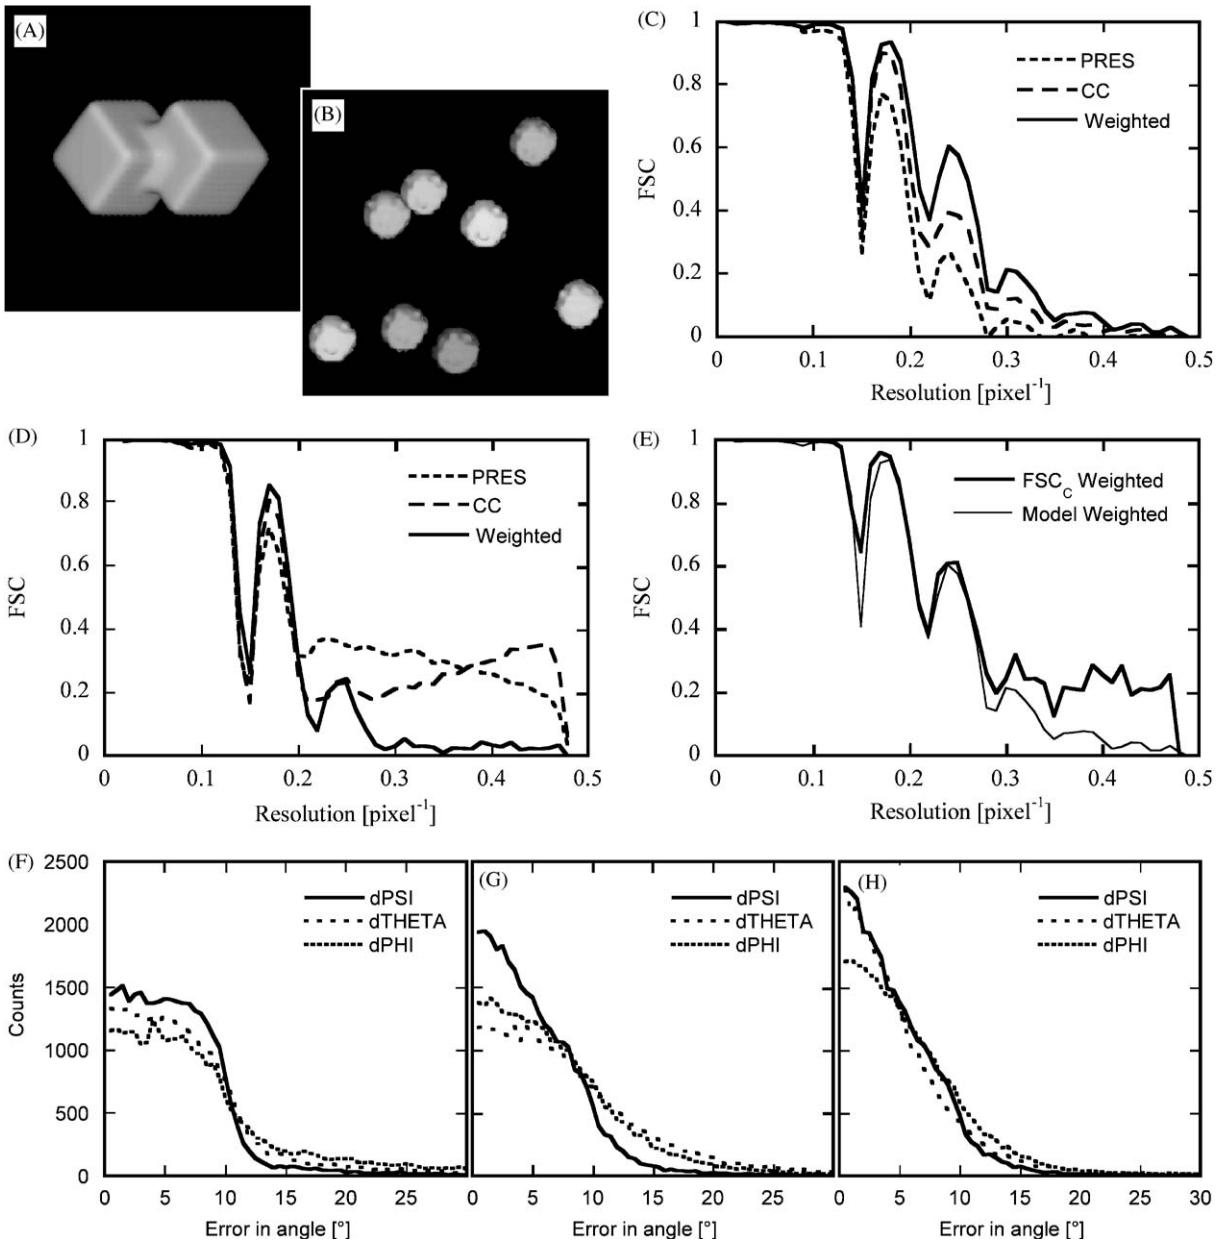

symmetry was never imposed during the refinement. A Gaussian low-pass filter $G(k) = \exp[-k^2/0.18]$ (SPIDER [25], command FQ, Gaussian filter with radius 0.3) was applied to the structure to attenuate the strong high-resolution ripples originating from the sharp edges of the cubes. The second structure is a collection of eight randomly positioned spheres, each with a diameter of five pixels. This structure has no symmetry, and no filter was applied to attenuate the high-resolution terms.

