

Rapid Solution of the Cryo-EM Reconstruction Problem by Frequency Marching*

Alex Barnett[†], Leslie Greengard[‡], Andras Pataki[§], and Marina Spivak[§]

Abstract. Determining the three-dimensional (3D) structure of proteins and protein complexes at atomic resolution is a fundamental task in structural biology. Over the last decade, remarkable progress has been made using “single particle” cryo-electron microscopy (cryo-EM) for this purpose. In cryo-EM, hundreds of thousands of two-dimensional (2D) images are obtained of individual copies of the same particle, each held in a thin sheet of ice at some unknown orientation. Each image corresponds to the noisy projection of the particle’s electron-scattering density. The reconstruction of a high-resolution image from this data is typically formulated as a nonlinear, nonconvex optimization problem for unknowns which encode the angular pose and lateral offset of each particle. Since there are hundreds of thousands of such parameters, this leads to a very CPU-intensive task—limiting both the number of particle images which can be processed and the number of independent reconstructions which can be carried out for the purpose of statistical validation. Moreover, existing reconstruction methods typically require a good initial guess to converge. Here, we propose a deterministic method for high-resolution reconstruction that operates in an *ab initio* manner—that is, without the need for an initial guess. It requires a predictable and relatively modest amount of computational effort, by marching out radially in the Fourier domain from low to high frequency, increasing the resolution by a fixed increment at each step.

Key words. cryo-EM, single particle reconstruction, protein structure, recursive linearization, frequency marching

AMS subject classifications. 92C55, 94A08, 65K10

DOI. 10.1137/16M1097171

1. Introduction. Cryo-electron microscopy (cryo-EM) is an extremely powerful technology for determining the three-dimensional (3D) structure of individual proteins and macromolecular assemblies. Unlike X-ray diffraction-based methods, it does not require the synthesis of high quality crystals, although there are still significant sample preparation issues involved. Instead, hundreds of thousands of copies of the same particle are held frozen in a thin sheet of ice at unknown, random orientations. The sample is then placed in an electron microscope and, using a weak scattering approximation, the resulting two-dimensional (2D) transmission images are interpreted as noisy projections of the particle’s electron-scattering density, with some optical aberrations that need to be corrected. Much of the recent advance in imaging quality is a result of hardware improvements (direct electron detectors), motion

*Received by the editors October 5, 2016; accepted for publication (in revised form) May 30, 2017; published electronically August 1, 2017.

<http://www.siam.org/journals/siims/10-3/M109717.html>

[†]Flatiron Institute, Simons Foundation, New York, NY 10010, and Department of Mathematics, Dartmouth College, Hanover, NH 03755 (ahb@math.dartmouth.edu).

[‡]Flatiron Institute, Simons Foundation, New York, NY 10010, and Courant Institute, NYU, New York, NY 10012-1185 (lgreengard@simonsfoundation.org).

[§]Flatiron Institute, Simons Foundation, New York, NY 10010 (apataki@simonsfoundation.org, mspivak@simonsfoundation.org).

correction, and the development of new image reconstruction algorithms. Resolutions at the 2–4 Å level are now routinely achieved.

There are a number of excellent overviews and texts in the literature to which we refer the reader for a more detailed introduction to the subject [8, 32, 27, 30, 51, 52]. One important issue to keep in mind, however, is that the field is currently lacking a rigorous method for assessing the accuracy of the reconstructed density map, although there are a variety of best practices in use intended to avoid overfitting [8, 37, 19, 21].

In this paper, we concentrate on the reconstruction problem, that is, converting a large set of noisy experimental cryo-EM images into a 3D map of the electron-scattering density. As a rule, this is accomplished in two stages: the generation of a low-resolution initial guess, either experimentally or computationally, followed by an iterative refinement step seeking to achieve the full available resolution permitted by the data. It is generally the refinement step which is CPU-intensive. Although our *ab initio* method does not require an initial guess, we briefly review some of the existing approaches for the sake of completeness.

For methods requiring a good, low-resolution initial guess in which the user has not supplied one experimentally, many methods are based on the *common lines* principle. This exploits the projection-slice theorem (see Theorem 2.3 and Figure 1) and was first introduced in [50, 49]. In the absence of noise, it is possible to show that any three randomly oriented experimental images can be used to define a coordinate system (up to mirror image symmetry). Given that coordinate system, the angular assignment for all other images can be easily computed. Unfortunately, this scheme is very sensitive to noise. Significant improvement was achieved in the common lines approach in [35], where all angles are assigned simultaneously by minimizing a global error. However, this requires a very time-consuming calculation, as it involves searching in an exponentially large parameter space for all possible orientations of all projections. More recently, the SIMPLE algorithm [12] used various optimization strategies, such as simulated annealing and differential evolution optimizers, to accelerate this search process. Singer, Shkolnisky, and coworkers [45, 9, 46, 42] proposed a variety of rigorous methods based on the eigendecomposition of certain sparse matrices derived from common line analysis and/or convex relaxation of the global least-squares error function. Wang, Singer, and Wen [55] subsequently proposed a more robust self-consistency error, based on the sum of absolute residuals, which permits a convex relaxation and solution by semidefinite programming.

Another approach, introduced in the late 1980s, is the method of moments [14, 15], which is aimed at computing second order moments of the density from second order moments of the 2D experimental images. Around the same time, Provencher and Vogel [36, 53] proposed representing the electron-scattering density as a truncated expansion in orthonormal basis functions in spherical coordinates and estimating the maximum likelihood of this density. Recently, a machine-learning approach was suggested, using a sum of Gaussian “pseudo-atoms” with unknown locations and radii to represent the density, and Markov chain Monte Carlo sampling to estimate the model parameters [23].

Once a low-resolution initial guess is established, the standard reconstruction packages EMAN2 [2, 47], SPIDER [13], SPARX [22], RELION [39, 40], and FREALIGN [17, 18, 26] all use some variant of iterative projection matching in either physical or Fourier space. Some, like RELION, SPARX, and EMAN2, make use of soft matching or a regularized version of the maximum likelihood estimation (MLE) framework introduced by Sigworth [43, 44]. For this,

a probability distribution for angular assignments is determined for each experimental image. Others, like FREALIGN, assign unique angular assignments to each image. Two drawbacks of these methods are that they can be time-consuming, especially those based on MLE, and that it can be difficult to determine if and when a global optimum has been reached.

In order to accelerate the MLE-based methods, Dvornek, Sigworth, and Tagare suggested in [11] that particle images and structure projections be represented in low-dimensional subspaces that permit rotation, translation, and comparison by defining suitable operations on the subspace bases themselves. They demonstrated 300-fold speedups in reconstruction. Recently, Punjani, Brubaker, and Fleet [5] introduced a new scheme based on a probabilistic generative model, marginalization over angular assignments, and optimization using stochastic gradient descent and importance sampling. They demonstrated that their method was both efficient and insensitive to the initial guess. Finally, there has been significant effort aimed at harnessing high-performance computing hardware including GPUs for the most computation-intensive tasks in the cryo-EM reconstruction pipeline [25].

In this paper, we will focus on a new method for refinement, assuming that all particle images are drawn from a homogeneous population. One of our goals is the creation of a method for refinement that is sufficiently fast that it can be run multiple times on the same data, opening up classical jackknife, bootstrap, or cross-validation statistics to be used for validation and resolution assessment. At present, the gold standard is based on methods such as “Fourier shell correlation” [20, 33, 37], which can be viewed as a jackknife with two samples.

The basic intuition underlying our approach is already shared with many of the standard software packages, such as RELION and FREALIGN, as well as the stochastic optimization method of [5], namely, that the best path to a refined structure is achieved by *gradually increasing resolution*. The main purpose of the present paper is to propose a deterministic and mathematically precise version of this idea, carried out in the Fourier domain. Unlike existing schemes, it involves no global optimization. Instead, we use resolution (defined by the maximum frequency content of the current reconstruction) as a homotopy path. For each small step along that path, we solve only *uncoupled* projection matching problems for each experimental image. More precisely, let $k \in [0, K]$ denote the band-limit of the model at the current step of the reconstruction algorithm, where K denotes the maximum resolution we seek to achieve. For the model, we generate a large number of “templates”—simulated projections of the model—at a large number of orientations. For each experimental image, we then find the template that is a best match and assign the orientation of that template to the image. Given the current angular assignments of all experimental images, we solve a linear least-squares problem to build a *new* model at resolution $k + \Delta k$ and repeat the process until the maximum resolution K is reached (see section 6).

With M images, each at a resolution of $K \times K$ pixels, our scheme requires $O(MK^4)$ or $O(MK^5)$ work, depending on the cost of template-matching. If this is done by brute force, the second estimate applies. If a hierarchical but local search strategy is employed for template-matching, then the cost of the least-squares procedures dominates and the $O(MK^4)$ complexity is achieved. (The memory requirements are approximately $8MK^2 + 8K^3$ bytes.)

Remark 1. Some structures of interest have nontrivial point group symmetries. As a result, there may not be a unique angular assignment for each experimental image. In the

most extreme case, one could imagine imaging perfect spheres, for which angular assignment makes no sense at all. Preliminary experiments with noisy projection data have been successful in this case, using the randomized assignment scheme discussed in section 4.2. We believe that our procedure is easily modified to handle point group symmetries as well but have not explored this class of problems in detail (see [38] for further discussion).

Our frequency marching scheme was inspired by the method of recursive linearization for acoustic inverse scattering, originally introduced by Chen [6, 7] (also see [1]). We show here that high resolution and low errors can be achieved systematically, with a well-defined estimate of the total work required. As noted above, unlike packages such as RELION and EMAN2, we do not address the issues that arise when there is heterogeneity in the data sets due to the presence of multiple quasi-stable conformations of the particles being imaged. We also assume that the experimental images have known in-plane translations. Although fitting for such translations would be desired in a production code, we believe that this can be included in our algorithm with only a small constant factor increase in computation time (see the concluding section for further discussion). We do, however, include in our algorithm a known *contrast transfer function* (CTF) which models realistic aberrations of each experimental image.

