Published in final edited form as:

Conf Rec Asilomar Conf Signals Syst Comput. 2011 ; : 1026–1030. doi:10.1109/ACSSC.2011.6190167.

Surface-Constrained 3D Reconstruction in Cryo-EM

Andrew C. Barthel,

Department of Biomedical Engineering, Yale University, New Haven, Connecticut, 06520

Hemant Tagare, and

Department of Biomedical Engineering, Department of Diagnostic Radiology, Yale University, New Haven, Connecticut, 06520

Fred J. Sigworth

Department of Biomedical Engineering, Department of Cellular and Molecular Physiology, Yale University, New Haven, Connecticut, 06520

Andrew C. Barthel: andrew.barthel@yale.edu; Hemant Tagare: hemant.tagare@yale.edu; Fred J. Sigworth: fred.sigworth@yale.edu

Abstract

Random spherically-constrained (RSC) reconstruction is a new form of single particle reconstruction (SPR) using cryo-EM images of membrane proteins embedded in spherical lipid vesicles to generate a 3D protein structure. The method has many advantages over conventional SPR, including a more native environment for protein particles and an initial estimate of the particle's angular orientation. These advances allow us to determine structures of membrane proteins such as ion channels and derive more reliable structure estimates. We present an algorithm that relates conventional SPR to the RSC model, and generally, to projection images of particles embedded with an axis parallel to the local normal of a general 2D manifold. We illustrate the performance of this algorithm in the spherical system using synthetic data.

I. Introduction

Classical cryo-EM reconstructs 3D structures of biological macromolecules by freezing purified macromolecules in ice and obtaining a transmission electron microscope image of the particles. A major limitation of this technique is that the particles are not in their natural environment and need not assume their natural conformation. This is especially problematic when imaging transmembrane proteins, which are purified and imaged without the membrane.

Surface-constrained cryo-EM first embeds these particles in a surface (typically a lipid membrane) before freezing and imaging [1]. Besides providing a membrane, this allows the conformation of channel proteins to be controlled by electrical or chemical gradients across the membrane [2].

The surface-constrained cryo-EM process is illustrated in fig. 1. The particle is embedded in the (e.g. spherical) surface in such a way that the particle's z-axis points along the surface normal, with an arbitrary rotation of the particle around the normal. An image of the embedded particle is then obtained by projecting along the vesicle's z-axis. The user specifies the approximate center of the particle in the image. Using multiple images and associated center coordinates, the particle structure is reconstructed. In the spherical case, each particle center is recorded relative to the center of the vesicle it's embedded in along with the vesicle radius. Vesicle images are well understood in cryo-EM. Their radius and center can be found quite accurately, and they can be numerically subtracted from the collected images, leaving only the macromolecule particles and image noise [3]. The

relationship between the particle location and orientation yields extra information in the surface-constrained data collection process.

Some ambiguities remain in the surface-constrained cryo-EM process. Whether a particle points in or out of the vesicle is not controlled in the vesicle formation process, adding a two-fold ambiguity in the particle's angular orientation. In addition, there is an angular ambiguity of whether the particle is on the upper (as in fig. 1) or lower hemisphere relative to the image projection. These ambiguities are resolved in the reconstruction algorithm.

This paper poses the surface-constrained cryo-EM reconstruction problem in a maximum-likelihood (ML) framework. A reconstruction algorithm using the ML framework is also described.

A. Notation and conventions

We begin by defining our notation and conventions. Most of the mathematical development that follows is coordinate free.

1) Spaces—We use the standard 2D and 3D spaces, $\mathbb{R}^2$ and $\mathbb{R}^3$. Points in these spaces are denoted bold face as $\mathbf{u}$, $\mathbf{v}$ etc. When we need coordinates, we write $\mathbf{u} = (u_x, u_y)$ or $\mathbf{u} = (u_x, u_y, u_z)$.

Below, we will need the projection of 3D points onto the x, y plane. We take $p_z: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ to be projection operator that projects along the z-axis, $p_z(\mathbf{u}) = (u_x, u_y)$.