The test structures are used to generate projections in random orientations and with Gaussian-distributed positional displacements from the image center that had a standard deviation of 2 pixels. Each data set contained 30,000 projections in images of 100×100 pixels. A contrast transfer function (CTF) was applied, in analogy to images obtained from an electron microscope. The CTF did not have an envelope attenuating its amplitude towards high resolution. Defocus values for the CTF varied between 40 and 60 in generalized units [26]. Finally, Gaussian noise was added to the image to give an SNR of $\frac{1}{50}$.

To test the performance of the algorithms in a refinement, the parameters describing the orientations and positions of each particle in the simulated data sets were perturbed by an addition of random, uniformly distributed angles with a standard deviation of about 10° and random, uniformly distributed displacements with a standard deviation of about two pixels. During the refinement, the particles were masked using a circular mask with a radius of 37 pixels. Fig. 1C shows the results of 9 cycles of refinement, as described above, for the data set derived from the two cubes. The resolution of the final structure was measured by the FSC between the refined structure and the original, noise-free model used to generate each data set.

For both data sets, all refinement runs improved the resolution of the structure significantly over the initial structure calculated using the perturbed particle parameters. All refinement algorithms had converged after about 5 cycles, and there was no further change in the reconstruction. The different refinement algorithms differ, however, in the resolution they are able to achieve, with the algorithm using the phase residual (Eq. (3)) the worst and that using the weighted correlation coefficient (Eq. (19)) the best.

As explained above, we expect the weighted correlation coefficient to produce less over-fitting than the linear correlation coefficient or the phase residual. To assess the degree of over-fitting, the resolution of the final structure is also determined using the commonly used method of calculating the FSC between reconstructions from two halves of the data sets. The FSC plots for each refinement run are plotted in Fig. 1D. In contrast to the true resolution calculated against the noise-free original structure shown in the previous plots, the indicated resolution is now highest for the algorithm using the phase residual (Eq. (3)), and lowest for that using the weighted correlation coefficient (Eq. (19)). Clearly, the weighted correlation coefficient leads to the smallest degree of over-fitting. A more quantitative measurement of the degree of over-fitting can be obtained by converting the FSC between the two reconstructions from two halves of a data set into an expected $\mathrm{FSC_C}$ a reconstruction from the entire data set would have with a perfect reference structure. Rosenthal and Henderson [27] showed that this conversion is given by

$$\mathrm{FSC_C} = \sqrt{2\mathrm{FSC}/(1+\mathrm{FSC})}. \tag{20}$$

They argue that a realistic resolution cut-off for a structure is a value for the FSC of 0.143, corresponding to 0.5 for the $\mathrm{FSC_C}$, or an SNR of the final structure of 1. The value of 0.143 replaces the often quoted cut-off value of 0.5 for the FSC. Fig. 1E compares plots of the $\mathrm{FSC_C}$, as well as the actual FSC curves evaluated between the final reconstructions and the noise-free model with the two cubes used to generate the data set, for the refinement using the weighted correlation coefficient. The agreement is striking, suggesting that there is very little over-fitting using the weighted correlation coefficient.

To confirm the improved alignment of the particles using the weighted correlation coefficient, histograms of the absolute differences between the true angles for each particle and the angles after 9 refinement cycles are shown in Figs. 1F–H, for the refinement runs using the phase residual, linear correlation coefficient and weighted correlation

coefficient, respectively. As expected, the histogram is widest for the phase residual and narrowest for the weighted correlation coefficient.

As a final refinement test using the two cubes, the angles were completely randomized, and 9 cycles of refinement were carried out using as the first reference the reconstruction resulting from the randomized particle orientations. Although this first reconstruction looked almost spherical, reasonable structures emerged (not shown) when using the linear or weighted correlation coefficient, but not with the phase residual. The FSC curves evaluated between these structures and the noise-free model (not shown) gave a value of 0.5 at a resolution of 0.14 for both correlation functions although the weighted correlation coefficient gave slightly higher FSC values. As before, FSC curves between reconstructions from two halves of the data set indicated significant over-fitting for the linear correlation coefficient and little over-fitting for the weighted correlation coefficient. The randomization experiment shows that the chance of converging onto a point that is completely off the global optimum depends strongly on the choice of the target function used in the refinement.

In an attempt to better understand how the over-fitting occurs, one cycle of refinement was performed on the data set originating from the two cubes, using the two cubes as a noise-free reference. The noise-free reference ensures that any error in the alignment is solely due to the noise in the data set. The FSC between the structure resulting from this single step of the refinement and the noise-free reference is shown in Fig. 2A for each refinement algorithm. Surprisingly, the refinement using the phase residual (Eq. (3)) gives the highest resolution whereas that using the linear and the weighted correlation coefficients give the poorest resolution. The very nature of over-fitting means, however, that part of the good agreement between the refined structures and the reference comes from the agreement between the reference and the noise component of the data.