In sections 2–5, we introduce the notation necessary to describe the algorithm in detail, as well as the various computational kernels that will be needed. The frequency marching procedure is described in section 6. Section 7 presents our numerical results, which use simulated data derived from known atomic positions for three relevant protein geometries of interest. The use of simulated data allows us to investigate the effect of signal-to-noise ratio (SNR) on reconstruction quality. We draw conclusions in section 8.

2. Mathematical preliminaries. We begin by establishing some notation. Throughout this paper, the unknown electron-scattering density will be denoted by $f(\mathbf{x})$, where $\mathbf{x} = (x, y, z)$ in Cartesian coordinates. We assume, without loss of generality, that the unit of length is chosen so that the particle (support of f) fits in the unit ball at the origin. The Fourier transform of $f(\mathbf{x})$ will be denoted by $F(\mathbf{k})$, where $\mathbf{k} = (k_1, k_2, k_3)$ in Cartesian coordinates and $\mathbf{k} = (k, \theta, \phi)$ in spherical coordinates, with $k_1 = k \sin \theta \cos \phi$, $k_2 = k \sin \theta \sin \phi$, and $k_3 = k \cos \theta$. Since we will use a spherical discretization of Fourier space, we write the standard Fourier relations in the following form:

$$(2.1) \quad F(k, \theta, \phi) = \int_{-1}^1 \int_{-1}^1 \int_{-1}^1 f(x, y, z) e^{ik(x \sin \theta \cos \phi + y \sin \theta \sin \phi + z \cos \theta)} dx dy dz$$

and

$$(2.2) \quad f(x, y, z) = \frac{1}{(2\pi)^3} \int_0^\infty \int_0^{2\pi} \int_0^\pi F(k, \theta, \phi) e^{-ik(x \sin \theta \cos \phi + y \sin \theta \sin \phi + z \cos \theta)} k^2 \sin \theta d\theta d\phi dk.$$

We will make use of Clenshaw–Curtis (Chebyshev) quadrature in $\cos \theta$ for the inner integral (see section 2.2), since this is spectrally accurate for smooth integrands and results in nodes that are equispaced in the parameter θ , which simplifies the task of interpolation.




Figure 1. Illustration of the projection-slice theorem: the 2D Fourier transform of the projection of a compactly supported function $f(\mathbf{x})$ equals a central slice of the Fourier transform $F(\mathbf{k})$ through the origin, with normal vector $(1, \alpha, \beta)$.

2.1. Rotation and projection operators. Let $(1, \alpha, \beta)$ denote the spherical coordinates of a point on the unit sphere, which we will also view as an *orientation vector*. With a slight abuse of notation, we will often identify (α, β) with the vector $(1, \alpha, \beta)$. Rather than assuming the electron beam orientation is along the z -axis and that the particle orientations are unknown, it is convenient to imagine that the particle of interest is fixed in the laboratory frame and that each projection obtained from electron microscopy corresponds to an electron beam in some direction (α, β) .

Definition 2.1. Let $\mathbf{d} = (1, \alpha, \beta)$ be an orientation vector, and let (r, ψ, s) denote a cylindrical coordinate system in \mathbb{R}^3 , where (r, ψ) are polar coordinates in the projection plane orthogonal to \mathbf{d} and s is the component along \mathbf{d} . The projection of the function $f(\mathbf{x})$ in the direction \mathbf{d} is denoted by $P_{\alpha, \beta}[f]$, with

$$P_{\alpha, \beta}[f](r, \psi) = \int_{-\infty}^{\infty} f(r, \psi, s) ds.$$

There is a simple connection between the projection of a function f , namely $P_{\alpha, \beta}[f]$, and its Fourier transform F , given by Theorem 2.3 and illustrated in Figure 1. For this, we will need to define a central slice of F .

Definition 2.2. Let $\mathbf{d} = (1, \alpha, \beta)$ be an orientation vector, and let $F(\mathbf{k})$ denote a function in the 3D Fourier transform domain. Then the central slice $S_{\alpha, \beta}[F]$ is defined to be the restriction of F to the plane through the origin with normal vector \mathbf{d} .

Theorem 2.3 (the projection-slice theorem). Let $f(\mathbf{x})$ denote a compactly supported function in \mathbb{R}^3 , let $F(\mathbf{k})$ denote its Fourier transform, and let (α, β) denote an orientation vector.

Let $P_{\alpha,\beta}[f]$ denote the corresponding projection of f , and let $\widehat{P_{\alpha,\beta}[f]}$ denote its 2D Fourier transform (with normalization analogous to (2.1)–(2.2)). Then $\widehat{P_{\alpha,\beta}[f]}$ corresponds to an equatorial slice through $F(\mathbf{k})$, with orientation vector (α, β) . That is,

$$\widehat{P_{\alpha,\beta}[f]} = S_{\alpha,\beta}[F].$$

The proof of this theorem is straightforward (see, for example, [29]). There is a third angular degree of freedom which must be taken into account in cryo-EM, easily understood by inspection of Figure 1. In particular, any rotation of the image plane about the orientation vector (α, β) could be captured during the experiment.

To take this rotation into account, we need to be more precise about the polar coordinate system (r, ψ) in the image plane. We will denote by ψ the angle subtended in the image plane relative to the projection of the negative z -axis in the fixed laboratory frame. (In the case of projections with directions passing through the poles $\alpha = 0$ and π , the above definition becomes ambiguous, and one may set ψ to be the standard polar angle in the xy plane.) Then γ specifies in-plane rotation of a projection. Thus, the full specification of an arbitrary projection takes the form

$$(2.3) \quad P_{\alpha,\beta,\gamma}[f](r, \psi) = P_{\alpha,\beta}[f](r, \psi - \gamma).$$

Remark 2. The angles (α, β, γ) correspond to a particular choice of Euler angles that define an arbitrary rotation of a rigid body in three dimensions. The notation introduced here is most convenient for the purposes of our method (corresponding to an extrinsic rotation of α about the y -axis, an extrinsic rotation of β about the z -axis, and an intrinsic rotation of γ about the new z -axis).

We will denote the set of experimental images obtained from electron microscopy (the input data) by \mathfrak{M} , with the total number of images given by $M = |\mathfrak{M}|$. Each image $\mathcal{M}^{(m)} \in \mathfrak{M}$ has support in the unit disc. We will often omit the superscript (j) when the context is clear. Representing the image in Cartesian coordinates, we let $\widehat{\mathcal{M}}^{(m)}$ denote its 2D Fourier transform in polar coordinates as follows:

$$(2.4) \quad \widehat{\mathcal{M}}^{(m)}(k, \psi) = \int_{-1}^1 \int_{-1}^1 \mathcal{M}^{(m)}(x, y) e^{ik(\cos \psi x + \sin \psi y)} dx dy.$$

2.2. Discretization. We discretize the full 3D Fourier transform domain (\mathbf{k} -space) in spherical coordinates as follows. We choose N_r equispaced quadrature nodes in the radial direction, between zero and a maximum frequency K . The latter sets the achievable resolution. On each sphere of radius k , we form a 2D product grid from N_ϕ equispaced points in the ϕ direction, and the following N_θ equispaced points in the θ direction:

$$\theta_j = \frac{(2j-1)\pi}{2N_\theta}, \quad 1 \leq j \leq N_\theta.$$

Writing the volume element $k^2 \sin \theta d\theta d\phi dk$ as $k^2 d\mu d\phi dk$, where $\mu = \cos \theta$ is the scaled z -coordinate, we note that the scaled z -coordinates of the nodes $\mu_j = \cos \theta_j$ are thus located at

the classical (first kind) Chebyshev nodes on $[-1, 1]$. The advantage of this particular choice for θ (or μ) nodes is that they will be convenient later for local interpolation in θ near the poles.

The complete spherical product grid of $N_r N_\theta N_\phi$ nodes is used for \mathbf{k} -space quadrature, with the weight associated with each node being $(2\pi/N_\phi)k^2 w_j$, where w_j is the weight corresponding to the node μ_j on $[-1, 1]$. (Here the μ quadrature is sometimes known as Fejér's first rule [56].) In particular, this quadrature will be used in computing the final density in physical space by means of the inverse Fourier transform (2.2).

It is important to note that because the support of f is in the unit ball, we have very precise bounds on the *smoothness* of $F(\mathbf{k})$. Namely, F , as a function of \mathbf{k} , is band-limited to unit "frequency" (note that here, and only here, we use "frequency" in the reverse sense to indicate the rate of oscillation of F with respect to the Fourier variable \mathbf{k}). Furthermore, for targets \mathbf{x} in the unit ball, the same is true of the exponential function in (2.2). Thus, the integrand is band-limited to a "frequency" of 2. This means that the above quadrature scheme is superalgebraically convergent with respect to N_ϕ and N_θ , due to well-known results on the periodic trapezoid and Chebyshev-type quadratures [56]. Furthermore, for oscillatory periodic band-limited functions, the periodic trapezoid rule reaches full accuracy once "one point per wavelength" is exceeded for the integrand [48] (note that this is half the usual Nyquist criterion). These results are expected to carry over to the spherical sections of 3D band-limited functions that we use, allowing for superalgebraic error terms. Since the most oscillation with respect to θ and ϕ occurs on the largest k sphere, the above sampling considerations imply $N_\phi \geq 2K$ and $N_\theta \geq K$. In practice, to ensure sufficient accuracy, we choose values slightly larger (i.e., up to a factor of 1.5) than these bounds.

Remark 3. In practice, the sphere can be sampled more uniformly and more efficiently by choosing N_θ to vary with k and by reducing the number of azimuthal points N_ϕ near the poles. For simplicity of presentation, we keep them fixed here (see section 8 for further discussion).

Quadrature in the k direction is only second order accurate, as we are relying on the trapezoidal rule with an equispaced grid. Spectral convergence could be achieved using Gauss–Legendre quadrature, but we find that accuracy with a regular grid is sufficient using a node spacing $\delta k := K/N_r \approx 2$, so that $N_r = O(K)$.