Unit coordinate vectors in $\mathbb{R}^2$ and $\mathbb{R}^3$ are $\mathbf{i}, \mathbf{j}$ and $\mathbf{i}, \mathbf{j}, \mathbf{k}$ respectively.

2) Rotations and translations—Rotations in $\mathbb{R}^3$ can be parameterized in many different ways. We adopt a parametrization that is convenient for our problem. Any rotation maps the unit coordinate vectors $\mathbf{i}, \mathbf{j}, \mathbf{k}$ to another set of orthonormal vectors $\mathbf{i}', \mathbf{j}', \mathbf{k}'$. This rotation can be taken to be a composition of two simpler rotations. The first is around the rotation axis $\mathbf{k} \times \mathbf{k}'$ (fig. 2) and takes $\mathbf{k}$ to $\mathbf{k}'$. The second is around $\mathbf{k}'$ takes $\mathbf{i}, \mathbf{j}$ to $\mathbf{i}', \mathbf{j}'$. Thus every 3D rotation can be parameterized by a unit vector $\mathbf{k}'$ and an angle ϕ . To exhibit this dependence, we write the rotation operator as $R_{\mathbf{k}, \mathbf{k}', \phi}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. The notation suggests that the rotation first aligns $\mathbf{k}$ to $\mathbf{k}'$ and then rotates around $\mathbf{k}'$ by ϕ .

The rotation operator on $\mathbb{R}^2$ is denoted as R_θ where θ is angle of rotation. The number of subscripts distinguishes the 2D rotation operator from the 3D rotation operator.

The translation operator is denoted $T_{\mathbf{v}}$ where $\mathbf{v}$ is a vector in $\mathbb{R}^2$ or $\mathbb{R}^3$. This operator maps the origin to $\mathbf{v}$.

The inverse of the translation operator is $(T_{\mathbf{v}})^{-1} = T_{-\mathbf{v}}$. The inverse of 2D rotation is $(R_\theta)^{-1} = R_{-\theta}$. The inverse of a 3D rotation does not have a simple form. However, it is easy to see that $R_{\mathbf{k}, \mathbf{k}', \phi}^{-1} = R_{\mathbf{k}', \mathbf{k}, \zeta}$ for some angle ζ .

3) Particle structure—By a *particle structure*, or simply a *structure*, we mean a function $S: \mathbb{R}^3 \rightarrow \mathbb{R}$. Its value $S(\mathbf{u})$ gives the electron scattering density of the physical particle at the point $\mathbf{u} \in \mathbb{R}^3$.

As defined above, rotation and translation operate on points of $\mathbb{R}^2$ or $\mathbb{R}^3$. There is a natural extension by which rotation and translation act on functions, i.e. structures. Let L be a rotation or translation operator and let $S: \mathbb{R}^n \rightarrow \mathbb{R}$, $n = 2, 3$ be a structure (function). Then $S \circ L^{-1}$ is the structure S “rotated” or “translated” in the same way as L acting on points of $\mathbb{R}^2$ or

$\mathbb{R}^3$. In cryo-EM it is conventional to talk of “rotating or translating the structure” suggesting that rotation and translation act on the left of S . To mimic this, we introduce a notation. The dual operator L^* acts on S by

$$L^* \circ S = S \circ L^{-1}$$

so that L^* rotates or translates S in the same way that L maps points of $\mathbb{R}^2$ or $\mathbb{R}^3$. Thus, $R_{\mathbf{k}, \mathbf{k}', \zeta}^* \circ S$ rotates the structure S in the same way that $R_{\mathbf{k}, \mathbf{k}', \zeta}$ maps points of $\mathbb{R}^3$.

4) Images—An *image* is a function from $\mathbb{R}^2$ to $\mathbb{R}$. Images can be rotated and translated by using the duals of 2D rotation and translation operators.

5) Ray projection—The *ray projection* operator Σ_z gives the integral of a structure along vertical lines. For any structure S , the operation $\Sigma_z \circ S$ gives an image whose value at $\mathbf{u} \in \mathbb{R}^2$ is $(\Sigma_z \circ S)(\mathbf{u}) = \int S(\mathbf{u} + z\mathbf{k})dz$.

