Published in final edited form as:

J Struct Biol. 2013 September ; 183(3): . doi:10.1016/j.jsb.2013.07.005.

Likelihood-based classification of cryo-EM images using FREALIGN

Dmitry Lyumkis^{1,†}, Axel F. Brilot^{2,†}, Douglas L. Theobald², and Nikolaus Grigorieff^{2,3,*}

¹National Resource for Automated Molecular Microscopy, Department of Integrative Structural and Computational Biology, The Scripps Research Institute, La Jolla, CA 92037, USA

²Department of Biochemistry, Rosenstiel Basic Medical Sciences Research Center, Brandeis University, MS029, 415 South Street, Waltham, MA 02454, USA

³Howard Hughes Medical Institute, Brandeis University, MS029, 415 South Street, Waltham, MA 02454, USA

Abstract

We describe an implementation of maximum likelihood classification for single particle electron cryo-microscopy that is based on the FREALIGN software. Particle alignment parameters are determined by maximizing a joint likelihood that can include hierarchical priors, while classification is performed by expectation maximization of a marginal likelihood. We test the FREALIGN implementation using a simulated dataset containing computer-generated projection images of three different 70S ribosome structures, as well as a publicly available dataset of 70S ribosomes. The results show that the mixed strategy of the new FREALIGN algorithm yields performance on par with other maximum likelihood implementations, while remaining computationally efficient.

Keywords

Electron microscopy; maximum likelihood; classification; single particle; protein structure

1. Introduction

Electron cryo-microscopy (cryo-EM) is a versatile technique to visualize the three-dimensional (3D) structure of macromolecules and their assemblies. Since its inception in the 1970s and 80s (Adrian et al., 1984; Taylor and Glaeser, 1974), methods for sample preparation, microscopy instrumentation and image processing have all been significantly improved. Research in the past decade has focused especially on developing the single-particle technique, which can be applied to non-crystalline, isolated molecules. The single-particle technique has recently been applied to a dataset of 80S ribosomes to obtain a resolution of about 4.5 Å (Bai et al., 2013). The data was recorded with a direct electron detector that registers electrons directly rather than via a scintillator. This new detector

© 2013 Elsevier Inc. All rights reserved.

*Correspondence should be sent to: Nikolaus Grigorieff, Brandeis University, MS029, 415 South Street, Waltham, MA 02454, USA, Tel: (781) 736-2444; fax: (781) 736-2419; niko@brandeis.edu.

[†]These authors contributed equally.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

technology is also capable of recording movies that capture sample motions occurring during beam exposure, allowing for motion correction (Brilot et al., 2012; Campbell et al., 2012). The work on the 70S ribosome further highlighted the importance of new image processing algorithms, most notably the application of Bayesian statistics and the related maximum likelihood (ML) approach for the estimation of 3D structures from the noisy images obtained using low-dose imaging. These methods are less susceptible to initial model bias and over-refinement, especially at low signal-to-noise ratio (SNR) (Sigworth, 1998; Sigworth et al., 2010).

Originally introduced for the processing of images of viruses and the ribosome (Doerschuk and Johnson, 2000; Provencher and Vogel, 1988; Vogel and Provencher, 1988), ML was later developed as a more general approach for single particle structure estimation (Scheres et al., 2005a; Scheres et al., 2005b; Sigworth, 1998). A general difficulty to overcome when implementing ML and Bayesian methods is the representation of the noise in cryo-EM data. Implementation of a Gaussian prior for the noise is straightforward and leads to efficient maximization algorithms. Therefore, current implementations assume that data points (pixels) are independent of each other either in real space or reciprocal space (Scheres, 2010). Assuming independence in real space, one cannot accommodate correlations due to the contrast transfer function (CTF) of the electron microscope. Conversely, assuming independence in reciprocal space prohibits real-space masking of noise surrounding each particle to boost overall SNR. The RELION software (Scheres, 2012b) implements an ML approach with a hierarchical prior to impose smoothness in the estimation of single particle structures from cryo-EM images. It performs most calculations in Fourier space, thus taking CTF effects into account. Using RELION, the authors demonstrate superior results compared to most other software that does not make use of the ML formalism. Specifically, they show improved resolution in the final 3D reconstruction, reduced refinement artifacts such as inflated resolution estimation (overfitting), and better separation of structurally distinct particles in heterogeneous datasets (mixtures).

The FREALIGN software (Grigorieff, 2007) used in the present study is designed to refine single particle reconstructions when an initial reconstruction at lower resolution is already available. While FREALIGN is not based on ML or Bayesian statistics, it also implements a series of measures to reduce overfitting and over-estimation of resolution (Chen et al., 2009; Grigorieff, 2007; Sindelar and Grigorieff, 2012; Stewart and Grigorieff, 2004), leading to a refinement performance similar to RELION. However, unlike RELION, FREALIGN cannot currently be used for classification of heterogeneous datasets. Here, we describe an extension of FREALIGN that allows classification using an ML approach that is related to that described previously (Scheres, 2012a). The new algorithm works together with the previously implemented FREALIGN refinement scheme and is therefore computationally more efficient than implementations that are entirely based on ML. We apply the new algorithm to a computer-generated dataset of 70S ribosomes, as well as an experimental (Baxter et al., 2009) dataset, both of which contain a heterogeneous mixture of conformationally and compositionally variable single particles.

2. Theory

Our new algorithm is an ML procedure: we seek to find the values of our model parameters that maximize the probability of our data. We will first describe a simple likelihood model for EM data from a single structure, connecting our data model to earlier work, and then expand the model to include data generated from mixtures of multiple different structures.

2a. A Gaussian statistical data model

Following seminal work by Sigworth (Sigworth, 1998), we assume a multivariate Gaussian model for our EM image data. We consider each image to be a two-dimensional (2D) projection of a randomly translated and rotated 3D structure, to which independent white Gaussian noise has been added to each pixel. This data model can be represented by:

$$\mathbf{X}_i = \mathbf{P}(\varphi_i, \mathbf{A}) + \sigma_i G_i \quad (1)$$

where $\mathbf{X}_i$ is the i th image in a dataset of N images, $\varphi_i = (\alpha, \beta, \gamma, x, y)_i$ is a set of transformation parameters (α, β, γ particle Euler angles and x, y coordinates) for image $\mathbf{X}_i$, $\mathbf{P}$ is the projection operator to produce a transformed 2D projection of $\mathbf{A}$ according the Euler angles, coordinates φ_i and determined CTF (we assume that image defocus is determined in an independent step using, for example the CTFIND3 software (Mindell and Grigorieff, 2003) that is not part of this formalism), σ_i is the standard deviation of the noise in the image, and G_i is a “noise image” of independent pixels with values sampled from a standard Gaussian distribution (zero mean and unit variance).

2b. The likelihood function

Given the data model in (1), the probability of observing a single image $\mathbf{X}_i$ is given by the Gaussian probability density function (PDF):

$$p(\mathbf{X}_i | \varphi_i, \Theta) = \left(\frac{1}{\sqrt{2\pi}\sigma_i} \right)^M \exp \left[-\frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{2\sigma_i^2} \right] \quad (2)$$

where M is the number of pixels in image $\mathbf{X}_i$ (in practice, M includes only pixels within an appropriate mask that is often applied to the particle image to reset image densities to a constant value outside the mask), $\|\mathbf{X}\|^2 = \mathbf{X}^T \mathbf{X}$ denotes the squared Frobenius norm of $\mathbf{X}$ (the squared inner product of $\mathbf{X}$, so that $\|\mathbf{X} - \mathbf{P}(\varphi_i, \mathbf{A})\|^2$ is the sum of squared differences between each pixel in $\mathbf{X}$ and the transformed $\mathbf{A}$), $\Theta = \{\mathbf{A}, \varphi\}$ is a set of parameters associated with structure $\mathbf{A}$, and $\varphi = \{\varphi_1, \dots, \varphi_N\}$ is vector of φ_i for all images. Here Θ trivially contains only the true structure $\mathbf{A}$ and φ , but Θ will be expanded later with more complex models. The notation for the conditional probability $p(x|y)$ is read as “the probability of x , given parameter y ”, and both x and y may be scalars, vectors, or matrices.

Because our error model assumes statistical independence of pixels and images, the joint PDF for a set of images is simply given by the product of PDFs for each individual image:

$$p(\mathbf{X} | \varphi, \Theta) = \prod_i^N p(\mathbf{X}_i | \varphi_i, \Theta) \quad (3)$$

$$= \left(\frac{1}{\sqrt{2\pi}} \right)^{MN} \prod_i^N \sigma_i^{-M} \exp \left[-\frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{2\sigma_i^2} \right] \quad (4)$$

where now $\{\mathbf{X}_i\}$ are each vectors of length N .