In our model, rather than using the spherical grid points themselves to sample $F(\mathbf{k})$, we will, for the most part, make use of a spherical harmonic representation. That is, for each fixed radial value k , we will represent $F(k, \theta, \phi)$ on the corresponding spherical shell in the form

$$(2.5) \quad F(k, \theta, \phi) = \sum_{n=0}^{p(k)} \sum_{m=-n}^n f_{nm}(k) Y_n^m(\theta, \phi).$$

Here,

$$(2.6) \quad Y_n^m(\theta, \phi) = S_n^m P_n^{|m|}(\cos \theta) e^{im\phi},$$

where $P_n(x)$ denotes the standard Legendre polynomial of degree n , the associated Legendre

functions P_n^m are defined by the Rodrigues formula

$$P_n^m(x) = (-1)^m (1-x^2)^{m/2} \frac{d^m}{dx^m} P_n(x),$$

and

$$(2.7) \quad S_n^m = \sqrt{\frac{2n+1}{4\pi}} \sqrt{\frac{(n-|m|)!}{(n+|m|)!}}.$$

The functions Y_n^m are orthonormal with respect to the L_2 inner product on the unit sphere

$$(2.8) \quad \langle f, g \rangle = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) g(\theta, \phi) \sin \theta \, d\theta \, d\phi.$$

Remark 4. Note that the degree of the expansion in (2.5) is a function of k , denoted by $p(k)$. It is straightforward to show that spherical harmonic modes of degree greater than k are exponentially decaying on a sphere of radius k given that the original function $f(x, y, z)$ is supported in the unit ball in physical space. This is the same argument used in discussing grid resolution above for N_r , N_ϕ , and N_θ . In short, an expansion of degree $O(k)$ is sufficient, and, throughout this paper we simply fix the degree $p(k) = k + 2$.

Remark 5. The orthogonality of the Y_n^m allows us to obtain the coefficients of the spherical harmonic expansion f_{nm} above via projection. Moreover, separation of variables permits all the necessary integrals to be computed in $O(p(k)^3)$ operations.

3. Template generation. In any projection matching procedure, a recurring task is that of *template generation*. That is, given a current model defined by $f(\mathbf{x})$ or $F(\mathbf{k})$, we must generate a collection of projection images of the model for a variety of orientation vectors. We will denote those orientations by (α_i, β_j) , leaving (θ, ϕ) to refer to the angular coordinates in Fourier space for a fixed frame of reference. From the projection-slice theorem (Theorem 2.3), we have

$$\widehat{P_{\alpha_i, \beta_j}[f]} = S_{\alpha_i, \beta_j}[F].$$

In fact, our projection matching procedure will work entirely in 2D Fourier space, i.e., using $\widehat{P_{\alpha_i, \beta_j}[f]}(k, \psi)$, where (k, ψ) are polar coordinates in the projection plane, so our task is simply to compute the central slice $S_{\alpha_i, \beta_j}[F]$ on a polar grid for each i and j (see Figure 2). The real-space projections $P_{\alpha_i, \beta_j}[f]$ will not be needed.

In order to generate all templates efficiently, let a discrete set of desired orientation vectors be denoted by $\{(\alpha_i, \beta_j) \mid i = 1, \dots, N_\theta, j = 1, \dots, N_\phi\}$. We have found that, to achieve acceptable errors at realistic noise levels, these grid sizes N_θ and N_ϕ can be chosen to be the same as for the quadrature points via the Nyquist criterion in the previous section. Each α corresponds to an initial rotation about the y -axis in the laboratory frame, while each β corresponds to a subsequent rotation about the z -axis in the laboratory frame. (The fact that projection matching can be simplified by creating a sufficiently fine grid on the sphere was already discussed and used in [31, 34].)

To generate the points at a given radius k within each slice, we must parametrize points in a polar coordinate system on the equatorial plane in the laboratory frame that have been

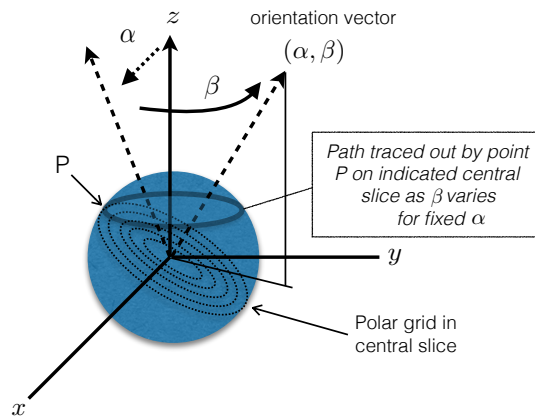


Figure 2. To generate templates for projection matching, we need to sample the central slice $S_{\alpha,\beta}[F]$ on a polar grid for every orientation vector (α, β) . The process is accelerated by observing that the z -coordinate for any such target point depends only on α and not on β . This permits the effective use of separation of variables (equation (3.1)).

rotated through the preceding actions. Using Cartesian coordinates for the moment, it is easy to see that

$$\begin{pmatrix} k \cos \psi \\ k \sin \psi \\ 0 \end{pmatrix} \xrightarrow{\alpha \text{ rotation}} \begin{pmatrix} k \cos \alpha \cos \psi \\ k \sin \psi \\ -k \sin \alpha \cos \psi \end{pmatrix} \xrightarrow{\beta \text{ rotation}} \begin{pmatrix} k[\cos \beta \cos \alpha \cos \psi - \sin \beta \sin \psi] \\ k[\sin \beta \cos \alpha \cos \psi + \cos \beta \sin \psi] \\ k[-\sin \alpha \cos \psi] \end{pmatrix}.$$

Note that the z -coordinate (and therefore the spherical coordinate θ) in the laboratory frame is independent of β , while the x -, y -coordinates in the laboratory frame depend on α , β , and ψ .

For each sampled normal orientation vector (α_i, β_j) for the slice, let $\{\psi_l | l = 1, \dots, N_\phi\}$ denote equispaced points on $[0, 2\pi]$. From above, the sample points with radius k , which are located at $(k \cos \psi_l, k \sin \psi_l, 0)$ before rotation, move to

$$\begin{pmatrix} x_{ijl} \\ y_{ijl} \\ z_{il} \end{pmatrix} = \begin{pmatrix} k[\cos \beta_j \cos \alpha_i \cos \psi_l - \sin \beta_j \sin \psi_l] \\ k[\sin \beta_j \cos \alpha_i \cos \psi_l + \cos \beta_j \sin \psi_l] \\ k[-\sin \alpha_i \cos \psi_l] \end{pmatrix}.$$

Fixing the radius k for now, we denote the spherical coordinates of these points $(x_{ijl}, y_{ijl}, z_{il})$ by $(k, \theta_{il}, \phi_{ijl})$. To generate the template data, we now need to evaluate the spherical harmonic expansion at these points, namely,

$$\begin{aligned} F(k, \theta_{il}, \phi_{ijl}) &= \sum_{n=0}^{p(k)} \sum_{m=-n}^n f_{nm}(k) Y_n^m(\theta_{il}, \phi_{ijl}) \\ (3.1) \qquad \qquad \qquad &= \sum_{m=-n}^n A_{iml} e^{im\phi_{ijl}} \end{aligned}$$

for $i = 1, \dots, N_\theta$ and $j = 1, \dots, N_\phi$, where

$$A_{iml} = \sum_{n=|m|}^{p(k)} f_{nm}(k) S_n^m P_n^m(\cos \theta_{il}).$$

Note that only the degrees at least as large as the magnitude of the order $|m|$ contribute. For each radius k , the cost of computing the set $\{A_{iml}\}$ is $O(N_\theta N_\phi p(k)^2)$, and the cost of the subsequent evaluation of all values $F(k, \theta_{il}, \phi_{ijl})$ is $O(N_\theta N_\phi^2 p(k))$. Since $p(k) = O(K)$, both terms have the asymptotic complexity $O(K^4)$. This generates $O(K^3)$ template data points on a single spherical shell. Then, summing over all k shells, the cost to generate the full spherical grid templates up to a resolution of K requires $O(K^5)$ work and generates $O(K^4)$ data.

Letting $k_q \in \{k_1, \dots, k_{N_r}\}$ denote the discretization in k , the above procedure evaluates the $O(K^2)$ samples comprising the polar grid (k_q, ψ_l) for each of the $O(K^2)$ central slices. However, it is more convenient to store these templates in terms of their angular Fourier series for each slice, i.e.,

$$(3.2) \quad S_{\alpha_i, \beta_j}[F](k_q, \psi) = \sum_{n=-N_q}^{N_q} S_n^{ij}(k_q) e^{in\psi},$$

with $2N_q + 1$ Fourier modes used on the q th ring. By band-limit considerations of the spherical harmonics, one need only choose $N_q = p(k_q)$, the maximum degree for each k shell.

The coefficients $S_n^{ij}(k_q)$ are evaluated by applying the fast Fourier transform (FFT) to the template data along the ψ grid direction, requiring $O(K^4 \log K)$ work.

4. Projection matching. Given an experimental image $\mathcal{M} \in \mathfrak{M}$ (or, more precisely, its 2D Fourier transform $\widehat{\mathcal{M}}(k, \psi)$ discretized on a polar grid), we seek to rank the templates defined by (α_i, β_j) and the rotational degree of freedom γ in terms of how well they match the image. For this we need a generative model for images. We will then present an algorithm for matching γ .

4.1. Image model with contrast transfer function correction. For reasons having to do with the physics of data acquisition, each particle image $\mathcal{M}^{(m)}$ is not simply a projection of the electron-scattering density, but may be modeled as the projection of the electron-scattering density convolved with a CTF [28], plus noise. The CTF includes diffraction effects due to the particle's depth in the ice sheet, linear elastic and inelastic scattering, and detection effects. In the simplest case, the CTF is radially symmetric and its Fourier transform is real-valued and of the form $C^{(m)}(k)$, depending only on k . For the purposes of the present paper, we will assume that $C^{(m)}(k)$ is known for each experimental image. By the convolution theorem, the CTF acts multiplicatively on the Fourier transform image. Thus, the expectation for the m th Fourier image given by beam direction (α, β) and in-plane rotation γ , deriving from a \mathbf{k} -space scattering density F , is

$$\widehat{\mathcal{M}}^{(m)}(k, \psi - \gamma) \approx C^{(m)}(k) S_{\alpha, \beta}[F](k, \psi).$$

To this is added measurement noise, which is assumed to be Gaussian and i.i.d. (independent and identically distributed) on each pixel in the image.