6) Contrast transfer function of the electron microscope—The *CTF operator* C applies the contrast transfer function to an image I , i.e. it convolves the image I with the CTF kernel.

7) Linearity—Dual rotations, translations, ray projection, and the CTF operator are all linear operators.

B. Surface-constrained cryo-EM

Although our experiments use a sphere for the particle-embedding surface, we develop the theory of surface-constrained cryo-EM for any compact surface. In keeping with modern differential geometric tradition, we develop the theory in a coordinate-free manner.

As in fig. 3, let $\mathbf{u}$ be a point on the surface and let the surface normal at $\mathbf{u}$ be $\mathbf{n}(\mathbf{u})$. Embedding a structure S at $\mathbf{u}$ means first rotating the structure so that the z -axis of the structure points along $\pm \mathbf{n}(\mathbf{u})$ and then translating the rotated structure to $\mathbf{u}$. The rotation of the structure also includes an arbitrary rotation around $\mathbf{n}(\mathbf{u})$, ϕ . The particle may insert into the membrane inside-out, which we model with a variable that can take values of 1 or -1 , λ .

Thus the embedded structure is $T_{\mathbf{u}}^* \circ R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \phi}^* \circ S$. Projecting this along the z -axis, and applying the CTF gives the image I . Thus,

$$I = C \circ \sum_z \circ T_{\mathbf{u}}^* \circ R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \phi}^* \circ S.$$

This equation can be simplified by using the following properties. The properties are straightforward to prove.

1. $\sum_z \circ T_{\mathbf{u}}^* = T_{p_z(\mathbf{u})}^* \circ \sum_z$. This property simply says that translating a structure in 3D followed by ray projection (the left hand side) is identical to a ray projection of the untranslated structure followed by a translation of the 2D projection (by the x, y components) of the 3D translation (the right hand side).
2. $C \circ T_{\mathbf{v}}^* = T_{\mathbf{v}}^* \circ C$, i.e. the CTF operator commutes with image translation.

Using these properties we get:

$$\begin{aligned} I &= C \circ \sum_z T_{\mathbf{u}}^* \circ R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \varphi}^* \circ S \\ &= C \circ T_{p_z(\mathbf{u})}^* \circ \sum_z R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \varphi}^* \circ S \\ &= T_{p_z(\mathbf{u})}^* \circ C \circ \sum_z R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \varphi}^* \circ S. \end{aligned}$$

In reality, the measured image is noisy, i.e.

$$I = T_{p_z(\mathbf{u})}^* \circ C \circ \sum_z R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}), \varphi}^* \circ S + \varepsilon_I, \quad (1)$$

where ε_I is the noise.

Besides the image, our data includes the center coordinates of the particle as selected by the user. Let $\mathbf{c}$ be the point in the image clicked by the user. Then

$$\mathbf{c} = p_z(\mathbf{u}) + \varepsilon_c, \quad (2)$$

where ε_c is noise.

Equations (1–2) relate the measurements to the structure. The fact that both equations are related by a similar term, $p_z(\mathbf{u})$, is what differentiates RSC from conventional SPR, where both equations are unrelated.

II. Maximum likelihood Formulation

Several images and center coordinates are obtained in a cryo-EM study. Each image has its individual embedding of the particle on the surface, CTF, and the user-selected center point. Letting $i = 1, \dots, N$ index the number of images,

$$\begin{aligned} I_i &= T_{p_z(\mathbf{u}_i)}^* \circ C_i \circ \sum_z R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}_i), \varphi_i}^* \circ S + \varepsilon_{I_i}, \\ \mathbf{c}_i &= p_z(\mathbf{u}_i) + \varepsilon_{c_i}, \end{aligned}$$