When we have fixed, observed data $\mathbf{X}_i$, and the parameters Θ in Eq. (3) are unknown, the probability $p(\mathbf{X} | \Theta)$ is a function only of the parameters, and Eq. (3) is called the *likelihood* of the parameters. The term “likelihood function” is used to emphasize that in the likelihood function the variables are the parameters, in contrast to the PDF in which the data values are the variables. In the method of maximum likelihood, the objective is to choose the

parameter values that maximize the likelihood, i.e. we find the parameter values that assign the highest possible probability to our observed data. Maximizing the likelihood over the parameters is equivalent to maximizing the log of the likelihood (Eq. (3)), which is usually more convenient to work with. The log-likelihood for Eq. (3) is:

$$\ell(\Theta, \varphi | \mathbf{X}) = -\frac{MN}{2} \ln(2\pi) - \frac{1}{2} \sum_i^N \frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{\sigma_i^2} \quad (5)$$

Often the maximum likelihood estimate of a parameter can be found analytically by taking the first derivative of the log-likelihood with respect to the parameter, setting it to zero, and solving for the unknown parameter.

2c. Cross-correlation maximization and least squares as ML

After some algebraic rearrangement, it can be shown that maximizing the log-likelihood (5) over φ and $\mathbf{A}$ with a constant σ_i^2 is equivalent to maximizing either

$$\sum_i^N \sum_j^N \left[\mathbf{R}(\varphi_i, \mathbf{X}_i)^T \mathbf{R}(\varphi_j, \mathbf{X}_j) \right] \quad (6)$$

or

$$\sum_i^N \left[\mathbf{R}(\varphi_i, \mathbf{X}_i)^T \hat{\mathbf{A}} \right] \quad (7)$$

where

$$\hat{\mathbf{A}} = \sum_i^N \mathbf{R}(\varphi_i, \mathbf{X}_i) \quad (8)$$

and $\mathbf{R}(\varphi, \mathbf{X})$ is a “back-projection” operator that puts the 2D image back into 3D. These equations can be recognized as restatements of the classical least-squares solution to finding the optimal image registration (Frank et al., 1988), based on maximizing cross-correlations. Thus, when assuming a simple Gaussian model with a common σ_i^2 for all pixels and images, ML is equivalent to least-squares. If each image is further assumed to have its own respective σ_i^2 , then the sums in Eqns. (6) – (8) are simply weighted by the inverse of the σ_i^2 , a procedure equivalent to weighted least-squares.

2d. Nuisance parameters and hierarchical priors

In most scientific estimation problems, certain parameters are of central interest while other parameters are only used as a means to an end. In statistics these “uninteresting” parameters are called *nuisance parameters*. Sigworth (Sigworth, 1998) treated the EM transformation variables (e.g., the rotations and translations that align images) as nuisance parameters, since they are only used transiently to estimate the parameter of interest, the reference structure. When SNR is low, the estimates of nuisance parameters can be highly uncertain. Since ultimately we do not care about the particular values of nuisance parameters, it would be useful to somehow account for, and perhaps mitigate, the uncertainty in their values.

A general statistical method for dealing with nuisance parameters is to treat them as random variables with their own PDF. For example, Sigworth recognized that the image transformation variables could be considered to be random variables themselves, and he proposed a bivariate Gaussian distribution for the x, y coordinate (translation) transformation variables:

$$p(\varphi_i | \sigma_x, \sigma_y, \hat{x}, \hat{y}) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left[-\frac{\|x_i - \hat{x}\|^2}{2\sigma_x^2} - \frac{\|y_i - \hat{y}\|^2}{2\sigma_y^2} \right] \quad (9)$$

where $\{x, y, \sigma_x, \sigma_y\}$ are the means and standard deviations of the x, y coordinates. Model parameters for the Euler angles can also be introduced if their distribution is non-uniform.

A PDF for parameters is called a *prior*. In this case Eq. (9) is specifically referred to as a *hierarchical prior*, since we now have a statistical model with a hierarchy of distributions — a PDF for the data, given certain parameters, supplemented by a higher level PDF for some of the parameters. The parameters of the hierarchical prior (e.g., σ_x and σ_y in Eq. (9)) may be called hierarchical parameters, to distinguish them from the parameters of the pure likelihood function.

Given a hierarchical prior for the φ_i parameters, the likelihoods in Eqs. (2) and (4) can then be augmented to construct an *extended likelihood* function $p(\mathbf{X}_i | \varphi_i)$ by multiplying the normal likelihood by the hierarchical prior:

$$p(\mathbf{X}_i, \varphi_i | \Theta) = p(\mathbf{X}_i | \Theta, \varphi_i) p(\varphi_i | \sigma_x, \sigma_y, \hat{x}, \hat{y}) \quad (10)$$

$$= \left(\frac{1}{\sqrt{2\pi}\sigma_i} \right)^M \exp \left[-\frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{2\sigma_i^2} \right] p(\varphi_i | \sigma_x, \sigma_y, \hat{x}, \hat{y}) \quad (11)$$

where the $\Theta = \{\mathbf{A}, \sigma_x, \sigma_y, \hat{x}, \hat{y}\}$ is the augmented set of all model parameters associated with reference structure $\mathbf{A}$. Note that Eqs. (10) and (11) correspond to 3D versions of Eqs. (11) and (12) of Sigworth, using our notation. The full hierarchical joint likelihood of a set of images is thus:

$$p(\mathbf{X}, \varphi | \Theta) = \left(\frac{1}{\sqrt{2\pi}} \right)^{MN} \prod_i^N \left(\sigma_i^{-M} \exp \left[-\frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{2\sigma_i^2} \right] p(\varphi_i | \sigma_x, \sigma_y, \hat{x}, \hat{y}) \right) \quad (12)$$

with corresponding log-likelihood:

$$\ln [p(\mathbf{X}, \varphi | \Theta)] = -\frac{MN}{2} \ln(2\pi) - \frac{1}{2} \sum_i^N \frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_i, \mathbf{A})\|^2}{\sigma_i^2} + \sum_i^N \ln [p(\varphi_i | \sigma_x, \sigma_y, \hat{x}, \hat{y})] \quad (13)$$

Other hierarchical priors can be added in a similar fashion to describe the distribution of other parameters, for example defocus (Chen et al., 2009) and magnification. The form of the additional distributions is often assumed to be Gaussian. Other authors may refer to an extended likelihood as a *regularized likelihood*, *penalized likelihood*, or a *hierarchical likelihood*. An extended likelihood as in Eq. (11) is also a *joint likelihood*, as it is equivalent to the joint PDF of the data and the parameters given the hyperparameters.

Given a hierarchical statistical model and a corresponding extended likelihood, there are several different ways to proceed with parameter estimation. When the hyperparameters of

the hierarchical prior distributions are estimated from the data using variants of ML methodology, such techniques are referred to as extended likelihood or *empirical Bayes*. There are two main ML variants: (a) to maximize the extended likelihood directly, and (b) to maximize the marginal likelihood, in which the nuisance parameters have been integrated out.

2e. Maximization of the joint extended likelihood

The extended likelihood can be maximized over all unknown parameters simultaneously, including both the parameters and the hyperparameters in the optimization. In practice, this is usually done using an iterative algorithm, in which each parameter is maximized in turn, conditional on the current optimal values of all other parameters. However, any multi-parameter optimization method may be used.

Maximization of the joint extended likelihood aims to find the joint point estimates of the “best” values for all parameters simultaneously. However, this method may not work well when the hierarchical prior is diffuse or multimodal. Direct maximization of the joint likelihood works best when the prior PDF for the hyperparameters is smooth and highly peaked.

2f. Maximization of the marginal likelihood

Alternatively, when the nuisance parameters are highly uncertain, it may be desirable to completely eliminate them from the analysis, while taking into account the uncertainty in their values. This is accomplished by integrating them out of the extended likelihood, resulting in a marginal likelihood function. For example, we can eliminate φ_i from the extended likelihood function in Eqs. (10) and (11) by integrating over its distribution:

$$p(\mathbf{X}_i|\Theta) = \int_{\varphi_i} p(\mathbf{X}_i|\Theta, \varphi_i) p(\varphi_i|\sigma_x, \sigma_y, \hat{x}, \hat{y}) d\varphi_i \quad (14)$$

which results in a marginal PDF that is independent of φ_i . This is the approach taken by Sigworth (Sigworth, 1998), where he integrates out the transformation parameters and maximizes the marginal likelihood function over $\mathbf{A}$ and $\mathbf{B}$.