Since the test data set was computer-generated, unlike experimental data, both the noise component and the signal component for each image are known exactly. Using this knowledge, the 3D reconstruction can be disassembled into the contributions from the noise and the signal, which add together to form the reconstructed model. For an unbiased alignment and reconstruction, the noise should be unaligned and bear no correlation with the signal. By using the alignments from the test data set we can reconstruct a 3D model using only the signal from the data set images and another 3D model using only the noise that was added to the images in the data set. Figs. 2B–E show cross-sections through the reconstructions from the noise images for each of the alignment algorithms, together with slices through the noise-free reference structure for comparison. As expected for the alignment of noisy data, the reconstructions from the images containing the noise component clearly display some of the features of the reference structure. This is most visible for the alignment using the phase residual (Eq. (3)) and least visible for the alignment using the weighted correlation coefficient (Eq. (19)). Closer inspection of the noise reconstruction reveals that there is some amplification of edges in the features, indicating that the over-fitting is more severe at high frequencies where the SNR is low. Fig. 2F shows FSC plots between the noise reconstructions and the noise-free reference, confirming the different degrees of over-fitting for the three alignment algorithms at high resolution. At low resolution, there is a significant noise bias also for the alignment that used the weighted correlation coefficient. As explained above, this is expected since at low resolution the absolute value in Eqs. (18) and (19) will not prevent bias due to the strong signal. The same analysis is shown in Fig. 2G for the signal component of the aligned data set. The signal component shows the opposite trend: The alignment based on the phase residual shows the lowest FSC whereas that based on the weighted correlation coefficient shows the best FSC across the entire resolution range. Therefore, even though the reconstruction from the aligned images that contain signal and noise showed the best agreement with the reference when using the phase residual, and the worst agreement when using the weighted correlation coefficient, the accuracy of the alignment was worst with the former and best with the latter.