Assuming that noise ε_{I_i} is i.i.d. at every pixel of every image and is normally distributed with zero mean and std. dev. σ_1 , and assuming that noise ε_{c_i} is i.i.d. and normal with zero mean and std. dev. σ_2 ,

$$\begin{aligned} &\log p(\{I_i\} | S, \{\mathbf{u}_i\}, \{\varphi_i\}, \sigma_1, \sigma_2) \\ &= \sum_i -\|I_i - T_{p_z(\mathbf{u}_i)}^* \circ C_i \circ \sum_z R_{\mathbf{k}, \lambda \mathbf{n}(\mathbf{u}_i), \varphi_i}^* \circ S\|^2 / 2\sigma_1^2 \\ &\quad - P \log \sigma_1 \\ &\quad - \|\mathbf{c}_i - p_z(\mathbf{u}_i)\|^2 / 2\sigma_2^2 - \log \sigma_2, \end{aligned} \quad (3)$$

where P is the number of pixels in each image.

Maximum-likelihood estimates of the structure and other parameters are

$$\hat{S}, \{\hat{\mathbf{u}}_i\}, \{\hat{\varphi}_i\}, \hat{\sigma}_1, \hat{\sigma}_2 = \arg \max_{S, \{\mathbf{u}_i\}, \{\varphi_i\}, \sigma_1, \sigma_2} \log p(\{I_i\} | S, \{\mathbf{u}_i\}, \{\varphi_i\}, \sigma_1, \sigma_2).$$

The maximization is carried out by maximizing one group of variables at a time, in an iterative fashion. This is typical of cryo-EM algorithms [4], [5]:

Maximize the log-likelihood

1. **Initialize:** Set $n = 0$, and $\{\hat{\mathbf{u}}_i\}, \{\hat{\varphi}_i\}, \hat{\sigma}_1, \hat{\sigma}_2$ to initial values.
2. **Update Estimates:**
 - a. **Update noise variance:**

$$\hat{\sigma}_1^{n+1}, \hat{\sigma}_2^{n+1} = \arg \max_{\sigma_1, \sigma_2} \log p(\{I_i\} | \hat{S}^n, \{\hat{\mathbf{u}}_i\}^n, \{\hat{\varphi}_i\}^n, \sigma_1, \sigma_2).$$

- b. **Update embedding:**

$$\{\hat{\mathbf{u}}_i\}^{n+1}, \{\hat{\varphi}_i\}^{n+1} = \arg \max_{\{\mathbf{u}_i\}, \{\varphi_i\}} \log p(\{I_i\} | \hat{S}^n, \{\hat{\mathbf{u}}_i\}, \{\hat{\varphi}_i\}, \hat{\sigma}_1^{n+1}, \hat{\sigma}_2^{n+1}).$$

- c. **Update Structure:**

$$\hat{S}^{n+1} = \arg \max_S \log p(\{I_i\} | S, \{\hat{\mathbf{u}}_i\}^{n+1}, \{\hat{\varphi}_i\}^{n+1}, \hat{\sigma}_1^{n+1}, \hat{\sigma}_2^{n+1}).$$

3. **Test for Convergence:** If the parameters have not converged, set $n = n + 1$, go to step 2.

The updates (steps 2(a)–2(c) in the above algorithm) are calculated as follows.

1) Update noise variance—The noise variance update has a closed-form solution:

$$\begin{aligned} \hat{\sigma}_1^{n+1} &= \text{sqrt} \left\{ \frac{1}{NP} \sum_i \|I_i - T_{p_z(\hat{\mathbf{u}}_i^n)}^* \circ C_i \circ \sum_z \circ R_{\mathbf{k}, \lambda \mathbf{n}(\hat{\mathbf{u}}_i^n), \hat{\varphi}_i^n}^* \circ \hat{S}^n\|^2 \right\}, \\ \hat{\sigma}_2^{n+1} &= \text{sqrt} \left\{ \frac{1}{N} \sum_i \|\mathbf{c}_i - p_z(\hat{\mathbf{u}}_i^n)\|^2 \right\}. \end{aligned}$$