4.2. Fitting the best orientation for each image. It follows from the above noise model that the correct quantity to minimize when searching for the best match would be the L_2 norm of the difference between the template and the image, which is equivalent to the L_2 norm of their difference in the Fourier image plane. This may be interpreted as minimizing the negative log likelihood, i.e., finding the maximum likelihood. Let $\widehat{\mathcal{M}} = \widehat{\mathcal{M}}^{(m)}$ be the current experimental image in question, and let $C = C^{(m)}$ be its corresponding known CTF. We will denote by $\widehat{\mathcal{M}}_\gamma$ the rotated Fourier image

$$\widehat{\mathcal{M}}_\gamma(k, \psi) := \widehat{\mathcal{M}}(k, \psi - \gamma) .$$

For the template index pair (i, j) , a quantity to be minimized over the rotation γ would thus be

$$(4.1) \quad \|\widehat{\mathcal{M}}_\gamma - CS_{\alpha_i, \beta_j}[F]\|_2,$$

where, for a function $g(k, \gamma)$ in polar coordinates, the L_2 norm (up to a band-limit of K) is defined by

$$(4.2) \quad \|g\|^2 := \int_0^K \int_0^{2\pi} |g(k, \psi)|^2 d\psi k dk .$$

However, typically there is uncertainty about the overall normalization (a multiplicative prefactor) associated with experimental images (as discussed, for instance, by Scheres et al. [41]). It is straightforward to check that minimizing (4.1) over i, j, γ including an unknown normalization factor for the image is equivalent to maximizing over i, j, γ the normalized inner product

$$(4.3) \quad \frac{\langle CS_{\alpha_i, \beta_j}[F], \widehat{\mathcal{M}}_\gamma \rangle}{\|CS_{\alpha_i, \beta_j}[F]\| \|\widehat{\mathcal{M}}_\gamma\|} .$$

Here the inner product corresponds to the norm (4.2), and (4.3) may be interpreted as the cosine of the ‘‘angle’’ between the image and the template in the abstract vector space with norm (4.2). We refer to maximizing (4.3) over i, j , and γ as *projection matching*.

For each $k_q \in \{k_1, \dots, k_{N_r}\}$, we precompute a Fourier series representation of the image, with $2N_q + 1$ Fourier modes on the q th ring:

$$\widehat{\mathcal{M}}(k_q, \psi) = \sum_{n=-N_q}^{N_q} \widehat{\mathcal{M}}_n(k_q) e^{in\psi} .$$

This requires $O(K^2 \log K)$ work, done once for each experimental image.

To rank the template matches, we loop over all projection directions indexed by i, j . For each of these directions (α_i, β_j) , since the denominator of (4.3) is fixed, we need to maximize the inner product as a function of γ ,

$$\langle CS_{\alpha_i, \beta_j}[F], \widehat{\mathcal{M}}_\gamma \rangle := \int_0^K \int_0^{2\pi} \widehat{\mathcal{M}}(k, \psi - \gamma) \overline{C(k) S_{\alpha_i, \beta_j}[F](k, \psi)} d\psi k dk .$$

Elementary Fourier analysis shows that

$$\langle CS_{\alpha_i, \beta_j}[F], \widehat{\mathcal{M}}_\gamma \rangle = \sum_n c_n(K) e^{-in\gamma},$$

where

$$(4.4) \quad c_n(\kappa) := \int_0^\kappa \widehat{\mathcal{M}}_n(k) \overline{C(k) S_n^{ij}(k)} k dk.$$

This integral is approximated using the existing k grid up to the maximum frequency $\kappa = O(K)$, requiring $O(K^2)$ effort to compute the $2N_q + 1$ coefficients. Since there are $O(K^2)$ pairs i, j , this totals $O(K^4)$ work per experimental image. Finally, for each pair i, j , we will compute the best value for γ by tabulating $\sum_n c_n(K) e^{-in\gamma}$ on a uniform grid of N_ϕ values, using the FFT, and taking the index with the maximum modulus. This requires $O(K^3 \log K)$ work per image. Thus, the complexity for angle fitting all images is $O(MK^4)$.

Definition 4.1. *The best values of (α, β, γ) for image $\mathcal{M}^{(m)}$ will be denoted by $\mathcal{A}_m = (\alpha_m, \beta_m, \gamma_m)$.*

In practice we implement two useful adjustments to the above procedure as follows:

(1) We search first over a coarse grid of template directions (with angle resolution five times coarser than the grid defined by N_θ and N_ϕ), and then only search over the set of grid points that are within one coarse grid point of the global maxima found on the coarse grid. This improves efficiency by a constant factor, and in our tests does not degrade accuracy noticeably.

(2) We define a “randomization parameter” f_{rand} . If $f_{rand} = 0$, we return the single best-fitting orientation: $(\alpha_m, \beta_m, \gamma_m)$ defined above. However, if $f_{rand} > 0$, we instead choose \mathcal{A}_m randomly from the set of all discrete orientations (α, β, γ) that produce a normalized inner product (4.3) greater than $1 - f_{rand}$. Typically we choose f_{rand} small, e.g., 0.02. This uniformizes the distribution of orientations over the set which fit the image almost equally well. However, it is sometimes beneficial to increase f_{rand} for improved convergence rate in the least-squares procedure of the following section; see also section 7.3.

It should be noted that the idea of using Fourier methods to find the optimal third Euler angle γ can be found in [10], and the complexity of various alignment schemes using polar grids in physical space is discussed in [24].

Remark 6. In the simplest version of the frequency marching scheme that we present, the c_n coefficients in (4.4) are evaluated using all frequencies from zero to κ , since F , and hence the templates, get updated over this frequency range at each iteration. We believe that updating the model F only in the current k shell could result in a faster algorithm without sacrificing accuracy. In that case, the c_n values are already known for a smaller κ , and they can be incremented according to the formula

$$c_n(\kappa_2) = c_n(\kappa_1) + \int_{\kappa_1}^{\kappa_2} \widehat{\mathcal{M}}_n(k) \overline{C(k) S_n^{ij}(k)} k dk.$$

5. Reconstruction from particle images with known angular assignments. Suppose now that we have the collection of Fourier transforms of all M particle images, which we denote by

$$\widehat{\mathfrak{M}} := \{\widehat{\mathcal{M}}^{(1)}(k, \psi), \widehat{\mathcal{M}}^{(2)}(k, \psi), \dots, \widehat{\mathcal{M}}^{(M)}(k, \psi)\},$$

together with the corresponding angular assignments $\mathcal{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_M\}$, where $\mathcal{A}_m = (\alpha_m, \beta_m, \gamma_m)$.

Remark 7. In our present implementation, the values of $(\alpha_m, \beta_m, \gamma_m)$ are selected from the set of computed template angles, as discussed in the previous section. A more sophisticated projection matching procedure could generate “off-grid” values for α, β , and γ . The reconstruction procedure below works in either case.

We seek to reconstruct a (new) model $F(k, \theta, \phi)$ that is consistent with all of the data in a least-squares sense. For each fixed spherical shell of radius k , we use the representation (2.5), and data $\widehat{\mathcal{M}}_{\gamma_m}^{(m)}(k, \psi)$ which, if the angle assignments are correct, is expected to approximate the CTF-corrected central slice $C^{(m)}(k)S_{\alpha_m, \beta_m}[F](k, \psi)$ for the new model. The matching between the new model and the image data on the k shell is done at a set of N_ϕ discrete angles ϕ_l on each great circle. The Cartesian coordinates of those points are

$$(5.1) \quad \begin{pmatrix} k[\cos \beta_m \cos \alpha_m \cos \psi_l - \sin \beta_m \sin \psi_l] \\ k[\sin \beta_m \cos \alpha_m \cos \psi_l + \cos \beta_m \sin \psi_l] \\ k[-\sin \alpha_m \cos \psi_l] \end{pmatrix},$$

from which their spherical coordinates (θ_{ml}, ϕ_{ml}) are easily computed. Thus, for the Fourier sphere of radius k , each image contributes N_ϕ points, for a total of $N_{tot} = MN_\phi$ points on this sphere.

At this stage, we simply collect all image data with given k into a right-hand side vector \mathbf{d} , defined by its N_{tot} entries

$$d_{(m-1)N_\phi+l} = \widehat{\mathcal{M}}_{\gamma_m}^{(m)}(k, \psi_l) \quad \text{for } m = 1, \dots, M, \quad l = 1, \dots, N_\phi.$$

From now we shall use a single index $i = 1, \dots, N_{tot}$ to reference these entries, and use (θ_i, ϕ_i) for the spherical coordinates for the corresponding point on the great circle. (In our present implementation we ignore the use of image normalization factors that could be extracted from the orientation fitting algorithm presented in section 4.2. This is justified since our synthetic noisy images will be generated without varying normalization factors.)

Let \mathbf{f} denote the vector of “unrolled” spherical harmonic coefficients in the representation for F on the current k sphere,

$$\mathbf{f} = \{f_{0,0}(k), f_{1,-1}(k), f_{1,0}(k), f_{1,1}(k), f_{2,-2}(k), \dots, f_{p(k),p(k)}(k)\},$$

so that \mathbf{f} is a complex vector of length $(p(k) + 1)^2$. Then we define the complex-valued matrix S , of dimension $N_{tot} \times (p(k) + 1)^2$, by

$$(5.2) \quad (S\mathbf{f})_i = \sum_{n=0}^{p(k)} \sum_{m=-n}^n f_{nm}(k) Y_n^m(\theta_i, \phi_i).$$

Thus, S evaluates the spherical harmonic expansion with coefficients \mathbf{f} at all sphere points. We also let $C_i(k)$ be the corresponding CTF for each image at frequency k (note that $C_i(k)$ is the same for all i belonging to the same image) and let \mathcal{C} be the diagonal matrix with diagonal entries $C_i(k)$. We find the desired solution by solving the problem

$$\mathcal{C}S\mathbf{f} = \mathbf{d}$$

in a least-squares sense. Note that the use of the l^2 norm corresponds to the same assumption of i.i.d. Gaussian noise on the images as in the previous section.