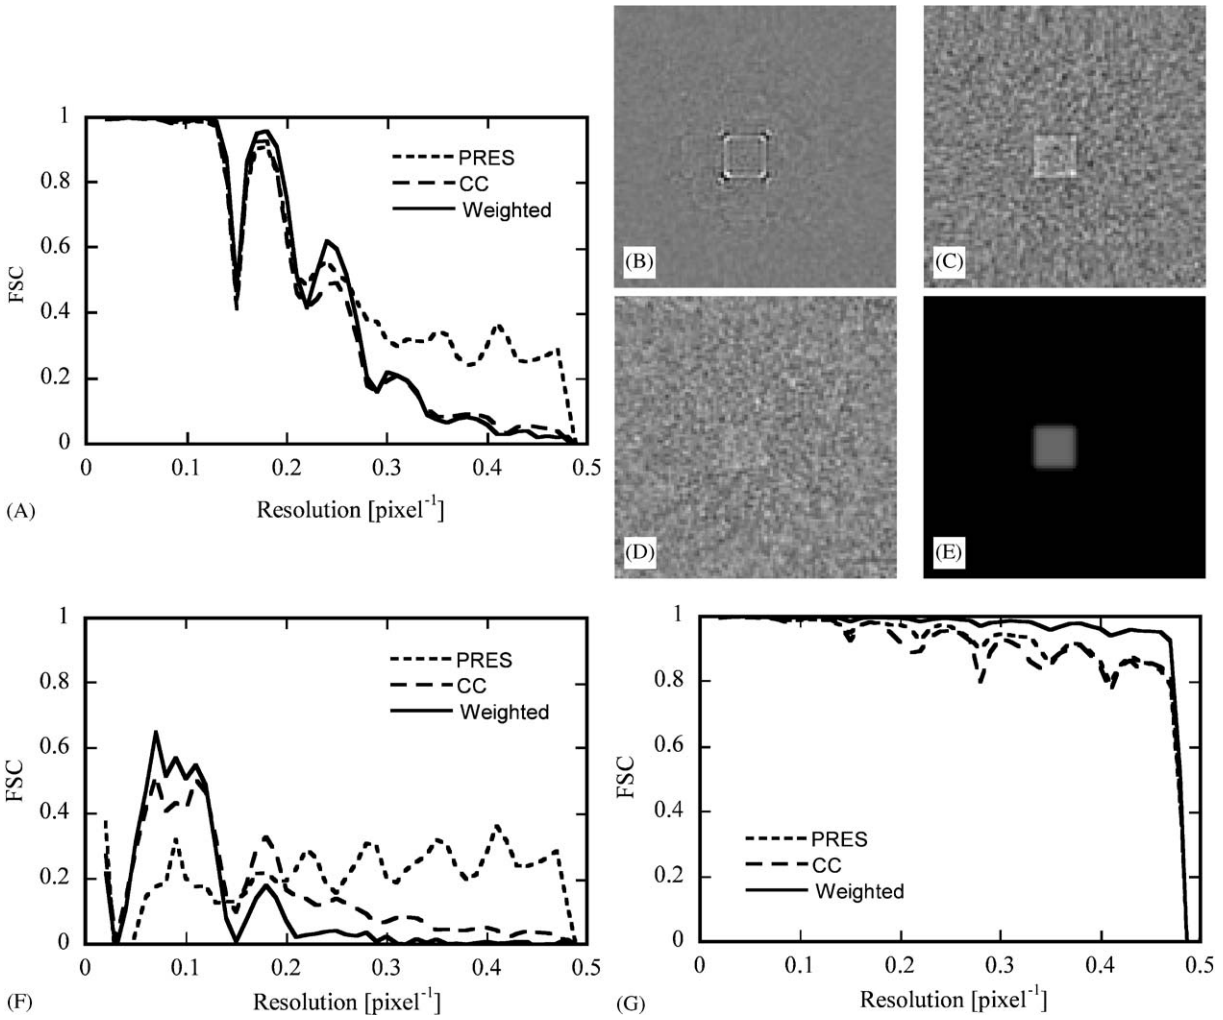

Fig. 2. Analysis of the noise bias. Panel A shows FSC curves between refined structures and the noise-free reference (the two cubes) after one cycle of refinement against the noise-free reference, using the three different refinement functions. Cross-sections through the 3D reconstructions calculated from the aligned noise images are shown in panels B, C and D, using the phase residual, linear correlation coefficient and weighted correlation coefficient, respectively, in the alignment. The slice in panel B shows mainly the edges of the reference structure, indicating over-fitting especially at high resolution. The noise replica in panel C also has low-resolution components. The replica in panel D is weakest, indicating the least over-fitting. Panel E shows a cross-section through the noise-free reference, for comparison. FSC curves between noise reconstructions and the noise-free reference are shown in panel F. The over-fitting of the noise at high resolution is strongest for the refinement using the phase residual (PRES), and weakest for that using the weighted correlation coefficient (weighted). At low resolution where the signal in the test data set is strong, this order is reversed. Panel G shows FSC curves between the corresponding signal reconstructions and the noise-free reference. The over-fitting of the noise in the refinement using the phase residual (PRES) and linear correlation coefficient (CC) leads to a larger error in the image alignment, resulting in a lower resolution of the reconstruction, compared with the alignment using the weighted correlation coefficient (weighted).

The simulations described for the test structure containing the two cubes was repeated with the structure containing the eight small spheres (not shown). The distribution of the signal across resolution differs between the two structures, and the second structure served as a validation for the results obtained for the first structure. As expected, the weighted correlation coefficient

performed best, with the highest resolution obtained in the refinement tests, and the least over-fitting visible in the noise reconstructions.

8. The Fourier shell correlation as a resolution measure

As discussed above, the commonly used FSC as a resolution measure for structures derived from single particles can display a significant bias towards artificially high correlation values. The plots in Fig. 1E, comparing the FSC_C with the true FSC between the refined structure and the noise-free reference, suggest that the reduced over-fitting by the weighted correlation coefficient avoids significant bias in the FSC calculated between reconstructions from two halves of the data set.

Additional simulations were carried out using the eight spheres as a test structure. A band-pass filter was applied to the structure, setting all Fourier terms to zero in the resolution shell between 0.24 and 0.33. As before, a data set containing 30,000 randomly oriented projections with an SNR of $\frac{1}{50}$ were calculated. Due to the band pass filter, the images contained only noise in the masked resolution shell; signal was only present outside this mask. This type of test was introduced by Shaikh et al. [28]. Again, 9 cycles of refinement were carried out using the phase residual, linear correlation coefficient and weighted correlation coefficient. The FSC plots between reconstructions from two halves of the data set are shown in Fig. 3A, and a plot of