In practice there are several choices for accomplishing the marginalization. In the simplest cases an analytical solution can be obtained. Usually we are not so lucky and must resort to numerical methods such as brute force integration, the Expectation-Maximization algorithm, or some combination of the two (Scheres, 2012a; Sigworth, 1998).

2g. Expectation-Maximization of the marginal likelihood

The Expectation-Maximization algorithm (normally abbreviated as EM, but we will avoid that here) finds the parameter values that maximize the marginal distribution using a mathematical trick that only requires the (non-integrated) joint likelihood. In its most general form, the algorithm cycles between two steps: (a) the “expectation step”, in which one finds the expected logarithm of the joint likelihood function, where the expectation is taken over the nuisance parameters (e.g., φ_i in Eq. (11)), conditional on the current values of the other parameters and the data, and (b) the “maximization step”, in which one maximizes the expected log likelihood function found in (a), as usual, over the other (non-nuisance) parameters of interest (e.g., $\mathbf{A}$ in Eq. (11)). While it may not be obvious why the Expectation-Maximization algorithm works, it can be shown that the algorithm increases the marginal likelihood at each step, and thus it is guaranteed to find a local maximum of the marginal likelihood.

Often the expectation in the first step can be determined analytically. However, in the present case no analytical solutions exist for the expectations of Eq. (13) needed to maximize the marginal likelihood in (14), and so they must be found by numerical integration, making the procedure computationally expensive (Scheres et al., 2005a; Scheres et al., 2005b; Sigworth, 1998). Expectation-Maximization of the marginal likelihood delivers superior performance compared to maximizing the joint likelihood in Eq. (12) when the likelihood function is multimodal or does not display clear peaks due to a low SNR. However, as explained below, in most practical cases we expect the likelihood function to exhibit a single peak close to the correct particle alignment parameters φ_i . If so, the model parameters that maximize the joint likelihood function (12) will also approximately maximize the marginalized likelihood (14).

2h. Extended likelihood, hierarchical priors, and Bayesian methods

The extended likelihood is reminiscent of Bayesian methodology, in which the likelihood is multiplied by a prior for each parameter. However, there is an extremely important methodological and ideological difference. In Bayesian methods, the values of the prior parameters are assumed constants, based on prior knowledge. In contrast, in the ML methods discussed here the hyperparameters of the prior are unknown parameters that ultimately are estimated from the data via a maximization procedure.

A related Bayesian approach seeks the model parameters with the maximum probability given the data and any prior information or assumptions (known as maximum *a posteriori* probability or MAP). However, like all Bayesian methods, MAP requires explicit Bayesian prior distributions with fixed parameters that are independent of the observed data. Scheres (Scheres, 2012a) has described an extended likelihood procedure which includes a hierarchical zero-mean Gaussian prior on the reciprocal space voxels of each image (which can be considered a “smoothing” technique that regularizes the voxel values). While Scheres describes this method as MAP, the parameters of his hierarchical prior are in fact estimated from the data by maximizing the extended likelihood function, similar to Sigworth’s method (Sigworth, 1998) and others, including our ML classification method discussed below. Hence, Scheres’ methodology is also a variant of empirical Bayes, which is purely likelihood based.

2i. Mixtures of particles with distinct structures

We now extend this EM data model to the case where the sample contains a heterogeneous mixture of structures in different, distinct conformations. We assume that there are K classes of structures that may be observed in a data set of images. We imagine that the molecule under consideration adopts a particular conformation $\mathbf{A}_k$ with probability p_k and then an image of this structure is generated according to the data model presented in Eq. (1). Each particular class will have its own respective set of parameters, e.g. φ_{ik} , and σ_{ik} (however, note that we presently assume that a given image has the same φ_i regardless of which class it is in). Each image also has an additional integer “indicator” parameter z_i that holds the value of the index of the structural class $\mathbf{A}_k$ that generated the image. For example, if the i th image belongs to class 1, then $z_i = 1$. The likelihood of an individual image, given that it is from class k , is thus:

$$p(\mathbf{X}_i, \varphi_{ik} | \Theta_k, z_i = k) = \left(\frac{1}{\sqrt{2\pi}\sigma_i} \right)^M \exp \left[-\frac{\|\mathbf{X}_i - \mathbf{P}(\varphi_{ik}, \mathbf{A}_k)\|^2}{2\sigma_i^2} \right] p(\varphi_{ik} | \Theta_k, z_i = k) \quad (15)$$

where

$$p(\varphi_{ik}|\Theta_k, z_i=k) = \frac{1}{2\pi\sigma_{xk}\sigma_{yk}} \exp \left[-\frac{\|x_{ik}-\hat{x}_k\|^2}{2\sigma_{xk}^2} - \frac{\|y_{ik}-\hat{y}_k\|^2}{2\sigma_{yk}^2} \right]. \quad (16)$$

Initially, imagine that we have a set of N data images $\mathbf{X}_i$ from K different classes, and that we know to which class each image belongs (i.e., we know the value of z_i). The joint probability of a single image being from class k and having a specific orientation, translation, and pixel values, is:

$$p(z_i=k, \mathbf{X}_i, \varphi_{ik}|\Theta_k) = p(\mathbf{X}_i, \varphi_{ik}|\Theta_k, z_i=k) p(z_i=k) \quad (17)$$

where it can be seen that

$$p(z_i=k) = \pi_k \quad (18)$$

The total likelihood of this set of images is:

$$p(z=k, \mathbf{X}, \varphi|\Theta) = \prod_i^N p(z_i=k, \mathbf{X}_i, \varphi_{ik}|\Theta_k) \quad (19)$$

Note that we can equivalently write Eq. (19) as

$$p(z, \mathbf{X}, \varphi|\Theta) = \prod_i^N \prod_k^K p(z_i=k, \mathbf{X}_i, \varphi_{ik}|\Theta_k)^{I(z_i=k)} \quad (20)$$

where $I(z_i=k)$ in the exponent is an indicator function whose value is 1 if $z_i=k$ and is 0 otherwise. If the image does not belong to class k , then $I(z_i=k)=0$ and the total product in Eq. (20) is left unchanged. Careful inspection will verify that Eq. (20) reduces to Eq. (19) when all $I(z_i=k)$ are known integers. However, the central problem of classification is that we do *not* know to which class each image belongs.

2j. Expectation-Maximization for classifying particle mixtures

For a mixture of different particles, the joint log-likelihood is thus given by the logarithm of Eq. (20):

$$\ell(\Theta, z, \varphi|\mathbf{X}) = \sum_{i=1}^N \sum_{k=1}^K I(z_i=k) \ln [\pi_k p(z_i=k, \mathbf{X}_i, \varphi_{ik}|\Theta_k)] \quad (21)$$

where Θ contains all parameters for all classes, including $\mathbf{A}_1 \dots \mathbf{A}_K$ that represent the underlying structures for each class. The log-likelihood in (21) is also a hierarchical likelihood, as π_k can be considered a hierarchical prior for both the z_i indicator values and the indicator function $I(z_i=k)$. Of course, in practice we do not know the values of z_i and $I(z_i=k)$, and so they are unknown parameters. Directly maximizing (21) is difficult, and the z_i may be considered to be nuisance parameters that are only used as a means to find the k reference structures. Hence, we would prefer to integrate them out and maximize the marginal likelihood $p(\mathbf{X}, \varphi|\Theta)$:

$$p(\mathbf{X}, \varphi|\Theta) = \int_z p(\mathbf{X}, \varphi, z|\Theta) dz \quad (22)$$

We cannot accomplish this integration analytically, but it is relatively easy to maximize the marginal likelihood in (22) with the Expectation-Maximization algorithm. We simply find the expectation of the joint log-likelihood (21) with respect to z , and maximize that instead.