For this, we use conjugate gradient (CG) iteration on the normal equations

$$S^H\mathcal{C}^H\mathcal{C}S\mathbf{f} = S^H\mathcal{C}^H\mathbf{d},$$

where S^H is the Hermitian transpose of S . Applying the diagonal matrices \mathcal{C} and \mathcal{C}^H is trivial. Direct computation of the matrix-vector product with S or S^H would require $O(MN_\phi(p(k)+1)^2) = O(MK^3)$ work. The application of S is easily accelerated, however, by first evaluating the spherical harmonic expansion induced by \mathbf{f} on a *regular* grid with N_ϕ equispaced points in the ϕ direction, and N_θ points in the θ direction whose cosines are classical Chebyshev nodes (see section 2.2). Using separation of variables, this requires $O(K^3)$ work, since $N_\phi = N_\theta = O(K)$. The value of $S\mathbf{f}$ at an arbitrary point can then be computed by local q th order Lagrange interpolation in ϕ and θ from this regular grid, at a cost of q^2 operations per target. It is sufficient to use q in the range 5–10. We denote by S_{reg} the mapping from spherical harmonic coefficients to values on the $N_\phi \times N_\theta$ grid, and by T the (sparse) interpolation matrix from the $N_\phi \times N_\theta$ grid points to the N_{tot} arbitrary locations. In other words, up to interpolation error,

$$S\mathbf{f} \approx TS_{reg}\mathbf{f},$$

and applying S in this manner requires only $O(Mq^2K + K^3)$ work. Clearly,

$$S^H\mathcal{C}^H\mathcal{C}S \approx S_{reg}^HT^H\mathcal{C}^H\mathcal{C}TS_{reg},$$

where S_{reg}^H can be applied by projection from a regular grid to a spherical harmonic expansion, which is also $O(K^3)$. Thus each CG iteration requires only $O(Mq^2K + K^3)$ work. As long as the N_{tot} points have reasonable coverage of the sphere (specifically, there are no large ‘‘pockets’’ on the sphere which are empty of points), the system is well conditioned and requires only a modest number of iterations, around 20–50, to achieve several digits of accuracy. It is worth noting (as in other Fourier-based schemes, such as FREALIGN [17, 18, 26]) that this least-squares procedure is the step in the reconstruction that is responsible for denoising. The more experimental images that are available (with poor SNRs but correctly assigned angles), the more accurately we are able to estimate \mathbf{f} and hence $F(\mathbf{k})$ and the electron-scattering density $f(\mathbf{x})$.

The above describes the computation of F on a single k shell. One advantage of our representation is that to build the complete new model F , the least-squares solve for each k shell in the grid $\{k_1, \dots, k_{N_r}\}$ may be performed independently. This gives an overall complexity for the reconstruction of $O(Mq^2K^2 + K^4)$.

Definition 5.1. Let \mathbf{F} indicate the entire set of model coefficients $\{\mathbf{f}(k_1), \dots, \mathbf{f}(k_{N_r})\}$. We will refer to its least-squares approximation (on all k shells) by \mathbf{F}^* . We summarize this by

$$(5.3) \quad \mathbf{F}^* = L(\widehat{\mathfrak{M}}, \mathcal{A}) := \arg \min_{\mathbf{F}} \sum_{m=1}^M \|\mathcal{C}^{(m)} S_{\mathcal{A}_m}(\mathbf{F}) - \widehat{\mathcal{M}}^{(m)}\|_2^2,$$

where $\mathcal{C}^{(m)}$ is the CTF for the m th image and, by analogy with (2.3), $S_{\mathcal{A}_m}(\mathbf{F})$ denotes the slice

$$S_{\mathcal{A}_m}[F] = P_{\alpha_m, \beta_m, \gamma_m}[f].$$

Here, F is determined from the coefficients \mathbf{F} via the spherical harmonic representation (2.5).

6. The full inverse problem. As discussed in the previous section, if the angular assignments $\mathcal{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_M\}$ of the experimental images were given, \mathbf{F} would be recovered by solving the least-squares problem (5.3). Since the angles \mathcal{A} are unknown, however, it is standard to write the cryo-EM reconstruction problem in terms of the nonlinear and nonconvex optimization task in the joint set of unknowns:

$$(6.1) \quad \{\mathbf{F}^*, \mathcal{A}^*\} = \arg \min_{\{\mathbf{F}, \mathcal{A}\}} \sum_{m=1}^M \|\mathcal{C}^{(m)} S_{\mathcal{A}_m}(\mathbf{F}) - \widehat{\mathcal{M}}^{(m)}\|_2^2.$$

As noted above, for simplicity of presentation of the scheme, we omit image normalization prefactors discussed in section 4.2.

One well-known approach to solving (6.1) is to start with an initial guess $\mathbf{F}^{(0)}$ for the representation \mathbf{F} , and then to iterate as follows:

Classical iterative refinement.

Set $i = 1$. While (convergence criterion has not been met),

1. Compute $\mathcal{A}^{(i)}$ from $\mathbf{F}^{(i-1)}$ by projection matching (section 4).
2. Set $\mathbf{F}^{(i)} = L(\widehat{\mathfrak{M}}, \mathcal{A}^{(i)})$ by least-squares solution (section 5).
3. $i \leftarrow i + 1$.

Compute F from the final \mathbf{F} via (2.5); then take the inverse Fourier transform (2.2) to recover f .

This iteration can be viewed as coordinate descent, alternating between fitting the best angle assignment and fitting the best model density. It will converge, but to a local minimum—not necessarily the correct solution [8]. Various attempts to overcome this convergence failure have been proposed, including annealing strategies and stochastic hill climbing (see, for example, [12, 8]), but robustness has remained an issue.

6.1. Frequency marching (recursive linearization). As in the iteration above, we alternate between projection matching to obtain estimates for the angles \mathcal{A} and solving least-squares problems to determine the best set of density coefficients \mathbf{F} . However, by continuously increasing the resolution—measured in terms of the maximal spatial frequency used in the representation for \mathbf{F} —we bypass the difficulties associated with multiple minima in existing attempts at iterative refinement. In the language of optimization, this can be viewed

as a *homotopy method* using the maximum spatial frequency (resolution) as the homotopy parameter. The basic intuition underlying our scheme is motivated by the success of recursive linearization in inverse acoustic scattering [1, 4, 6, 7].

More precisely, let $\widehat{\mathfrak{M}}([0, k])$ denote the set of Fourier transforms of all experimental images restricted to the disk of radius k , and let $\mathbf{F}([0, k])$ denote the density coefficients only up to frequency k , i.e., using the shells for which $k_q \leq k$. The full objective function minimization restricted to maximum frequency k is

$$(6.2) \quad \{\mathbf{F}^*([0, k]), \mathcal{A}^*\} = \arg \min_{\{\mathbf{F}([0, k]), \mathcal{A}\}} \sum_{m=1}^M \|\mathcal{C}^{(m)} S_{\mathcal{A}_m}(\mathbf{F}([0, k])) - \widehat{\mathcal{M}}^{(m)}([0, k])\|_2^2,$$

which is still a nonlinear and nonconvex optimization problem. However, if $\mathbf{F}([0, k])$ is known and we only seek to find $\mathbf{F}([0, k + \delta k])$ for sufficiently small δk , then the solution can be reached by a suitable linearization. Moreover, at low frequency, say for $k \leq k_1 = 2$, the landscape is extremely smooth so that the global minimum is easily located. Thus, we propose simply assigning random angles $\mathcal{A}_m = (\alpha_m, \beta_m, \gamma_m)$ to the images at $k_{min} = k_1 = 2$ and iterating as follows:

Solution by frequency marching.

On input, we define a sequence of frequency steps from k_1 to $k_{N_r} = K$ with a step of $\delta k = k_{i+1} - k_i$.

- Set $\mathcal{A}^{(0)}$ to uniform random values over the allowed ranges.
- Set $\mathbf{F}([0, k_1]) = L(\widehat{\mathfrak{M}}([0, k_1]), \mathcal{A}^{(0)})$ by least squares (section 5).

For $i = 1, 2, \dots, N_r - 1$,

1. Compute $\mathcal{A}^{(i)}$ from $\mathbf{F}([0, k_i])$ by projection matching (section 4).
2. Set $\mathbf{F}([0, k_{i+1}]) = L(\widehat{\mathfrak{M}}([0, k_{i+1}]), \mathcal{A}^{(i)})$ by least squares (section 5).

Compute F from $\mathbf{F}([0, K])$ via (2.5); then take the inverse Fourier transform (2.2) to recover f .

Note that the procedures of sections 4 and 5 restrict naturally to any frequency range $[0, k]$. The key feature of refinement by recursive marching is that it is a deterministic procedure involving only linear solves and angular assignments of images. The overall complexity, combining the estimates from sections 4 and 5, and summing over k , is $O(MK^5 + K^6 + q^2MK^3)$, where the first two terms come from angle matching and the last from the least-squares solution. Since in current cryo-EM applications, $K \sim 10^2$, while $M \sim 10^5$ to 10^6 , and $q^2 < K$, the dominant cost is $O(MK^5)$, assuming we use the global angle matching procedure described above. However, we find that in our examples, due to the number of CG iterations, the time for least-squares fitting is actually quite similar to that for angle matching, i.e., they are close to balanced.

While we are not able to provide a proof in the general case, we believe that under fairly broad conditions this iteration will converge with high probability for sufficiently small δk (the radial frequency grid spacing). Informally speaking, we believe that a solution near the global minimum is *often* reached at low k in the first few iterations and that a path to the global minimum at K is then reached as k increases continuously.

In the experiments below we show that, even with very noisy data, a harsh random start at $k_{min} = 2$ is sufficient. We will return to this question in the concluding section.

In practice, we have implemented a further acceleration to the above scheme: if the mean absolute change in angles between $\mathcal{A}^{(i+1)}$ and $\mathcal{A}^{(i)}$ is less than 10^{-3} , the next increment of the index i is set to five rather than its usual value of one. This greatly reduces the number of steps in the marching procedure once the majority of angles has locked in with sufficient accuracy.