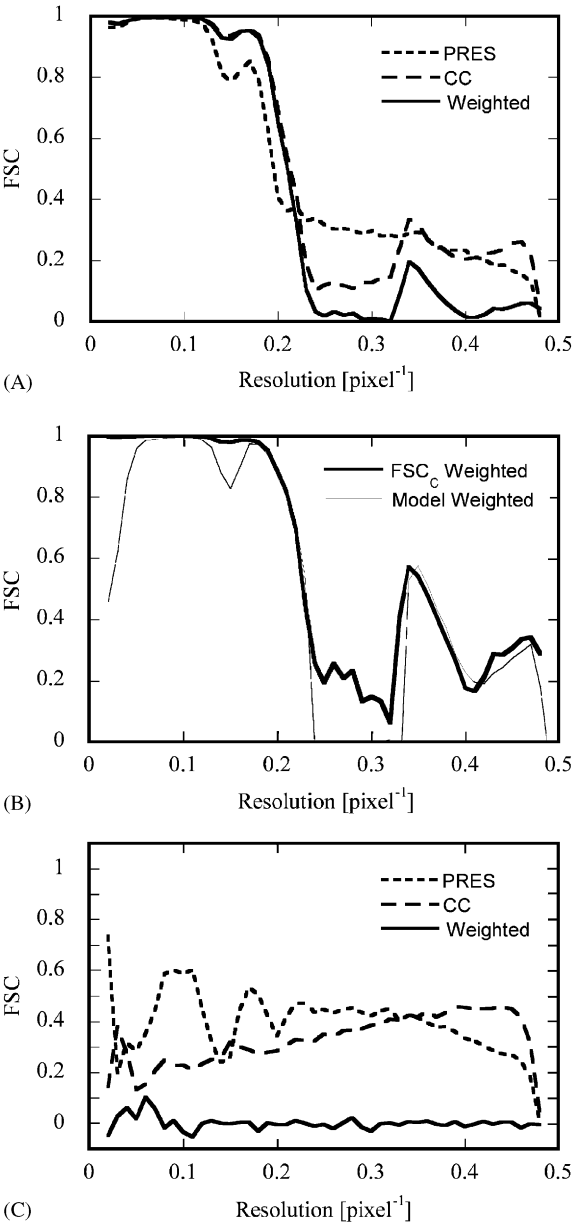

Fig. 3. FSC as a resolution measure. Panel A shows FSC curves between two 3D reconstructions from two halves of the data set derived from the eight spheres (Fig. 1B), after refinement using the three different refinement functions. The data set was band-pass filtered such that no signal was present at a resolution between 0.24 and 0.33. The absence of signal in this resolution interval is clearly reflected in the FSC curve resulting from the refinement using the weighted correlation coefficient (weighted). The other alignment functions lead to significant over-fitting in this resolution interval. In panel B, the FSC curve between the refined structure and the noise-free reference (labeled model weighted) is compared with the predicted FSC curve (labeled FSC_C weighted, Eq. (20)) based on the FSC curve between the two 3D reconstructions from two halves of the data set in panel A. Panel C shows FSC curves between two 3D reconstructions from two halves of the data set containing pure noise, after refinement using the three different refinement functions. The curves indicate significant over-fitting for the refinement using the phase residual (PRES) and the linear correlation coefficient (CC). The FSC curve resulting from the refinement using the weighted correlation coefficient (weighted) remains close to zero across the entire resolution range, as would be required for a reliable resolution measure.

the FSC_C and the true FSC between the refined structure using the weighted correlation coefficient and the noise-free reference are shown in Fig. 3B. As before, the FSC between reconstructions from two halves of the data set indicate significant overfitting for the refinement runs using the phase residual (Eq. (3)) and linear correlation coefficient (Eq. (2)). The FSC plot for the refined structure using the weighted correlation coefficient, however, indicates no significant signal in the masked resolution shell, consistent with its true FSC between the refined structure and the noise-free reference.

Ideally, a resolution measure such as the FSC would indicate the absence of signal in data that contains only noise. A further test was carried out with a data set containing 30,000 images of pure Gaussian-distributed noise. Fig. 3C shows plots of the FSC between reconstructions from two halves of the data set using the three different alignment functions. The plots for the refinement runs using the phase residual (Eq. (3)) and linear correlation coefficient (Eq. (2)) show significant correlation across the entire resolution range. This result was also obtained in refinement runs in 2D using simulated noise [10]. The FSC plot for the refinement using the weighted correlation coefficient (Eq. (19)) indicates essentially random correlation, as would be expected from a reliable resolution measure.

9. Comparison with the maximum likelihood algorithm

One of the most promising alignment algorithms that is capable of dealing with extremely small SNRs is the direct maximum likelihood algorithm, as demonstrated for alignments in a plane by Sigworth [17]. The main drawback of this algorithm is the necessity to integrate over the entire parameter space which is feasible for alignments within a plane (three parameters) but currently impractical for alignments in 3D (five parameters). The refinement algorithms discussed in the previous section were tested on one of the data sets used in Sigworth's study to compare their performance at a very small SNR. The data set used contained 4000 images of a projection of a bacteriorhodopsin trimer, sampled at 1.5 Å/pixel, and with an SNR of $\frac{1}{200}$ ($z = 0.07$ in Sigworth's study [17]). In-plane rotations of the trimer in the images were uniformly distributed, and their displacement from the center of the image had a Gaussian-distribution with a standard deviation of 10 pixels. The noise-free bacteriorhodopsin trimer and one member of the noisy data set are shown in Figs. 4A and B, respectively. As shown in Ref. [17], an alignment using a linear correlation coefficient (Eq. (2)) does not lead to any recognizable features in the final averaged structure. We found the same to be true when using the phase residual as the alignment function (Eq. (3), not shown). Fig. 4C shows the refined structure of the bacteriorhodopsin trimer, after 3000 refinement cycles using the weighted correlation coefficient (Eq. (18)) and step filtering. In this refinement, it was necessary to perform 2 cycles of alignment using Eq. (18) without the absolute value to allow an initial rough centering of the bacteriorhodopsin trimers. This step is normally not necessary since, as discussed above, we assume that the refinement procedure starts with a rough structure that was determined by other means, and therefore, that the particles are already roughly centered. Furthermore, after centering, a mask was applied to the images to mask out noise near the image boundaries. The large number of refinement cycles, compared with only 274 used in the maximum likelihood procedure [17], that were needed before convergence was reached is partly due to the fact that the alignment in each cycle involved only a local maximization of the weighted correlation, contrasting with the integration over the entire parameter space in the maximum likelihood procedure.