A natural choice for the target function to maximize for the particle refinement with respect to model k is the probability distribution $p(\mathbf{X}_i | \Theta_k)$ (Eq. (2), (Sigworth, 1998)). The target function used in FREALIGN, which we will use in our present study, is a weighted correlation coefficient CC_w (Eq. (17) in Stewart and Grigorieff, 2004) modified with the hierarchical prior $p(\mathbf{X}_i | \Theta_k)$ in analogy to Eq. (21) in (Sigworth, 1998). Therefore (Chen et al., 2009),

$$\varphi_{ik}^m = \text{argmax} \left\{ \left\| \mathbf{X}_i \right\| \left\| P(\varphi_{ik}, \mathbf{A}_k) \right\| CC_w(\mathbf{X}_i, \Theta_k, \varphi_{ik}) + \sigma_i^2 \ln p(\varphi_{ik} | \Theta_k) \right\}. \quad (23)$$

2k. Algorithm

Classification is performed using Expectation-Maximization with fixed alignment parameters φ_{ik}^m found in the refinement. The expected probability of a particle i belonging to class k is given by

$$q_{ik} = p(z_i = k | \Theta, \mathbf{X}) = \frac{p(\mathbf{X}_i, \varphi_{ik}^m | \Theta_k) \pi_k}{\sum_{k=1}^K p(\mathbf{X}_i, \varphi_{ik}^m | \Theta_k) \pi_k} \quad (24)$$

which will be set to $1/K$ as the starting estimate. We will refer to these probabilities as particle occupancies. Other parameters are updated in each classification cycle as

$$\pi_k = \frac{1}{N} \sum_{i=1}^N q_{ik} \quad (25)$$

$$\mathbf{A}_k = \frac{\sum_i^N \frac{q_{ik}}{\sigma_i^2} R(\varphi_i, \mathbf{X}_i)}{\sum_{i=1}^N \frac{q_{ik}}{\sigma_i^2}} \quad (26)$$

$$\hat{x}_k = \frac{1}{\sum_{i=1}^N q_{ik}} \sum_{i=1}^N q_{ik} x_{ik}, \quad \hat{y}_k = \frac{1}{\sum_{i=1}^N q_{ik}} \sum_{i=1}^N q_{ik} y_{ik} \quad (27)$$

$$\sigma_{xk}^2 = \frac{1}{\sum_{i=1}^N q_{ik}} \sum_{i=1}^N q_{ik} (x_{ik} - \hat{x}_k)^2, \quad \sigma_{yk}^2 = \frac{1}{\sum_{i=1}^N q_{ik}} \sum_{i=1}^N q_{ik} (y_{ik} - \hat{y}_k)^2 \quad (28)$$

$$\sigma_i^2 = \frac{1}{M} \sum_{k=1}^K q_{ik} \|X_i - P(\varphi_i, \mathbf{A}_k)\|^2 \quad (29)$$

where R is the reconstruction operator used in FREALIGN (a Fourier inversion algorithm, Grigorieff, 2007). Rounds of non-ML parameter refinement to update φ_{ik} using Eq. (23) can be run between rounds of ML classification to benefit from the improved class averages.

In the initial stages of the refinement, a single reference can be used that represents the average of the mixture present in the dataset. For multi-reference refinement, the desired number of seeds can be generated either by calculating reconstructions from randomly sampled subsets of the data (Penczek et al., 2006; Spahn and Penczek, 2009), or by supplying reconstructions that represent known aspects of the conformations present in the dataset (supervised classification, van Heel and Stoffer-Meilicke, 1985). The number of required seeds is somewhat arbitrary, which is a recognized problem for K-means classification schemes like the one described here. Different numbers of seeds can be tested to find the largest number that produces classes with new features (Scheres, 2010).

2l. White noise assumption

In applying our ML method to EM data, we model the real space noise as independent, constant variance white Gaussian noise (Sigworth, 1998). However, it is well known that CTF effects and other factors affecting image amplitudes (envelopes) can potentially introduce strong correlations among neighboring pixels in real space, which could invalidate the white noise approximation. We make three major assumptions about the images that help validate the white noise model: 1) The contrast in an image is unrelated to the actual single particle density (signal) and is due to background generated by the sample support layer (ice and/or carbon), shot noise, and noise introduced by the detector/scanner (Baxter et al., 2009; Zeng et al., 2007). 2) Detector MTF effects are small in the resolution range of interest or can be corrected (Zeng et al., 2007). 3) Of the three sources of noise (i.e., background, shot noise, and detector noise), the shot noise is the most dominant, which should largely eliminate CTF induced correlations. This approximation appears to be valid for typical low-dose images of particles embedded in ice (Baxter et al., 2009), where shot noise is relatively large.

The effect of the white noise assumption can be gauged by inspecting the $M \times M$ Pearson correlation matrix $\mathbf{C}$ of the images. For white noise, the correlation matrix $\mathbf{C}$ is equivalent to the identity matrix, with ones on the diagonal and zero cross-correlations elsewhere. In principle, we could apply a more complicated Gaussian statistical model in which we assume an arbitrary correlation structure, where the correlations are estimated from the data simultaneously with the other model parameters. In this augmented statistical treatment, the likelihood is effectively multiplied by the reduced determinant of the correlation matrix $|\mathbf{C}|^{1/2M}$, which ranges from zero to one (corresponding to completely correlated data and non-correlated data, respectively). By ignoring correlations, then, one underestimates the actual sigma by a factor of $|\mathbf{C}|^{1/2M}$. In the case of the real ribosome data analyzed in this paper (see Results, Fig. 1a), the reduced determinant of the correlation matrix ranges from 0.916 to 0.928, depending on the defocus. These values are close to one, indicating largely uncorrelated data, and only a minor overestimation of the sigma, by roughly 7 – 9%. Taking all these factors into account, our white noise approximation appears appropriate.

2m. Justification for maximizing the joint likelihood instead of the marginal likelihood

As noted previously, the first exponential in Eq. (11) can usually be approximated by a delta function once the reconstructions $\mathbf{A}_k$ are reasonably close to their corresponding underlying

structures (Scheres, 2010). This will be true even if all reconstructions are set to a single reconstruction representing the average of the different underlying structures, provided the differences between the underlying structures are relatively small. The latter condition is fulfilled if the differences will not significantly affect the strong low-resolution signal (typically at 20 Å and lower), which follows the molecular transform of the particle and exceeds the signal at higher resolution usually by several orders of magnitude (Rosenthal and Henderson, 2003). We therefore propose a strategy in which initial refinement of a single reconstruction is performed using our established high-resolution (non-ML) protocols implemented in FREALIGN (Chen et al., 2009; Grigorieff, 2007) to a point where no further improvement is observed.

3. Implementation in FREALIGN

To implement the ML approach described in Section 2, we modified FREALIGN to calculate $\ln[p(\mathbf{X}_i, \mathbf{A}_k | \theta_k)]$ for each particle i and each class k . This can be done by running K instances of FREALIGN, each using its own set of alignment parameters θ_{ik} and model parameters $\mathbf{A}_k$. In a second step, particle occupancies q_{ik} and mixture frequencies θ_k are updated using Eq. (24) and (25). This second step is done using a small additional piece of software called CALC_OCC. This implementation therefore allows straightforward parallelization of the refinement of K classes. The computational load is proportional to the number of particles N and the number of classes K . Furthermore, because maximization of Eq. (23) is done per particle, only one reference reconstruction $\mathbf{A}_k$ and image $\mathbf{X}_i$ have to be kept in memory at a time. One cycle of alignment parameter refinement, classification and reconstruction of the experimental 70S dataset (see below; 10,000 particles with a box size of 130×130 pixels and four classes) took about 12 minutes of CPU time using an Intel Xeon X5690 3.47 GHz CPU. The memory requirements were about 550 MB.

4. Results

4a. Quantification of pixel correlations in experimental images

As explained in the Theory section, our implementation of the ML classification algorithm is based on the assumption that the noise present in different pixels of an image is uncorrelated. In practice, however, CTF modulations and various slowly varying functions (envelopes) describing the decay of image and noise amplitudes towards high resolution will introduce pixel correlations. To quantify the amount of pixel correlation and its spatial distribution, we calculated correlation matrices $\mathbf{C}$ for experimental cryo-EM images of 70S *E. coli* ribosome (Baxter et al., 2009) recorded at different defoci. We used sets of $N = 200$ images measuring 40×40 pixels that were excised from the ribosome images in areas that excluded signal from the particles. The images were recorded at 200 kV and digitized with a pixel size of 2.82 Å/pixel. The pair-wise pixel correlations were then calculated as

$$c_{mn} = \frac{\sum_i^N (x_{i,m} - \hat{x}_m)(x_{i,n} - \hat{x}_n)}{\sqrt{\sum_i^N (x_{i,m} - \hat{x}_m)^2 \sum_i^N (x_{i,n} - \hat{x}_n)^2}} \quad (30)$$

where c_{mn} are the elements of the correlation matrix $\mathbf{C}$, indicating the correlation coefficient between pixel locations m and n , i indicates the image in the set, $x_{i,m}$ are the pixel values and $\hat{x}_m$ are the pixel means calculated for the set. For white noise, the expectation value for c_{mn} is zero for $m \neq n$, leading to a correlation matrix that is similar to the identity matrix (except for random fluctuations affecting the correlation coefficients).