7. Numerical experiments with synthetic data. We implemented the above frequency marching algorithm and performed experiments using simulated images to reconstruct electron-scattering densities for three proteins.

7.1. Choice of ground-truth densities and error metric. The three proteins we used for numerical experiments were (see Figure 3) spinach rubisco enzyme (molecular weight 541kDa, 133 Å longest dimension), lipoxygenase-1 (weight 95kDa, 107 Å longest dimension), and scorpion protein neurotoxin (weight 7.2kDa, 33 Å longest dimension). Atomic locations in Å are taken from the protein data bank (PDB) entries [3], with codes 1RCX, 1YGE, 1AHO, respectively, but were shifted to have their centroid at the origin. Only in the last of the three are hydrogen atom locations included in the PDB. Thus, the number of atoms used were 37456, 6666, and 925, respectively. Note that the last of these, the neurotoxin, is much smaller than the ≈ 100 kDa lower weight limit that can be imaged through cryo-EM with current detector resolution and motion-correction technology; our point in including it is to show that, given appropriate experimental image resolution, reconstruction of small molecules does not pose an algorithmic problem.

For each protein, a ground-truth electron-scattering density $f(\mathbf{x})$ was produced using the following simple model. First, a nondimensionalization length D was chosen giving the number of physical Å per unit numerical length, such that the support of f in such numerical units lies within the unit ball. For rubisco, $D = 70$ Å; for lipoxygenase-1, $D = 60$ Å; and for the neurotoxin, $D = 25$ Å. Thus, all physical distances in what follows are divided by D in the numerical implementation. We summed 3D spherically symmetric Gaussian functions, one at each atomic location, giving each Gaussian a standard deviation $\sqrt{(0.5r)^2 + b^2}$, where r is the atomic radius for the type of atom (which varies from 0.42 Å to 0.88 Å). The factor 0.5 results in around 74% of the mass of the unblurred Gaussian falling within radius r , and b is a convolution (blurring) radius used to make f smooth enough to be accurately reconstructed using the image pixel resolution. The radii b used were 2.5Å for rubisco, 2.0Å for lipoxygenase-1, and 1.0Å for scorpion toxin, chosen to be around twice the simulated pixel spacing (see the next section). Thus, when we assess reconstruction errors, we are doing so against an appropriately smoothed ground-truth f . For simplicity, each such Gaussian was given a unit peak amplitude.

Our error metric for a reconstructed density \tilde{f} is the relative L_2 -norm,

$$(7.1) \quad \epsilon := \frac{\|\tilde{f} - f\|_{L_2([-1,1]^3)}}{\|f\|_{L_2([-1,1]^3)}},$$

which is estimated using quadrature on a sufficiently fine uniform 3D Cartesian grid (we used 100 points in each dimension).




Figure 3. Top row: the three protein structures (derived from X-ray crystallography). Bottom row: the corresponding atomic densities $f(\mathbf{x})$ used as a ground truth for the numerical experiments, shown via an isosurface at approximately half the maximum density. Note that the visual length scale for the three molecules is different.

Remark 8. Since \tilde{f} generally acquires an arbitrary rotation (α, β, γ) relative to the ground-truth f , we must first rotate it to best fit the original f . This is done by applying the procedure of section 4 to a small number (typically 10) of random projections. \tilde{f} is then rotated using the 3D nonuniform FFT [16] to evaluate its Fourier transform \tilde{F} at a rotated set of spherical discretization points as in section 2.2. A second nonuniform FFT is then applied to transform back to real space. Finally, the error (7.1) is evaluated.

Some researchers use a normalized cross-correlation metric to report errors (e.g., [23]). We note that, when errors are small, this metric is close to $1 - \epsilon^2/2$, with ϵ given by (7.1).

7.2. Generation of synthetic experimental images. For each of the three proteins, M (of order 50,000) synthetic experimental images were produced as follows. For each image m we

first defined a realistic radially symmetric CTF function $C^{(m)}(k)$ using the standard formulae [28, 54]

$$(7.2) \quad C(k) = B(\theta)[w_1 \sin \chi(\theta) - w_2 \cos \chi(\theta)], \quad \text{where } \chi(\theta) := \frac{1}{2}kz\theta^2 + kC_s\theta^4/8, \quad B(\theta) := e^{-\theta^2/\theta_0^2}.$$

Here χ is called the phase function; $\theta_0 = 0.002$ sets the microscope acceptance angle; $w_2 = 0.07$ controls the relative inelastic scattering, with $w_1^2 + w_2^2 = 1$; and the spherical aberration is $C_s = 2 \times 10^7$ Å. The defocus parameter z was different for each image, chosen uniformly at random in the interval $[1, 4] \times 10^4$ Å. Finally, in the above, angles θ are related to wavenumbers k in the numerical experiments via

$$(7.3) \quad \theta = \frac{\lambda k}{2\pi D},$$

where $\lambda = 0.025$ Å is the free-space electron wavelength (a typical value corresponding to a 200 keV microscope), and for convenience the distance scaling D has been included. In our setting, θ is of order 10^{-5} times the numerical wavenumber k .

The images were sampled on a uniform 2D grid at a standard resolution of 100×100 pixels covering the numerical box $[-1, 1]^2$. The pixel spacing thus corresponded to $D/50$, or between 0.5 Å for the neurotoxin (this is smaller than what is currently achievable experimentally) and 1.4 Å for rubisco. The noise-free signal images were produced by inverse Fourier transforming (via the 2D nonuniform FFT) slices taken at random orientations through the ground-truth Fourier density $F(\mathbf{k})$, after multiplication by the radial CTF $C^{(m)}(k)$ particular to each image. Finally, i.i.d. Gaussian noise was added to each pixel, with variance chosen to achieve a desired SNR. The standard definition of SNR in imaging is the ratio of the squared L_2 -norm of the signal to that of the noise (for this we used the domain $[-1, 1]^2$ since the molecule images occupy a large fraction of this area). We generated images at SNR values of ∞ (no noise), 0.5, 0.1, and 0.05. A typical value in applications is 0.1.

7.3. Results. All numerical experiments were run on desktop workstations with 14 cores, Intel Xeon 2.6GHz CPU, and 128 GB RAM. Our implementation is in Fortran, with OpenMP parallelization. We used frequency marching with step size $\delta k = 2$, starting with random angular assignment at $k_1 = 2$, and marching to full resolution at $K = 70$. This K was sufficient to resolve the decaying tails of F without significant truncation error, given our choice of blurring radius b at approximately two pixels. We used interpolation order $q = 7$, and we started with randomization factor $f_{rand} = 0.02$. The latter was changed adaptively during marching to balance solution time and accuracy as follows: if at some point the least-squares CG did not converge after 100 iterations, f_{rand} was doubled and the least-squares solve repeated. If instead CG required fewer than 50 iterations, f_{rand} was halved. The speed of our algorithm allowed us to carry out multiple runs for each protein, using either the same or a fresh set of synthetic images. In the present paper, we carried out five runs for each protein.

To assess how close our relative L_2 error ϵ was to the best achievable, given the image sampling, Fourier representation, and SNR, we computed for comparison the *best possible*

Table 1

Relative L_2 errors ϵ (see (7.1)) of the reconstructions using frequency marching, compared to the best possible reconstruction using known angles for the experimental images. The frequency marching results are given as averages over five runs, with estimated standard deviation.

		No noise	SNR 0.5	SNR 0.1	SNR 0.05
Rubisco	Frequency marching	0.06 ± 0.001	0.06 ± 0.001	0.08 ± 0.001	0.09 ± 0.004
	Known angles	0.05	0.06	0.07	0.08
Lipoxygenase-1	Frequency marching	0.05 ± 0.008	0.06 ± 0.002	0.09 ± 0.005	0.12 ± 0.003
	Known angles	0.04	0.05	0.08	0.11
Scorpion toxin	Frequency marching	0.05 ± 0.002	0.06 ± 0.003	0.06 ± 0.002	0.07 ± 0.004
	Known angles	0.04	0.04	0.05	0.06

ϵ achieved by reconstructing f through a single application of the least-squares procedure (section 5) with all image orientations set to their true values. Table 1 shows that the errors resulting from applying our proposed frequency marching algorithm exceed this best possible error by only around 10^{-2} for all molecules and noise levels.

Figures 4(c),(d), 5(c),(d), and 6(c),(d) show results of the reconstruction for the three models using experimental images at SNR 0.5 and 0.05, respectively. Out of the five runs, we show the reconstructions with the lowest error. Table 2 shows that it took approximately two hours to do the reconstructions using frequency marching on 14 cores for all three models. This table also shows that the time is not significantly affected by the noise levels in the data. In addition, we expect that the quality of the reconstruction will improve as we increase the number of experimental images. This effect is visible in Figure 7, where the median squared error of the runs is surprisingly well fit by the functional form of a constant (accounting for Fourier and image discretization errors) plus a constant times $1/M$, accounting for the usual reduction in statistical variance with a growing data set size M .

8. Conclusions. We have presented a fast, robust algorithm for determining the three-dimensional (3D) structure of proteins using “single particle” cryo-electron microscopy (cryo-EM) data. This problem is typically formulated in the language of global optimization, leading to a nonconvex objective function for the unknown angles and lateral offsets of each particle in a collection of noisy projection images. With hundreds of thousands of images, this leads to a very CPU-intensive task. By using a recursive method in the Fourier domain, we have shown that an essentially deterministic scheme is able to achieve high-resolution reconstruction with predictable and modest computing requirements; in our initial implementation, 50,000 images can be processed in about one hour on a standard multicore desktop workstation.

One important consequence is that, with sufficiently fast reconstruction, one can imagine making use of multiple runs on the same data, opening up classical jackknife, bootstrap, or cross-validation statistics for quality and resolution assessment. In this work, we illustrate the power of doing so with only five runs.