2) Update embedding—Maximizing the log-likelihood with respect to $\{\mathbf{u}_i\}$ and $\{\varphi_i\}$ does not have a closed form solution, therefore a numerical strategy is necessary. Because the log-likelihood function depends on $\{\mathbf{u}_i\}$ and $\{\varphi_i\}$ in a highly non-convex manner, numerical strategies such as gradient ascent can get trapped in local maxima. In traditional cryo-EM, this problem is overcome by a brute-force search. We follow the same strategy. The constraining surface is covered with a finite number of points. Because the surface is compact, a high sampling density can be obtained by a finite number of points. Let U be the set of these points. Further, the range of ϕ , $[0, 2\pi)$, is also sampled at a finite number of angles. Let Φ be the set of these angles. The maximization with respect to $\{\mathbf{u}_i\}, \{\varphi_i\}$ is carried out by a brute force search with $(\mathbf{u}_i, \varphi_i) \in U \times \Phi$:

$$\{\widehat{\mathbf{u}}_i\}^{n+1}, \{\widehat{\varphi}_i\}^{n+1} = \arg \max_{(\mathbf{u}_i, \varphi_i) \in U \times \Phi, i=1, \dots, N} \log p(\{I_i\} | \widehat{S}^n, \{\widehat{\mathbf{u}}_i\}, \{\widehat{\varphi}_i\}, \widehat{\sigma}_1^{n+1}, \widehat{\sigma}_2^{n+1}).$$

This maximization comes down to the following: For each image (i.e. each i), find the embedding location and angle pair in $U \times \Phi$ at which the CTF-filtered projection of the structure best matches the image.

3) Update Structure—Because dual rotation, ray projection, CTF, and dual translation are linear, the log-likelihood function of equation (3) is quadratic in S and the maximization has a closed form solution. But, because S is high-dimensional (represented with approximately 10^6 voxels), the closed form solution is not practical to implement. Instead we use the conjugate gradient algorithm to minimize the negative log-likelihood with respect to S . In practice, only a few iterations (7 iterations seem sufficient) of the conjugate gradient algorithm are needed. These iterations form an inner loop executed within step 2(c) of the maximization procedure.

III. Methods

We implemented this algorithm primarily in MATLAB, with some assistance from C-functions accessed with MEX-files. We simulated several sets of 3000 data images using a potassium channel, Kv1.2, particle density constructed from the coordinates of protein data bank entry 3LUT at 3Å/pixel resolution [6]. We mapped this density in a 90×90×90 volume and used 90×90 pixel images, fig. 4. A rotated and translated particle projection was convolved with a unique CTF for every image, then white Gaussian noise was added to every pixel. Unique CTFs were calculated with variations in both B-factor and defocus, uniform in ranges 50–350Å² and 1–6 microns respectively (for explanations of these parameters, please see e.g. [5]).

Along with each image, a location relative to the vesicle center was generated matching the rotational component of the particle with a unique vesicle radius; this location was perturbed by up to a pixel or two to model the user “click error” in particle picking. Vesicle radii varied uniformly between 200–500Å. We initialize the algorithm with a low-pass filtered structural density, filtered to about 20Å resolution, similar to low-resolution data typical at the start of a cryo-EM investigation, fig. 5. We used about 450 projection directions to match the particle orientations, resulting in approximately 6 degrees separation between any two nearest projections, and about 10 iterations of the entire algorithm.

Each particle image was matched with up to two directions of particle projections, with the angular orientations of the upper hemisphere right-side-out particle, and the upper hemisphere inside-out. For each of these directions, the entire range of ϕ , the self-rotation, was searched. These preliminary data do not include simulated particles from the lower hemisphere. Assignment errors across hemispheres can impose an additional symmetry axis not seen in the particle.

IV. Discussion

The simulated data both significantly increased the apparent resolution of the structure, and matched images with the correct projection directions in low noise simulations. At higher noise levels, the algorithm naturally started to show more errors. The Fourier shell correlation (FSC) curve is a commonly used metric for judging resolution [5]. It aims to judge the similarity between different structures at different spatial frequencies, but how best to interpret it is not uniformly accepted [7], [8]. Our use of it is atypical because we compare

with the ground truth structure, known to be accurate at high resolution. The FSC is usually used to estimate maximum resolution within a single dataset, by comparing two volumes independently generated from half-datasets [9]. The absolute resolution is limited by the SNR, number and distribution of projection directions, number of voxels, and number of data images.