Fig. 1 displays a correlation matrix calculated using images with an underfocus of 2.8 μm . The matrix displays a pattern of lines running parallel to the diagonal and separated by 40 pixels. These lines are due to the conversion of the 2D images into one-column vectors which leads to a 40 pixel periodicity in physical pixel distances. Therefore, these lines describe correlations between nearest neighbors, second nearest neighbors and so on, that arise from the CTF modulations and amplitude envelope. Lines up to about the fourth nearest neighbor are discernible. The average nearest neighbor correlation (average correlations seen in the line nearest the diagonal) observed in Fig. 1a is about 0.28, dropping to 0.12 for the second nearest neighbor. In between the regularly spaced diagonals the correlation is much smaller and fluctuates around zero. Therefore, apart from a few off-diagonal lines the matrix approximates an identity matrix that would be expected for white noise. The reduced determinant of the correlation matrix in Fig. 1a is 0.924, close to 1, the value for an identity matrix. We also calculated reduced determinants for images with an underfocus of 2.1 μm and 3.5 μm , giving 0.928 and 0.916, respectively. Therefore, the observed correlation matrix suggests that our assumption of white noise is justified. Repeating the correlation analysis with image patches that contained part of ribosomes (signal) reproduced essentially the same results as in Fig. 1a. This is expected as these images were not aligned with each other and, therefore, the signal in different images is not correlated.

To show that the shot noise in the ribosome dataset is essentially white, we also calculated an average power spectrum of 857 ribosome particle images with a defocus of 2.8 μm (Fig. 1b) and plotted the rotational average in Fig. 1c. Except at very low resolution (below 100 $\text{\AA}$) where image contrast is affected by inelastic scattering, the noise power at the CTF zeros follows a slowly decaying function, again justifying our assumption of white noise. Background noise added by the embedding ice is affected by the CTF and has been estimated for this data to have about 70% of the variance of the signal produced by the ribosomes. Except at very low resolution, therefore, the CTF-modulated portion of the noise is only a small fraction of the total noise, consistent with the small effect the CTF modulations have on the correlation matrix (Fig. 1a).

The influence of the observed local correlations is further reduced in practical applications of our algorithm because they will most strongly affect the small, high-resolution details in the images. The SNR in a typical cryo-EM image of a single particle is lowest at high resolution, obscuring the fine details of the signal due to the particle. The correlations introduced by CTF and amplitude envelope will therefore affect the likelihood function (12) more or less independently of the signal. The small errors that result from these correlations will change the terms on the right side on Eq. (24) in a similar way, cancelling out in the calculation of the occupancies. The other equations (Eqs. (25) – (29)) used to update the model in each iteration are largely unaffected by the correlations, since they do not directly depend on the likelihood function.

4b. Classification of a simulated 70S ribosome dataset

To quantitatively assess the performance of our new algorithm, we used it to analyze computer-generated image data (Fig. 2) that mimic experimental data of heterogeneous ribosome populations (Baxter et al., 2009). We selected three previously determined maps of 70S ribosomes – EMD-1798 (Ratje et al., 2010) (Fig. 2a), EMD-1799 (Ratje et al., 2010) (Fig. 2b), and EMD-5030 (Schuette et al., 2009) (Fig. 2c). EMD-1798 and EMD-1799 both contain the elongation factor EF-G, but differ conformationally by a slight ratcheting of the 30S body and swiveling of the head subunit, the extent of which can be appreciated with their overlays in Fig. 2d. EMD-5030 differs from the previous two compositionally by the presence of EF-Tu in place of EF-G, as well as both an A-site and P-site tRNA. Detailed procedures for generating the simulated data are described in the Methods section. The

general sequence of steps was described previously (Baxter et al., 2009) and is shown in Fig. 3. It is intended to simulate the effects of background noise (an image noise component that is unrelated to the object, yet CTF-dependent), shot noise, and digitization noise.

We analyzed five different datasets containing 10,000 particle images each that differed in their SNR to assess the robustness of the FREALIGN ML approach, including SNR values significantly below those typically observed in an experiment. Previously, an SNR of ~0.05 was estimated for experimentally obtained 70S ribosomes (Baxter et al., 2009). Thus, the final SNRs for the five analyzed datasets were 0.010, 0.050, 0.025, 0.013, and 0.006, respectively (Fig. 3e–i). For each of the five datasets, three different classification schemes were performed: (1) classification only, whereby the Euler angles and shifts were fixed according to their true values (the values used to generate the projections); (2) particle alignment and classification, starting with the true Euler angles and shifts, but allowing changes during refinement; and (3) particle alignment and classification starting with perturbed Euler angles and shifts. The perturbations had a Gaussian distribution with standard deviations of 2.5° and 4 Å for Euler angles and shifts, respectively, and reduced the initial resolution of the starting 3D maps (calculated using the perturbed parameters) to 15 – 20 Å. This last test was intended to evaluate the ability of the algorithm to simultaneously refine particle alignment and classification (occupancy) parameters. This design enabled us to assess the performance of the algorithm under conditions whereby an increasingly large parameter space must be simultaneously searched during 3D classification.

For each SNR (Fig. 3e–i) and for each of the three classification schemes, 100 refinement iterations of FREALIGN were performed (Fig. 4). To assess the convergence of each run, we monitored the particle composition of the resulting classes after each iteration, as well as the mean occupancy change per particle per class. Plots depicting the variation of total particle compositions of the classes and mean particle occupancy changes with consecutive iterations, as well as values for the final particle composition of each output model with respect to the EMDB map from which they originate, are all displayed in Fig. 4. At an SNR of 0.100, the algorithm could readily recover the true classification parameters and reconstruct the correct starting maps with >99% accuracy, regardless of the classification scheme employed (Fig. 4a–c). At an SNR of 0.050, the algorithm accurately recovered the true classification parameters for each map with >95% certainty when the alignment parameters started with the correct values (Fig. 4d–e). However, starting from perturbed alignment parameters, the algorithm failed to simultaneously refine Euler angles, shifts, and classification parameters for the two smaller classes (Fig. 4f). When the alignment parameters were correct and did not require large adjustments, then even at a lower SNR of 0.025 it was possible to recover true classification parameters with reasonable accuracy (compare Fig. 4g–h with Fig. 4i), although the convergence behavior was somewhat more complicated (e.g. Fig. 4g). At still lower SNRs, recovering all three classes became increasingly difficult using our particular dataset and refinement strategy (see Methods) (Fig. 4j–o). While the SNR values used in these last two datasets were well below experimentally observed values (Baxter et al., 2009), the structure with the biggest difference compared with the other two structures (EMD-5030) could still be recovered (Fig. 4j–m).

We repeated the simulation for all SNRs and refinement strategies with initial particle Euler angles and shifts that were determined *ab initio* using a common lines approach and class averages (see Methods). These simulations represent an experimental situation in which nothing is known about the structure of the particle or their alignments. Results are shown in Supplementary Fig. 1. Both the classification-only and alignment and classification approaches performed as good (SNRs 0.100, 0.013, and 0.006) or better (SNRs 0.050 and

0.025) than in the case where perturbed Euler angles and shifts were supplied (Supplementary Fig. 1, compare columns 1/2 with 3).

4c. Classification of an experimental 70S ribosome dataset

As a second test, we used a publicly available dataset (Baxter et al., 2009) consisting of 10,000 images of 70S *E. coli* ribosomes with bound tRNAs and EF-G co-factor. This dataset was also used to test RELION (Scheres, 2012a; Scheres, 2012b) which produced three distinct classes, one class containing 70S bound to EF-G and a single tRNA in the E site (about 20 Å resolution), the second class containing 70S bound to three tRNAs and no EF-G (again, about 20 Å resolution), and a small third class (about 7% of the particles) containing the 50S large ribosomal subunit (about 30 Å resolution). The first class (70S with tRNA and EF-G) was duplicated in the analysis which used four classes ($K = 4$, Eq. (15)). Fig. 5 shows the corresponding analysis using the new FREALIGN algorithm, again assuming $K = 4$ classes. For the processing with FREALIGN, particles were masked with a circular mask with a radius of 142 Å. We performed 35 cycles of alignment including data out to 17 Å, followed by 65 cycles of ML classification at 16 Å resolution.

Our analysis detected four distinct classes – three that are equivalent to those detected by RELION and one additional class containing 70S bound to two tRNAs (in the A and P sites) and a weaker density in the E site. The partial presence of E-site density suggests that this additional class is not entirely homogeneous but still contains a mixture of particles. While the 70S-EF-G bound class contains about 50% of the particles in agreement with the previous analyses, the 70S-EF-G devoid class appears to contain 70S complexes with two and three tRNAs at roughly equal proportion. The resolution of all classes except the small 50S class was estimated to be about 15 Å which we confirmed by comparison with an atomic model (Fig. 5). For the 50S class (10% of the particles), we obtained a resolution of only about 40 Å. The resolution of this class is clearly limited by the number of member particles.