Practical considerations. There are several features which need to be added to the algorithm above for it to be applicable to experimental data. Most critically, we need to extend the set of unknowns for each image to include the translational degrees of freedom (unknown lateral offsets for the particle “center”). We believe that frequency marching can again be used to great effect by defining a fixed search region (either a 3×3 or a 5×5 offset grid)




Figure 4. Results for Rubisco reconstruction. (a),(b) Examples of experimental images at SNR 0.5 and 0.05 (black indicates smallest values, white largest). (c),(d) Reconstructions by frequency marching using experimental images at SNR 0.5 and 0.05, shown as isosurfaces. (e),(f) Examples of a single slice through the reconstructed f (blue indicates smallest values, yellow largest).

whose spatial scale corresponds to the current resolution. That is, we define a spatial step in each lateral direction of the order $O(1/k)$, so that large excursions are tested at low resolution and pixel-scale excursions are tested at high resolution. This would increase the workload by a constant factor. If done naively, the increase would correspond to a multiplicative factor of 9 or 25 on a portion of the code that currently accounts for 50% of the total time. We suspect that improvements in the search strategy (using asymptotic analysis) can reduce this overhead by a substantial factor and that other code optimizations will further reduce the run time (described below). Thus, we expect only a modest net increase in cost over our current timings.

Code optimization. Significant optimizations are possible to further reduce the run time of our template-matching and least-squares reconstruction steps. These include the following:

1. *Using more uniform and sparser sampling in Fourier space.* At frequency k , only $O(k^2)$ templates are needed to resolve the unknown function $F(k, \theta, \phi)$. Our initial implementation uses $O(K^2)$ templates, where K is the maximum frequency of interest. A factor of order 5 speedup is available here. Moreover, by reducing the number of azimuthal points near the poles, the number of templates can be reduced by an additional factor of $\pi/2$.




Figure 5. Results for lipoyxygenase-1 reconstruction. (a),(b) Examples of experimental images at SNR 0.5 and 0.05. (c),(d) Reconstructions by frequency marching using experimental images at SNR 0.5 and 0.05, shown as isosurfaces. (e),(f) Examples of a single slice through the reconstructed f .

2. *Updating only the last shell in Fourier space via the least-squares solve.* Our initial implementation rebuilds the function $F(k, \theta, \phi)$ on all spheres out to radius k at every iteration (see Remark 6).
3. *Using coarser grids at low resolution in frequency marching.* At present, we use the same full resolution spherical grid for every k .
4. *Restricting the angle search in (α, β) to $O(1)$ angles in the neighborhood of the best guess at the previous frequency k .* This would reduce the overall complexity of the step from $O(MK^5)$ to $O(MK^3)$. Similar strategies have been found to be very effective in algorithms that marginalize over angles [5].

While it would increase the cost, it is also worth exploring the use of marginalization over angles, as in full expectation-maximization (EM) approaches, as well as the use of subgrid angle fitting in the projection matching step. One could also imagine exploring the effectiveness of multiple iterations of frequency marching (although preliminary experiments with noisy data did not show significant improvements).

Finally, as noted above, our random initialization is subject to occasional failure. We will investigate the possibility of enforcing convergence both by carrying out a more sophisticated nonlinear search for a low resolution initial guess and by finding metrics by which to rapidly discard diverging trajectories. On a related note, we also plan to develop metrics for discarding images that appear to correspond to erroneously included “nonparticles.” In frequency




Figure 6. Results for scorpion toxin reconstruction. (a),(b) Examples of synthetic experimental images at SNR 0.5 and 0.05. (c),(d) Reconstructions by frequency marching from experimental images at SNR 0.5 and 0.05, shown as isosurfaces. (e),(f) Examples of a single slice through the reconstructed f .

Table 2

Time in hours for the reconstructions on 14 cores, in the context of different levels of noise added to the experimental images. All the times are given as averages with estimated standard deviations, over five different runs.

	No noise	SNR 0.5	SNR 0.1	SNR 0.05
Rubisco	1.6 ± 0.5	1.4 ± 0.3	1.2 ± 0.1	1.4 ± 0.1
Lipoxygenase-1	1.1 ± 0.4	1.5 ± 0.6	1.1 ± 0.4	1.2 ± 0.3
Scorpion toxin	0.5 ± 0.1	0.5 ± 0.1	0.6 ± 0.1	0.7 ± 0.2

marching, we suspect that the hallmark of such nonparticles (at least in the asymmetric setting) will be the failure of template-matching to find more and more localized matches. On a more speculative note, we suspect that frequency marching will be able to handle structural heterogeneity without too many modifications, at least in the setting where there is a finite number of well-defined conformations, say R , present in the data set. The simplest approach would be to assign a conformation label $(1, \dots, R)$ as an additional parameter for each image and to reconstruct R Fourier space models in parallel, with random initialization of the labels, and assignments of the best label recomputed at each stage. We will report on all of the above developments at a later date.

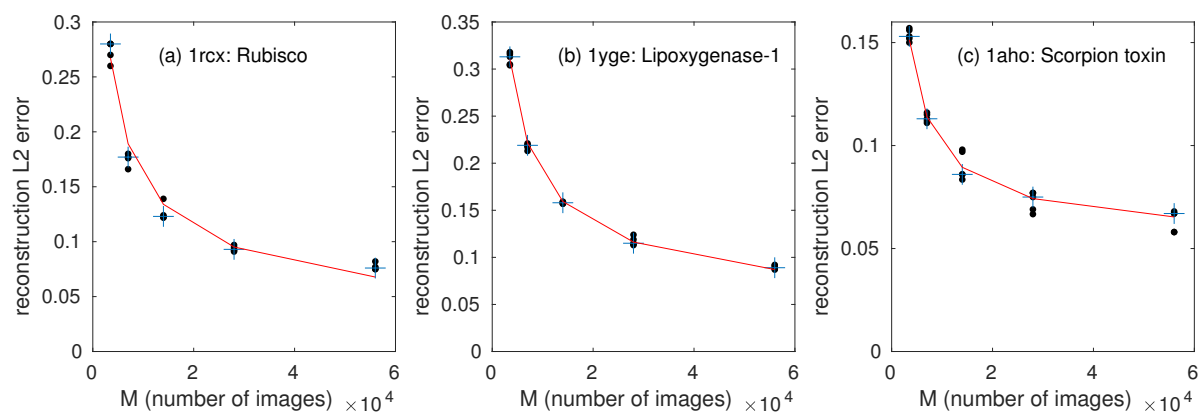


Figure 7. Relative L_2 error ϵ of the reconstructed model (see (7.1)) versus M , the number of simulated experimental images used for the reconstruction (black dots), using five runs for each of the three molecules. The SNR is 0.05. The error decreases as M increases. The median is also shown (plus signs), as is the best fit to the expected form $a_0 + a_1/M$ for the square of the median errors. (Two of the runs failed with our random initialization procedure and are not shown in this figure.)

Acknowledgments. We would like to thank David Hogg, Roy Lederman, Jeremy Magland, Christian Müller, Adi Rangan, Amit Singer, and Doug Renfrew for many useful conversations, and the anonymous referees for their helpful suggestions.

REFERENCES

- [1] G. BAO, P. LI, J. LIN, AND F. TRIKI, *Inverse scattering problems with multi-frequencies*, *Inverse Problems*, 31 (2015), 093001.
- [2] J. M. BELL, M. CHEN, P. R. BALDWIN, AND S. J. LUDTKE, *High resolution single particle refinement in EMAN2.1*, *Methods*, 100 (2016), pp. 25–34.
- [3] H. BERMAN, J. WESTBROOK, Z. FENG, G. GILLILAND, T. N. BHAT, H. WEISSIG, I. N. SHINDYALOV, AND P. E. BOURNE, *The protein data bank*, *Nucleic Acids Res.*, 28 (2000), pp. 235–242, <http://www.rcsb.org>.
- [4] C. BORGES, A. GILLMAN, AND L. GREENGARD, *High resolution inverse scattering in two dimensions using recursive linearization*, *SIAM J. Imaging Sci.*, 10 (2017), pp. 641–664, <https://doi.org/10.1137/16M1093562>.
- [5] M. A. BRUBAKER, A. PUNJANI, AND D. J. FLEET, *Building proteins in a day: Efficient 3D molecular reconstruction*, *IEEE Trans. Pattern Anal. Mach. Intel.*, 39 (2017), pp. 706–718.
- [6] Y. CHEN, *Recursive Linearization for Inverse Scattering*, Yale Research Report/DCS/RR-1088, Department of Computer Science, Yale University, New Haven, CT, 1995.
- [7] Y. CHEN, *Inverse scattering via Heisenberg’s uncertainty principle*, *Inverse Problems*, 13 (1997), pp. 253–282.
- [8] Y. CHENG, N. GRIGORIEFF, P. A. PENCZEK, AND T. WALZ, *A primer to single-particle cryo-electron microscopy*, *Cell*, 161 (2015), pp. 439–449.
- [9] R. COIFMAN, Y. SHKOLNISKY, F. J. SIGWORTH, AND A. SINGER, *Reference free structure determination through eigenvectors of center of mass operators*, *Appl. Comput. Harmon. Anal.*, 47 (2010), pp. 296–312.
- [10] Y. CONG, J. A. KOVACS, AND W. WRIGGERS, *2D fast rotational matching for image processing of biophysical data*, *J. Struct. Biol.*, 144 (2003), pp. 51–60.
- [11] N. C. DVORNEK, F. J. SIGWORTH, AND H. D. TAGARE, *SubspaceEM: A fast maximum-a-posteriori algorithm for cryo-EM single particle reconstruction*, *J. Struct. Biol.*, 190 (2015), pp. 200–214.