According to the Fourier ring correlation (FRC) in Fig. 4D, features of the trimer in Fig. 4C are resolved up to a resolution of about 0.11, corresponding to $0.07\,Å^{-1}$. The resolution after refinement using the maximum likelihood algorithm was significantly higher, about 0.2 ($0.13\,Å^{-1}$, [17]), indicating its superior performance to that using the weighted correlation coefficient. Part of this superior performance might be due to the additional information about the distribution of

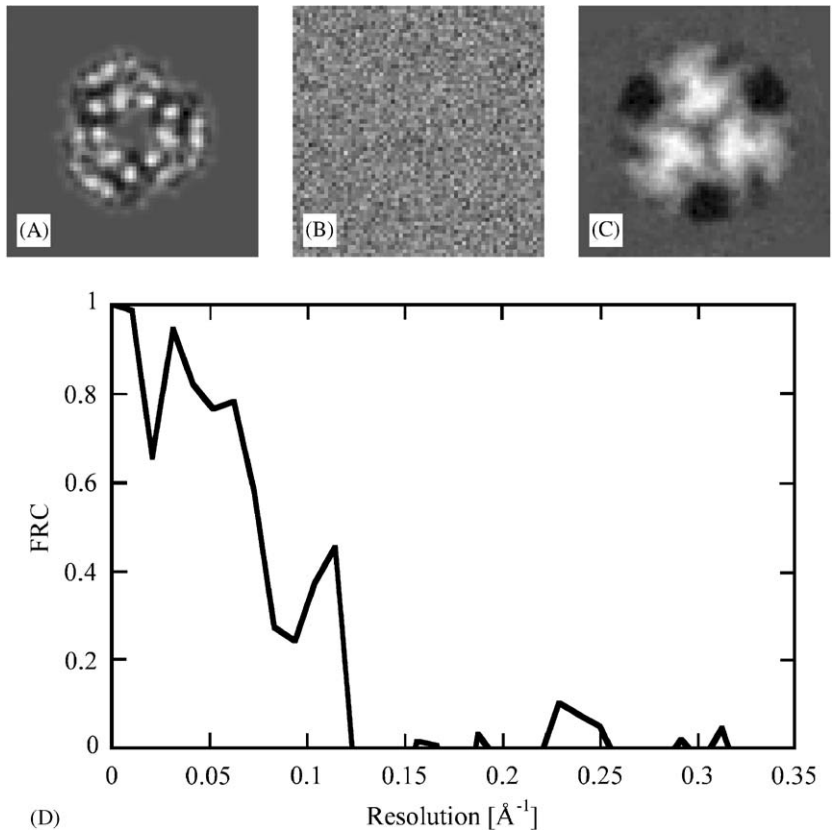

Fig. 4. Performance of the refinement algorithm using the weighted correlation coefficient and step filtering, when applied to data with extremely low SNR. Panel A shows the noise-free structure used to generate 4000 images with random displacements and rotations (courtesy of Fred Sigworth, [17]). Noise was added to these images to give an SNR of $\frac{1}{200}$. An example of the resulting images is shown in panel B. Panel C shows the refined structure after 3000 refinement cycles. Panel D shows the FRC curve between the refined structure and the noise-free reference.

the displacements of the trimers from the image center that was used by the maximum likelihood algorithm. Nevertheless, although the weighted correlation coefficient given in Eq. (18) is a significant improvement over the linear correlation coefficient and the phase residual, further improvement in refinement procedures can be expected with the implementation of the maximum likelihood algorithm in 3D.

10. Discussion

In the refinement of 3D structures obtained from images of single particles, one of the main challenges is to obtain the best possible alignment of all images. For an optimal alignment, it is important to find appropriate weights for parts of the data that share similar SNR values. The weighting of the data ensures that the alignment will follow the strong parts of the data, and it reduces the influence of the weaker data that is normally found at the high-resolution end. Because the weighting depends on the correlation between the image and the reference, and therefore is determined both by the actual SNR present in the data, as well as the accuracy of the alignment, it will change during the alignment of the image. If we assume that the alignment will lead to a better agreement between the image and the reference at

high resolution, the weight of this part of the data is increased during the refinement. Apart from the particle position and orientation, the CTF and magnification of a particle will also determine the agreement between the image and the reference, and they can be included in the refinement. The adjustment in the weighting is similar to a multi-resolution refinement algorithm where the resolution of the included data is gradually increased as the refinement progresses. Such as strategy has been shown to have a superior convergence radius, compared with a single resolution approach [29].