5. Discussion

Structural heterogeneity is found in most macromolecular complexes, albeit at varying degrees. While dynamic machines such as the spliceosome exhibit conformational and compositional variability that has prevented reconstruction of most of the splicing intermediates at a resolution higher than about 15 Å (Luhrmann and Stark, 2009), the variability in ribosomes is better understood and can be controlled more easily through careful biochemistry (Frank and Gonzalez, 2010). The highest resolution reconstructions are currently being obtained for icosahedral viruses (Grigorieff and Harrison, 2011), implying low variability compared with many asymmetric particles. The reduced variability is presumably partly the result of the viruses' symmetrical architecture which must restrain the conformational freedom of the subunits within these assemblies. To achieve similarly high resolution with asymmetrical assemblies, computational classification has to be applied to accommodate heterogeneity that cannot be reduced further through biochemical means (Bai et al., 2013). In some cases, such as the ribosome, analysis of the conformational variability may also lead to a deeper mechanistic understanding (Fischer et al., 2010; Mulder et al., 2010).

The presence of particles with multiple conformations or compositions increases the size of the dataset required to resolve each particle class at a given resolution. Therefore, although near-atomic resolution can now be achieved with only a few tens of thousands of particle images (Bai et al., 2013), a mixture of four or five different particle conformations would require four to five times as many images. This number increases still further if some of the classes are much smaller than others, and therefore, more particles need to be imaged to

obtain a sufficient number for the smallest class to be reconstructed at the desired resolution. The new algorithm implemented in FREALIGN combines computational speed with the superior convergence behavior of a maximum likelihood approach. Computational efficiency will help make the processing of datasets containing hundreds of thousands of particle images more feasible. FREALIGN is therefore ideally suited for refinement, classification and 3D reconstruction of large heterogeneous datasets at high resolution. Its performance in our test on a small ribosome dataset (10,000 particle images) is on par with that of RELION (Scheres, 2012a; Scheres, 2012b), as demonstrated by the results shown in Fig. 2.

In addition to the three classes obtained by RELION, a fourth class was detected that contains no EF-G and only two tRNAs (in the A and P sites). This fourth class is closely related to the class containing three tRNAs as it still shows some weak density in the third tRNA position. While our new algorithm was able to detect this fourth class, it was not capable of completely separating the particles between all the classes. This is presumably due to the relatively small difference between the two- and three-tRNA containing classes. Our new algorithm did not perform as well as RELION on the smallest class present in this dataset – the 50S subunit. The quality of the 50S map at about 40 Å is clearly worse than the 30 Å map obtained using RELION. In terms of its total signal in the dataset, the 50S class is the weakest due to the smallest number of class members and the lower molecular mass of 50S compared with the 70S particles. Therefore, in the limit of very low SNR, maximization of the marginal likelihood as implemented in RELION shows superior convergence. RELION also performs a more comprehensive and computationally expensive orientation search with each iteration that will further increase its success in correctly aligning 50S particles, which cannot be aligned correctly in the initial stages of the refinement when the reference structures represent mostly 70S particles.

6. Methods

6a. Generation of simulated data

We selected three previously determined cryo-EM maps of 70S ribosomes – EMD-1798 (Ratje et al., 2010) (Fig. 2a), EMD-1799 (Ratje et al., 2010) (Fig. 2b), and EMD-5030 (Schuette et al., 2009) (Fig. 2c). To reduce noise in these maps, binary masks enveloping each volume were generated, then Gaussian low-pass filtered to soften the edges, and multiplied by the corresponding map using the IMAGIC-5 suite (van Heel et al., 1996). Map EMD-1798 was low-pass filtered at 8 Å resolution using a Gaussian filter, and the other two maps were amplitude scaled against the filtered EMD-1798 map using DIFFMAP (<http://emlab.rose2.brandeis.edu/diffmap>). All subsequent procedures were automatically performed using the —create synthetic dataset functionality implemented within Appion (Lander et al., 2009), the details of which are described below. Projections using random Euler angles and shifts were generated using the project3d and proc2d functions of EMAN (Ludtke et al., 1999). We selected a pixel size of 2.52 Å/pixel, a box size of 160 × 160 pixels and assumed an acceleration voltage of 200 kV. 2,000, 3,000, and 5,000 projections of EMD-1798, EMD-1799, and EMD-5030, respectively were randomly distributed within a 10,000 particle dataset (Fig. 3a). Gaussian distributed noise was added using proc2d (Fig. 3b) to achieve an SNR (variance ratio of signal and noise) of 1.400 (Baxter et al., 2009). The initial addition of noise was then followed by a CTF that was randomly applied to each particle in the range of 2 – 4 µm underfocus using ace2correct (a variation of ACE1, (Mallick et al., 2005)) (Fig. 3c), and an experimentally obtained (Voss et al., 2010) envelope function (Fig. 3d). Finally the noise was brought down to five different levels based on the adjusted mean and standard deviation values of the dataset after CTF and envelope application, corresponding to a final SNR of 0.010, 0.050, 0.025, 0.013, and 0.006,

respectively (Fig. 3e–i). All particles were normalized to mean and standard deviation values of 0 and 1, respectively, prior to alignment parameter refinement and classification.

6b. Alignment parameter refinement and 3D classification of simulated data using FREALIGN

For each final level of noise (Fig. 3e–i) and for each of the three classification schemes, 100 iterations of FREALIGN were performed (Fig. 4). Each run initiated with randomized classification parameters that were obtained with the RSAMPLE program (distributed with FREALIGN). This produced three maps at iteration 0 that contained an approximately equal and random particle occupancy distribution, such that any differences between them were random. The resolution during refinement and classification was always limited to 20 Å. A mask, corresponding to a particle radius of 160 Å was applied and changes in occupancy between iterations were reduced to 50% to improve convergence. Otherwise, the refinement followed standard procedures in FREALIGN (Grigorieff, 1998; Grigorieff, 2007).

6c. Determination of *ab initio* Euler angle and shifts for 3D classification of the synthetic 70S ribosome dataset

To determine the Euler angles from scratch we utilized the following strategy. First, we used the 0.025 SNR dataset (due to its intermediate level of noise) that contains a mixture of emd-1798, emd-1799, and emd-5030 particles in randomized order to generate reference-free class averages using the ISAC method (Yang et al., 2012). This resulted in 265 — accounted class averages that were further processed using an automated procedure based on the angular reconstitution (common-lines) algorithm implemented in IMAGIC (van Heel et al., 1996). An initial model was obtained that was then used as a reference for projection-matching using Xmipp (Scheres et al., 2008). The Euler angles from the last iteration of projection-matching were converted to FREALIGN, followed by 5 cycles of alignment parameter refinement. The final angles from this procedure were used as starting points for iterative classification as described above.

6d. Estimating the resolution of a 3D reconstruction using atomic models

To obtain an unbiased resolution estimate of one of the 3D reconstructions obtained from the experimental 70S dataset (Baxter et al., 2009), we modeled the reconstruction using atomic models for the 30S subunit and three tRNAs (PDB code 2gy9), and 50S subunit (PDB code 2gya). The models were aligned with the 3D reconstruction from the test dataset using UCSF Chimera (Pettersen et al., 2004), converted into a single density map using Bsoft (Heymann and Belnap, 2007). The 3D reconstruction was masked using an envelope derived from a 60-Å low-pass filtered version of the map to remove noise from the solvent surrounding the particle. An FSC curve (FSC_{full}) was then calculated between the masked map and map generated by Bsoft using EMAN's proc3d program (Ludtke et al., 1999). To generate a curve that corresponds to an FSC curve calculated between two half data sets (FSC_{half}), we converted the curve calculated between the model and the map using

$$FSC_{half} = \frac{FSC_{full}^2}{2 - FSC_{full}^2}. \quad (31)$$

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are grateful to the Frank group for providing their ribosome test dataset to the cryo-EM community. The work was supported by National Institutes of Health Grants P01 GM62580 (awarded to N.G.), R01 GM094468 (awarded to D.L.T.), P41 GM103310 (supporting D.L.) and an NSERC Fellowship (awarded to A.B.). The synthetic ribosome data set used in this work can be downloaded at <http://nramm.scripps.edu>. The new version of FREALIGN is available for download at <http://emlab.rose2.brandeis.edu/>.