- [12] D. ELMLUND AND H. ELMLUND, *SIMPLE: Software for ab initio reconstruction of heterogeneous single-particles*, J. Struct. Biol., 180 (2012), pp. 420–427.
- [13] J. FRANK, B. SHIMKIN, AND H. DOWSE, *SPIDER—a modular software system for electron image processing*, Ultramicroscopy, 6 (1981), pp. 343–358.
- [14] A. B. GONCHAROV, *Methods of integral geometry and the determination of the mutual orientation of identical particles arbitrarily arranged in a plane from their projections on a line*, Dokl. Akad. Nauk SSSR, 293 (1987), pp. 355–358 (in Russian).
- [15] A. B. GONCHAROV AND M. S. GELFAND, *Determination of mutual orientation of identical particles from their projections by the moments method*, Ultramicroscopy, 25 (1988), pp. 317–328.
- [16] L. GREENGARD AND J.-Y. LEE, *Accelerating the nonuniform fast Fourier transform*, SIAM Rev., 46 (2004), pp. 443–454, <https://doi.org/10.1137/S003614450343200X>.
- [17] N. GRIGORIEFF, *FREALIGN: High-resolution refinement of single particle structures*, J. Struct. Biol., 157 (2007), pp. 117–125.
- [18] N. GRIGORIEFF, *Frealign: An exploratory tool for single-particle cryo-EM*, in The Resolution Revolution: Recent Advances in cryoEM, Methods in Enzymology 579, R. A. Crowther, ed., Academic Press, New York, 2016, pp. 191–226.
- [19] R. HENDERSON, *Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise*, Proc. Natl. Acad. Sci. U.S.A., 110 (2013), pp. 18037–18041.
- [20] R. HENDERSON, A. SALI, M. L. BAKER, B. CARRAGHER, B. DEVKOTA, K. H. DOWNING, E. H. EGELMAN, Z. FENG, J. FRANK, N. GRIGORIEFF, W. JIANG, S. J. LUDTKE, O. MEDALIA, P. A. PENCZEK, P. B. ROSENTHAL, M. G. ROSSMANN, M. F. SCHMID, G. F. SCHRÖDER, A. C. STEVEN, D. L. STOKES, J. D. WESTBROOK, W. WRIGGERS, H. YANG, J. YOUNG, H. M. BERMAN, W. CHIU, G. J. KLEYWEGT, AND C. L. LAWSON, *Outcome of the first electron microscopy validation task force meeting*, Structure, 20 (2012), pp. 205–214.
- [21] J. B. HEYMAN, *Validation of 3D EM reconstructions: The phantom in the noise*, AIMS Biophys., 2 (2015), pp. 21–35.
- [22] M. HOHN, G. TANG, G. GOODYEAR, P. BALDWIN, Z. HUANG, P. PENCZEK, C. YANG, R. GLAESER, P. ADAMS, AND S. LUDTKE, *Sparx, a new environment for cryo-em image processing*, J. Struct. Biol., 157 (2007), pp. 47–55.
- [23] P. JOUBERT AND M. HABECK, *Bayesian inference of initial models in cryo-electron microscopy using pseudo-atoms*, Biophysical Journal, 108 (2015), pp. 1165–1175.
- [24] L. JOYEUX AND P. A. PENCZEK, *Efficiency of 2D alignment methods*, Ultramicroscopy, 92 (2002), pp. 33–46.
- [25] D. KIMANIUS, B. O. FORSBERG, S. H. SCHERES, AND E. LINDEHL, *Accelerated cryo-EM structure determination with parallelisation using GPUs in RELION-2*, eLife, 5 (2016), e18722.
- [26] D. LYUMKIS, A. F. BRILLOT, D. L. THEOBALD, AND N. GRIGORIEFF, *Likelihood-based classification of cryo-EM images using FREALIGN*, J. Struct. Biol., 183 (2013), pp. 377–388.
- [27] J. L. S. MILNE, M. J. BORGNA, A. BARTESAGHI, E. E. H. TRAN, L. A. EARL, D. M. SCHAUDER, J. LENGYEL, J. PIERSON, A. PATWARDHAN, AND S. SUBRAMANIAM, *Cryo-electron microscopy—a primer for the non-microscopist*, FEBS J., 280 (2013), pp. 28–45.
- [28] J. A. MINDELL AND N. GRIGORIEFF, *Accurate determination of local defocus and specimen tilt in electron microscopy*, J. Struct. Biol., 142 (2003), pp. 334–347.
- [29] F. NATTERER, *The Mathematics of Computerized Tomography*, Classics Appl. Math. 21, SIAM, Philadelphia, 2001, <https://doi.org/10.1137/1.9780898719284>.
- [30] E. NOGALES, *The development of cryo-EM into a mainstream structural biology technique*, Nature Methods, 13 (2016), pp. 24–27.
- [31] P. PENCZEK, M. RADERMACHER, AND J. FRANK, *Three-dimensional reconstruction of single particles embedded in ice*, Ultramicroscopy, 40 (1992), pp. 33–53.
- [32] P. A. PENCZEK, *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*, Oxford University Press, Oxford, UK, 2006.
- [33] P. A. PENCZEK, *Resolution measures in molecular electron microscopy*, Methods Enzymol., 482 (2010), pp. 73–100.
- [34] P. A. PENCZEK, R. A. GRASSUCCI, AND J. FRANK, *The ribosome at improved resolution: New techniques for merging and orientation refinement in 3D cryo-electron microscopy of biological particles*, Ultramicroscopy, 53 (1994), pp. 251–270.

- [35] P. A. PENCZEK, J. ZHU, AND J. FRANK, *A common-lines based method for determining orientations for $N > 3$ particle projections simultaneously*, *Ultramicroscopy*, 63 (1996), pp. 205–218.
- [36] S. W. PROVENCHER AND R. H. VOGEL, *Three-dimensional reconstruction from electron micrographs of disordered specimens I. Method*, *Ultramicroscopy*, 25 (1988), pp. 209–221.
- [37] P. B. ROSENTHAL AND J. L. RUBINSTEIN, *Validating maps from single particle electron cryomicroscopy*, *Current Opinion Struct. Biol.*, 34 (2015), pp. 135–144.
- [38] E. SANZ-GARCÍA, A. B. STEWART, AND D. M. BELNAP, *The random-model method enables ab initio three-dimensional reconstruction of asymmetric particles and determination of particle symmetry*, *J. Struct. Biol.*, 171 (2010), pp. 216–222.
- [39] S. H. W. SCHERES, *A Bayesian view on cryo-EM structure determination*, *J. Mol. Biol.*, 415 (2012), pp. 406–418.
- [40] S. H. W. SCHERES, *RELION: Implementation of a Bayesian approach to cryo-EM structure determination*, *J. Struct. Biol.*, 180 (2012), pp. 519–530.
- [41] S. H. W. SCHERES, M. VALLE, P. GROB, E. NOGALES, AND J.-M. CARAZO, *Maximum likelihood refinement of electron microscopy data with normalization errors*, *J. Struct. Biol.*, 166 (2009), pp. 234–240.
- [42] Y. SHKOLNISKY AND A. SINGER, *Viewing direction estimation in cryo-EM using synchronization*, *SIAM J. Imaging Sci.*, 5 (2012), pp. 1088–1110, <https://doi.org/10.1137/120863642>.
- [43] F. J. SIGWORTH, *A maximum-likelihood approach to single-particle image refinement*, *J. Struct. Biol.*, 122 (1998), pp. 328–339.
- [44] F. J. SIGWORTH, P. C. DOERSCHUK, J.-M. CARAZO, AND S. H. W. SCHERES, *An introduction to maximum-likelihood methods in cryo-EM*, in *Cryo-EM, Part B: 3D Reconstruction*, *Methods in Enzymology* 482, G. J. Jensen, Academic Press, New York, 2010, pp. 263–294.
- [45] A. SINGER, R. R. COIFMAN, F. J. SIGWORTH, D. W. CHESTER, AND Y. SHKOLNISKY, *Detecting consistent common lines in cryo-EM by voting*, *J. Struct. Biol.*, 169 (2009), pp. 312–322.
- [46] A. SINGER AND Y. SHKOLNISKY, *Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming*, *SIAM J. Imaging Sci.*, 4 (2011), pp. 543–572, <https://doi.org/10.1137/090767777>.
- [47] G. TANG, L. PENG, P. BALDWIN, D. MANN, W. JIANG, I. REES, AND S. LUDTKE, *EMAN2: An extensible image processing suite for electron microscopy*, *J. Struct. Biol.*, 157 (2007), pp. 38–46.
- [48] L. N. TREFETHEN AND J. A. C. WEIDEMAN, *The exponentially convergent trapezoidal rule*, *SIAM Rev.*, 56 (2014), pp. 385–458, <https://doi.org/10.1137/130932132>.
- [49] B. VAINSHTEIN AND A. GONCHAROV, *Determination of the spatial orientation of arbitrary arranged identical particles of an unknown structure from their projections*, in *Proceedings of the 11th International Congress on Electron Microscopy*, Kyoto, Japan, 1986, pp. 459–460.
- [50] M. VAN HEEL, *Angular reconstruction: A posteriori assignment of projection directions for 3D reconstruction*, *Ultramicroscopy*, 21 (1987), pp. 111–123.
- [51] M. VAN HEEL, E. V. ORLOVA, G. HARAUZ, H. STARK, P. DUBE, F. ZEMLIN, AND M. SCHATZ, *Angular reconstitution in three-dimensional electron microscopy: Historical and theoretical aspects*, *Scanning Microscopy*, 11 (1997), pp. 195–210.
- [52] K. R. VINOTHKUMAR AND R. HENDERSON, *Single particle electron cryomicroscopy: Trends, issues and future perspective*, *Quart. Rev. Biophys.*, 49 (2016), e13.
- [53] R. H. VOGEL AND S. W. PROVENCHER, *Three-dimensional reconstruction from electron micrographs of disordered specimens II. Implementation and results*, *Ultramicroscopy*, 25 (1988), pp. 223–240.
- [54] R. H. WADE, *A brief look at imaging and contrast transfer*, *Ultramicroscopy*, 46 (1992), pp. 145–156.
- [55] L. WANG, A. SINGER, AND Z. WEN, *Orientation determination of cryo-EM images using least unsquared deviations*, *SIAM J. Imaging Sci.*, 6 (2013), pp. 2450–2483, <https://doi.org/10.1137/130916436>.
- [56] J. A. C. WEIDEMAN AND L. N. TREFETHEN, *The kink phenomenon in Fejér and Clenshaw–Curtis quadrature*, *Numer. Math.*, 107 (2007), pp. 707–727.