Our approach differs from a previously described approach [30] where a fixed cut-off for the data was introduced, based on the correlation of a modeled CTF with oscillations visible in the power spectrum of the raw images. A fixed cut-off has the advantage that noise that lies beyond the cut-off cannot introduce errors in the alignment. However, the CTF oscillations in an image are not always a good indicator for the signal preset in the data (see for example, the power spectrum from images of 2D crystals, [31]). Even if the signal is weak, and therefore, does not give rise to recognizable CTF oscillations in an image power spectrum, it can be extracted from the image provided the molecule or complex is properly aligned, such as in a 2D crystal. A procedure that uses a fixed cut-off would be limited in resolution by this cut-off.

The convergence of the single-step algorithm is much faster than that using the step filtering. Part of the reason for the slow convergence of the latter is scaling of the step by the inverse spatial frequency in our current procedure. A more sophisticated scaling method should also take into account the SNR of the refined 3D structure to allow faster convergence for parts of the data that are more reliable. This will be addressed in a future publication. The weighting of image data at the 3D reconstruction stage has not been explored here. Weighting schemes, again based on an SNR estimate, have been proposed [32] and should further improve the resolution and convergence of the refinement.

An important consequence of the weighting is the reduction of over-fitting of the data. Over-fitting results in sub-optimal refinement and produces an over-optimistic resolution estimate of the FSC and spectral SNR [33,34]. The FSC curves for the refinement using the weighted correlation coefficient (Eq. (19)) in Figs. 1E and 3B still show a small degree of over-fitting but the estimated resolution of the refined 3D structures according to an FSC cut-off of 0.143 does not differ much from the 0.5 cut-off for the FSC between the refined structure and the noise-free reference. We expect that a refinement using a maximum likelihood algorithm would eliminate this residual over-fitting and yield an even more accurate FSC resolution estimate.

The FSC provides only a measurement of the average resolution. If parts of the structure are disordered, undergo conformational changes, or the reconstructed volume has segments that do not actually contain any part of the particle but only noise, the FSC will underestimate the resolution in the ordered parts of the structure. Empty parts of the reconstructed volume can be masked to remove noise (solvent flattening). Flexible or disordered parts of the structure, however, cannot simply be masked but should be weighted according to their degree of disorder. Such a real-space weighting would not only improve the accuracy of the FSC resolution estimate but also serve as a real-space filter in the alignment, to enhance the influence of the well-ordered parts of the structure on the alignment. The real-space weights could be derived from a variance-based real-space SNR measurement in the structure.

11. Conclusions

The two key features of the algorithm described here are that the weighted correlation function for alignment is chosen to optimize the fit of the data to the model by adapting to the resolution dependent SNR at each iteration, and that the optimizer does not distinguish between positive and negative correlation, ensuring that purely noise-driven alignment with the model will contribute both negatively and positively to the reconstruction and so will not be reinforced in the next round.

The incorporation of the weighted correlation function for alignment into the FREALIGN program [14] has improved our ability to reach high resolution in structures derived from images of single particles, even if these are small (for example, the transferrin receptor [35]). The computer model experiments above show this to be due to lower sensitivity to noise in the alignment process. However, it is also clear from this work that an implementation of the direct maximum likelihood method for the full five-parameter alignment problem should do better. This is particularly relevant in attempts to improve resolution by combining very weak signals from data sets with tens or hundreds of thousands of images.

Acknowledgements

The authors like to thank David DeRosier and Duncan Sousa for their interest and criticism during discussions involving this work. They also thank Fred Sigworth for making available his very noisy data set for test purposes. N.G. gratefully acknowledges financial support from the National Institutes of Health, grant 1 P01 GM-62580.

Appendix A

Following Rosenthal and Henderson [27], we set $F1 = S + N1$ and $F2 = S + N2$, for the Fourier components of the particle image and reference, respectively. S is the signal component common to both, and $N1$ and $N2$ are the respective noise components. The correlation coefficient between the two is then given by

$$\mathrm{CC} = \frac{\sum (S+N1)(S+N2)^{\dagger}}{\sqrt{\sum |S+N1|^2 \sum |S+N2|^2}}. \tag{A.1}$$

For uncorrelated noise and both images scaled equally, this simplifies to

$$\begin{aligned}\mathrm{CC} &= \frac{\sum |S|^2}{\sum |S|^2 + \sum |N|^2}, \\ &= \frac{\mathrm{SNR}}{1+\mathrm{SNR}}\end{aligned} \tag{A.2}$$

with $\mathrm{SNR} = \sum |S|^2 / \sum |N|^2$ and N equal to one of the noise components. Eq. (A.2) is equivalent to Eq. (9) in the main text. If the reference does not contain noise, $N2 = 0$ and

$$\begin{aligned}\mathrm{CC} &= \frac{\sum (S+N1)S^{\dagger}}{\sqrt{\sum |S+N1|^2 \sum |S|^2}}, \\ &= \frac{\sqrt{\sum |S|^2}}{\sqrt{\sum |S+N1|^2}}, \\ &= \sqrt{\frac{\sum |S|^2}{\sum |S|^2 + \sum |N1|^2}}, \\ &= \sqrt{\frac{\mathrm{SNR}}{1+\mathrm{SNR}}},\end{aligned} \tag{A.3}$$

which is equivalent to Eq. (10) in the main text.