References

- Adrian M, Dubochet J, Lepault J, McDowell AW. Cryo-electron microscopy of viruses. *Nature*. 1984; 308:32–36. [PubMed: 6322001]
- Bai XC, Fernandez IS, McMullan G, Scheres SH. Ribosome structures to near-atomic resolution from thirty thousand cryo-EM particles. *eLife*. 2013; 2:e00461. [PubMed: 23427024]
- Baxter WT, Grassucci RA, Gao H, Frank J. Determination of signal-to-noise ratios and spectral SNRs in cryo-EM low-dose imaging of molecules. *J Struct Biol*. 2009; 166:126–132. [PubMed: 19269332]
- Brilot AF, Chen JZ, Cheng A, Pan J, Harrison SC, Potter CS, Carragher B, Henderson R, Grigorieff N. Beam-induced motion of vitrified specimen on holey carbon film. *J Struct Biol*. 2012; 177:630–637. [PubMed: 22366277]
- Campbell MG, Cheng A, Brilot AF, Moeller A, Lyumkis D, Veisler D, Pan J, Harrison SC, Potter CS, Carragher B, Grigorieff N. Movies of ice-embedded particles enhance resolution in electron cryo-microscopy. *Structure*. 2012; 20:1823–1828. [PubMed: 23022349]
- Chen JZ, Settembre EC, Aoki ST, Zhang X, Bellamy AR, Dormitzer PR, Harrison SC, Grigorieff N. Molecular interactions in rotavirus assembly and uncoating seen by high-resolution cryo-EM. *Proc Natl Acad Sci U S A*. 2009; 106:10644–10648. [PubMed: 19487668]
- Doerschuk PC, Johnson JE. Ab initio reconstruction and experimental design for cryo electron microscopy. *Ieee Transactions on Information Theory*. 2000; 46:1714–1729.
- Elad N, Clare DK, Saibil HR, Orlova EV. Detection and separation of heterogeneity in molecular complexes by statistical analysis of their two-dimensional projections. *J Struct Biol*. 2008; 162:108–120. [PubMed: 18166488]
- Elmlund D, Davis R, Elmlund H. Ab initio structure determination from electron microscopic images of single molecules coexisting in different functional states. *Structure*. 2010; 18:777–786. [PubMed: 20637414]
- Fischer N, Konevega AL, Wintermeyer W, Rodnina MV, Stark H. Ribosome dynamics and tRNA movement by time-resolved electron cryomicroscopy. *Nature*. 2010; 466:329–333. [PubMed: 20631791]
- Frank J, Gonzalez RL Jr. Structure and dynamics of a processive Brownian motor: the translating ribosome. *Annual review of biochemistry*. 2010; 79:381–412.
- Frank J, Radermacher M, Wagenknecht T, Verschoor A. Studying ribosome structure by electron microscopy and computer-image processing. *Methods in enzymology*. 1988; 164:3–35. [PubMed: 3071670]
- Grigorieff N. Three-dimensional structure of bovine NADH:ubiquinone oxidoreductase (complex I) at 22 Å in ice. *Journal of molecular biology*. 1998; 277:1033–1046. [PubMed: 9571020]
- Grigorieff N. FREALIGN: high-resolution refinement of single particle structures. *J Struct Biol*. 2007; 157:117–125. [PubMed: 16828314]
- Grigorieff N, Harrison SC. Near-atomic resolution reconstructions of icosahedral viruses from electron cryo-microscopy. *Current opinion in structural biology*. 2011; 21:265–273. [PubMed: 21333526]
- Heymann JB, Belnap DM. Bsoft: image processing and molecular modeling for electron microscopy. *J Struct Biol*. 2007; 157:3–18. [PubMed: 17011211]
- Lander GC, Stagg SM, Voss NR, Cheng A, Fellmann D, Pulokas J, Yoshioka C, Irving C, Mulder A, Lau PW, Lyumkis D, Potter CS, Carragher B. Appion: an integrated, database-driven pipeline to facilitate EM image processing. *J Struct Biol*. 2009; 166:95–102. [PubMed: 19263523]
- Lawson CL. Unified data resource for cryo-EM. *Methods in enzymology*. 2010; 483:73–90. [PubMed: 20888470]

- Lawson CL, Baker ML, Best C, Bi C, Dougherty M, Feng P, van Ginkel G, Devkota B, Lagerstedt I, Ludtke SJ, Newman RH, Oldfield TJ, Rees I, Sahni G, Sala R, Velankar S, Warren J, Westbrook JD, Henrick K, Kleywegt GJ, Berman HM, Chiu W. EMDataBank.org: unified data resource for CryoEM. *Nucleic acids research*. 2011; 39:D456–464. [PubMed: 20935055]
- Liao HY, Frank J. Classification by Bootstrapping in Single Particle Methods. *Proceedings / IEEE International Symposium on Biomedical Imaging: from nano to macro. IEEE International Symposium on Biomedical Imaging*. 2010; 2010:169–172.
- Ludtke SJ, Baldwin PR, Chiu W. EMAN: semiautomated software for high-resolution single-particle reconstructions. *J Struct Biol*. 1999; 128:82–97. [PubMed: 10600563]
- Luhrmann R, Stark H. Structural mapping of spliceosomes by electron microscopy. *Current opinion in structural biology*. 2009; 19:96–102. [PubMed: 19211241]
- Mallick SP, Carragher B, Potter CS, Kriegman DJ. ACE: automated CTF estimation. *Ultramicroscopy*. 2005; 104:8–29. [PubMed: 15935913]
- Mindell JA, Grigorieff N. Accurate determination of local defocus and specimen tilt in electron microscopy. *J Struct Biol*. 2003; 142:334–347. [PubMed: 12781660]
- Mulder AM, Yoshioka C, Beck AH, Bunner AE, Milligan RA, Potter CS, Carragher B, Williamson JR. Visualizing ribosome biogenesis: parallel assembly pathways for the 30S subunit. *Science*. 2010; 330:673–677. [PubMed: 21030658]
- Penczek PA, Yang C, Frank J, Spahn CM. Estimation of variance in single-particle reconstruction using the bootstrap technique. *J Struct Biol*. 2006; 154:168–183. [PubMed: 16510296]
- Pettersen EF, Goddard TD, Huang CC, Couch GS, Greenblatt DM, Meng EC, Ferrin TE. UCSF Chimera—a visualization system for exploratory research and analysis. *Journal of computational chemistry*. 2004; 25:1605–1612. [PubMed: 15264254]
- Provencher SW, Vogel RH. Three-dimensional reconstruction from electron micrographs of disordered specimens. I. Method. *Ultramicroscopy*. 1988; 25:209–221. [PubMed: 3212837]
- Ratje AH, Loerke J, Mikolajka A, Brunner M, Hildebrand PW, Starosta AL, Donhofer A, Connell SR, Fucini P, Mielke T, Whitford PC, Onuchic JN, Yu Y, Sanbonmatsu KY, Hartmann RK, Penczek PA, Wilson DN, Spahn CM. Head swivel on the ribosome facilitates translocation by means of intra-subunit tRNA hybrid sites. *Nature*. 2010; 468:713–716. [PubMed: 21124459]
- Rosenthal PB, Henderson R. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *J Mol Biol*. 2003; 333:721–745. [PubMed: 14568533]
- Scheres SH. Classification of structural heterogeneity by maximum-likelihood methods. *Methods Enzymol*. 2010; 482:295–320. [PubMed: 20888966]
- Scheres SH. A Bayesian view on cryo-EM structure determination. *J Mol Biol*. 2012a; 415:406–418. [PubMed: 22100448]
- Scheres SH. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *J Struct Biol*. 2012b; 180:519–530. [PubMed: 23000701]
- Scheres SH, Valle M, Carazo JM. Fast maximum-likelihood refinement of electron microscopy images. *Bioinformatics*. 2005a; 21(Suppl 2):ii243–244. [PubMed: 16204112]
- Scheres SH, Nunez-Ramirez R, Sorzano CO, Carazo JM, Marabini R. Image processing for electron microscopy single-particle analysis using XMIPP. *Nat Protoc*. 2008; 3:977–990. [PubMed: 18536645]
- Scheres SH, Valle M, Nunez R, Sorzano CO, Marabini R, Herman GT, Carazo JM. Maximum-likelihood multi-reference refinement for electron microscopy images. *J Mol Biol*. 2005b; 348:139–149. [PubMed: 15808859]
- Schuette JC, Murphy FVt, Kelley AC, Weir JR, Giesebrecht J, Connell SR, Loerke J, Mielke T, Zhang W, Penczek PA, Ramakrishnan V, Spahn CM. GTPase activation of elongation factor EF-Tu by the ribosome during decoding. *The EMBO journal*. 2009; 28:755–765. [PubMed: 19229291]
- Shatsky M, Hall RJ, Nogales E, Malik J, Brenner SE. Automated multi-model reconstruction from single-particle electron microscopy data. *J Struct Biol*. 2010; 170:98–108. [PubMed: 20085819]
- Sigworth FJ. A maximum-likelihood approach to single-particle image refinement. *J Struct Biol*. 1998; 122:328–339. [PubMed: 9774537]