References

[1] E. Nogales, N. Grigorieff, J. Cell. Biol. 152 (2001) F1.
[2] O. Sokolova, L. Kolmakova-Partensky, N. Grigorieff, Structure (Camb) 9 (2001) 215.
[3] M.S. Jurica, D. Sousa, M.J. Moore, N. Grigorieff, Nat. Struct. Mol. Biol. 11 (2004) 265.
[4] B. Bottcher, S.A. Wynne, R.A. Crowther, Nature 386 (1997) 88.
[5] J.F. Conway, N. Cheng, A. Zlotnick, P.T. Wingfield, S.J. Stahl, A.C. Steven, Nature 386 (1997) 91.
[6] Z.H. Zhou, M.L. Baker, W. Jiang, M. Dougherty, J. Jakana, G. Dong, G. Lu, W. Chiu, Nat. Struct. Biol. 8 (2001) 868.
[7] R. Matadeen, A. Patwardhan, B. Gowen, E.V. Orlova, T. Pape, M. Cuff, F. Mueller, R. Brimacombe, M. van Heel, Struct. Fold. Des. 7 (1999) 1575.
[8] M.G. Gomez-Lorenzo, C.M. Spahn, R.K. Agrawal, R.A. Grassucci, P. Penczek, K. Chakraburtty, J.P. Ballesta, J.L. Lavandera, J.F. Garcia-Bustos, J. Frank, EMBO J. 19 (2000) 2710.
[9] R. Henderson, Q. Rev. Biophys. 28 (1995) 171.
[10] N. Grigorieff, Acta Crystallogr. D Biol. Crystallogr. 56 (Pt 10) (2000) 1270.
[11] M. Radermacher, T. Wagenknecht, A. Verschoor, J. Frank, J. Microsc.—Oxford 146 (1987) 113.
[12] M. van Heel, G. Harauz, Scanning Microsc. 2 (Suppl.) (1988) 295.
[13] P.A. Penczek, R.A. Grassucci, J. Frank, Ultramicroscopy 53 (1994) 251.
[14] N. Grigorieff, J. Mol. Biol. 277 (1998) 1033.
[15] H.O. Hartley, Biometrics 14 (1958) 174.
[16] A.P. Dempster, N.M. Laird, D.B. Rubin, J. R. Statist. Soc. Ser. B 39 (1977) 1.

[17] F.J. Sigworth, J. Struct. Biol. 122 (1998) 328.
[18] W. Jiang, Z. Li, Z. Zhang, M.L. Baker, P.E. Prevelige Jr., W. Chiu, Nat. Struct. Biol. 10 (2003) 131.
[19] G. Harauz, M. van Heel, Optik 73 (1986) 146.
[20] M. Vanheel, M. Schatz, E. Orlova, Ultramicroscopy 46 (1992) 307.
[21] T. Walz, N. Grigorieff, J. Struct. Biol. 121 (1998) 142.
[22] J. Frank, L. Al-Ali, Nature 256 (1975) 376.
[23] J. Skilling, R.K. Bryan, Mon. Not. R. Astron. Soc. 211 (1984) 111.
[24] J.M. Carazo, C.O. Sorzano, E. Rietzel, R. Schröder, R. Marabini, Discrete tomography in electron microscopy, in: G.T. Herman, A. Kuba, eds.), Discrete Tomography, Birkhäuser, Boston, 1999, p. 405.
[25] J. Frank, M. Radermacher, P. Penczek, J. Zhu, Y. Li, M. Ladjadj, A. Leith, J. Struct. Biol. 116 (1996) 190.
[26] R.H. Wade, Ultramicroscopy 46 (1992) 145.
[27] P.B. Rosenthal, R. Henderson, J. Mol. Biol. 333 (2003) 721.
[28] T.R. Shaikh, R. Hegerl, J. Frank, J. Struct. Biol. 142 (2003) 301.
[29] J.Z. Chen, J. Furst, M.S. Chapman, N. Grigorieff, J. Struct. Biol. 144 (2003) 144.
[30] Z. Huang, P.R. Baldwin, S. Mullapudi, P.A. Penczek, J. Struct. Biol. 144 (2003) 79.
[31] R. Henderson, J.M. Baldwin, K.H. Downing, J. Lepault, F. Zemlin, Ultramicroscopy 19 (1986) 147.
[32] S.J. Ludtke, J. Jakana, J.L. Song, D.T. Chuang, W. Chiu, J. Mol. Biol. 314 (2001) 253.
[33] M. Unser, B.L. Trus, A.C. Steven, Ultramicroscopy 23 (1987) 39.
[34] M. Unser, B.L. Trus, J. Frank, A.C. Steven, Ultramicroscopy 30 (1989) 429.
[35] Y. Cheng, O. Zak, P. Aisen, S.C. Harrison, T. Walz, Cell 116 (2004) 565.