- Sigworth FJ, Doerschuk PC, Carazo JM, Scheres SH. An introduction to maximum-likelihood methods in cryo-EM. *Methods Enzymol.* 2010; 482:263–294. [PubMed: 20888965]
- Sindelar CV, Grigorieff N. Optimal noise reduction in 3D reconstructions of single particles using a volume-normalized filter. *J Struct Biol.* 2012; 180:26–38. [PubMed: 22613568]
- Spahn CM, Penczek PA. Exploring conformational modes of macromolecular assemblies by multiparticle cryo-EM. *Curr Opin Struct Biol.* 2009; 19:623–631. [PubMed: 19767196]
- Stewart A, Grigorieff N. Noise bias in the refinement of structures derived from single particles. *Ultramicroscopy.* 2004; 102:67–84. [PubMed: 15556702]
- Taylor KA, Glaeser RM. Electron diffraction of frozen, hydrated protein crystals. *Science.* 1974; 186:1036–1037. [PubMed: 4469695]
- Thon F. Zur Defokussierungsabhängigkeit des Phasenkontrastes bei der elektronenmikroskopischen Abbildung. *Zeitschrift für Naturforschung.* 1966; 21a:476–478.
- van Heel M, Stoffler-Meilicke M. Characteristic views of *E. coli* and *B. stearothermophilus* 30S ribosomal subunits in the electron microscope. *The EMBO journal.* 1985; 4:2389–2395. [PubMed: 3908096]
- van Heel M, Harauz G, Orlova EV, Schmidt R, Schatz M. A new generation of the IMAGIC image processing system. *J Struct Biol.* 1996; 116:17–24. [PubMed: 8742718]
- Vogel RH, Provencher SW. Three-dimensional reconstruction from electron micrographs of disordered specimens. II. Implementation and results. *Ultramicroscopy.* 1988; 25:223–239. [PubMed: 3212838]
- Voss NR, Lyumkis D, Cheng A, Lau PW, Mulder A, Lander GC, Brignole EJ, Fellmann D, Irving C, Jacovetty EL, Leung A, Pulokas J, Quispe JD, Winkler H, Yoshioka C, Carragher B, Potter CS. A toolbox for ab initio 3-D reconstructions in single-particle electron microscopy. *J Struct Biol.* 2010; 169:389–398. [PubMed: 20018246]
- Yang Z, Fang J, Chittuluru J, Asturias FJ, Penczek PA. Iterative stable alignment and clustering of 2D transmission electron microscope images. *Structure.* 2012; 20:237–247. [PubMed: 22325773]
- Zeng X, Stahlberg H, Grigorieff N. A maximum likelihood approach to two-dimensional crystals. *J Struct Biol.* 2007; 160:362–374. [PubMed: 17964808]

Figure 1. Pixel correlation in a typical cryo-EM image

(a) Correlation matrix calculated using 40×40 pixel patches excised from 200 cryo-EM images of 70S ribosomes. The images were recorded at 200 kV with an underfocus of 2.8 μm and digitized with a pixel size of 2.82 $\text{\AA}/\text{pixel}$. Lines parallel to the diagonal are visible that indicate correlations between pixels that are close to each other in the image. The correlations are due to modulations of part of the noise by the contrast transfer function (CTF) of the microscope and other effects (envelopes) leading to non-white noise. (b) Thon ring pattern (Thon, 1966) calculated from 857 ribosome images, corresponding to the imaging conditions in (a). (c) Rotational average of the pattern in (b).

Figure 2. Models used to generate simulated data

Simulated 2D images were generated from three previously determined EM maps that were deposited to the EMDB (Lawson, 2010; Lawson et al., 2011): (a–c) (top) views from the inter-subunit cleft of the complete starting maps and (bottom) from the top of the ribosome with the 30S subunit computationally removed (a) EMD-1798 (50S subunit is blue, 30S subunit yellow, EF-G red, and P/E-site tRNA of the PRE state green), (b) EMD-1799 (50S subunit is blue, 30S subunit yellow, EF-G red, and P/E-site tRNA of the POST state green), and (c) EMD-5030 (50S subunit is blue, 30S subunit yellow, EF-Tu orange, A/T tRNA pink, P-site tRNA blue, and E-site tRNA green). (d) EMD-1798 (light grey) and EMD-1799 (dark grey) contain minor differences, the majority of which can be attributed to slight ratcheting of the 30S subunit with respect to the 50S subunit, and swiveling of the head when aligned to the 30S body. Arrows indicate the differences that are evident in surface rendered maps of the 30S subunit when the two are aligned to each other (top – view from the mRNA entry site; bottom – view from the 50S subunit). Density maps are represented as isosurfaces by UCSF Chimera (Pettersen et al., 2004).

Figure 3. Generation of simulated data

(a) Projections were randomly calculated from each of the three previously determined EM maps, then randomly shifted. (b) Noise was added to the combined dataset such that the SNR was brought down to 1.400, which was followed by (c) application of a CTF in the range of 2 – 4 μm underfocus, and (d) an experimentally determined envelope. Finally, a last layer of noise was added, bringing the final SNR down to (e) 0.100, (f) 0.050, (g) 0.025, (h) 0.013, and (i) 0.006. The datasets used for 3D parameter refinement and classification are outlined in red.

Figure 4. Classification of the simulated 70S ribosome dataset using $K = 3$ classes

Tests were performed using different levels of noise (arranged vertically) and different classification schemes (arranged horizontally). Each run was initiated with randomized classification parameters, producing three maps at iteration 0 that contained an approximately equal and random particle occupancy distribution, such that any differences among them was essentially random. Five different levels of noise were tested, corresponding to a final SNR of (a-c) 0.100, (d-f) 0.050, (g-i) 0.025, (j-l) 0.013, and (m-o) 0.006. For each SNR level, three independent runs of FREALIGN were performed starting with (a, d, g, j, m) the correct Euler angles and shifts and disabling their refinement, (b, e, h, k, n) the correct Euler angles and shifts and enabling parameter refinement, and (c, f, i, l, o)

perturbed Euler angles and shifts, such that the resolution of the starting models at iteration 0 was between 15 – 20 Å. For each plot, the solid lines (left y-axis plotted against the x-axis) represent the total particle occupancy for each iteration within output model 1 (red), model 2 (green), and model 3 (blue). Dotted lines (right y-axis plotted against the x-axis) represent the mean occupancy change per particle per model for each iteration with regard to output model 1 (red), model 2 (green), and model 3 (blue). Both sets of lines are plotted against 100 performed iterations. The tables above each plot describe the particle composition of each output model at iteration 100 (expressed as a percentage) with respect to the original map in the EMDB. Thus, for example, at a SNR of 0.025 and for the perturbed alignment parameters and classification run (panel i), final models 1 & 2 generated by FREALIGN are both primarily composed of particles that originated from EMD-5030 (>99%), whereas model 3 is composed of particles representing the remaining data, and this remainder is divided approximately in accordance with its original 60% / 40% distribution within the full 10,000 particle dataset.

Figure 5. Classification of the experimental 70S ribosome dataset using $K = 4$ classes

Density for EF-G is shown in red, the 50S and 30S subunits are shown in blue and yellow, respectively, and the tRNAs are shown in green. Views with transparent surfaces to show density due to the tRNAs and EF-G are shown in (a), and cut-way views at a higher density threshold to highlight internal density features are shown in (b). Fourier Shell Correlation (FSC) curves for each class are shown in (c) to indicate the resolution of each map. The first two classes represent 70S ribosomes with and without bound EF-G and were also obtained in previous analyses of this test dataset (Elad et al., 2008; Elmlund et al., 2010; Liao and Frank, 2010; Scheres, 2012a; Shatsky et al., 2010). For the second class, a second FSC curve is shown in red, indicating the resolution as estimated using atomic models (PDB codes 2gy9 and 2gya, see Methods). The last class represents the 50S subunit and was also observed in a classification performed by RELION (Scheres, 2012a; Scheres, 2012b). A new class that was not observed before is shown in third position from the left and represents a 70S ribosome with tRNAs bound in the A and P sites. There is weaker density for a third tRNA in the E site, indicating that this class still contains some “contaminating” particles from the second class from the left. Density maps are represented as isosurfaces by UCSF Chimera (Pettersen et al., 2004).