Figure 6 shows the reconstructed volume from the noisiest of our tests, where the signal to noise ratio was approximately 1:25, an experimentally realistic value [10], [11]. We estimate this reconstruction to have a resolution of about 12Å.

Figure 7 shows several FSC curves comparing the estimated volumes to the true structure; initially (alternate length dashed line), after the first iteration (long dashed line), and after the final iteration (solid).

Figure 8 shows a strict error rate as a percentage of total mis-assigned images for each noise level tested (solid line). Though errors increased with noise variance, they generally did not yet affect the quality of the reconstruction. There are two potential reasons for this, first that our datasets were not noisy enough, and second that the mis-assignment is typically to a nearby projection; only a minor mis-alignment. These errors may not significantly affect the reconstruction, as they still have much of the same data as their neighboring projections. The error rate of significantly mis-aligned (i.e. greater than 6 degrees off) particles was quite low through all tests (fig. 8 dashed line).

These simple tests are but an initial demonstration of the worth of this algorithm. We intend further testing with higher noise to find the limits of its utility, and other surfaces (including using both the upper and lower hemispheres of spherical vesicles, better representing experimental data), as well as tests on experimental data. Modifications to the algorithm will need to be considered if particle embedding is not as simple as we described (i.e. the particle doesn't align perfectly with the surface normal, or bulges in or out of the surface plane).

Acknowledgments

The authors would like to thank Liguang Wang for discussions of methodology and practical implementation, and Yunhui Liu for programming assistance.

References

1. Jiang Q, Chester D, Sigworth F. Spherical reconstruction: a method for structure determination of membrane proteins from cryo-EM images. *Journal of structural biology*. 2001; 133(2–3):119–131. [PubMed: 11472084]
2. Wang L, Sigworth F. Structure of the BK potassium channel in a lipid membrane from electron cryomicroscopy. *Nature*. 2009; 461(7261):292–295. [PubMed: 19718020]
3. Wang L, Bose P, Sigworth F. Using cryo-EM to measure the dipole potential of a lipid membrane. *Proceedings of the National Academy of Sciences*. 2006; 103(49):18528.
4. Sigworth F. A maximum-likelihood approach to single-particle image refinement. *Journal of structural biology*. 1998; 122(3):328–339. [PubMed: 9774537]
5. Frank, J. Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state. Oxford University Press; USA: 2006.
6. Chen X, Wang Q, Ni F, Ma J. Structure of the full-length Shaker potassium channel Kv1. 2 by normal-mode-based X-ray crystallographic refinement. *Proceedings of the National Academy of Sciences*. 2010; 107(25):11352.
7. van Heel M, Schatz M. Fourier shell correlation threshold criteria. *Journal of structural biology*. 2005; 151(3):250–262. [PubMed: 16125414]

8. Rosenthal P, Henderson R. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *Journal of molecular biology*. 2003; 333(4): 721–745. [PubMed: 14568533]
9. Harauz G, van Heel M. Exact filters for general geometry three dimensional reconstruction. *OPTIK*. 1986; 73(4):146–156.
10. Fu J, Gao H, Frank J. Unsupervised classification of single particles by cluster tracking in multi-dimensional space. *Journal of structural biology*. 2007; 157(1):226–239. [PubMed: 16931050]
11. Baxter W, Grassucci R, Gao H, Frank J. Determination of signal-to-noise ratios and spectral SNRs in cryo-EM low-dose imaging of molecules. *Journal of structural biology*. 2009; 166(2):126–132. [PubMed: 19269332]

Fig. 1.
Simplified illustration of (spherical) surface-constrained cryo-EM.

Fig. 2.
3D rotation as a sequence of two rotations.

Fig. 3.
Model of surface-constrained cryo-EM image formation.

Fig. 4.
An isosurface representation of the Kv1.2 structure.

Fig. 5.
The initial low-pass filtered structural estimate.

Fig. 6.
The reconstructed volume.

Fig. 7. Fourier shell correlation progression by iteration, showing initial model (alternate length dashed line), first iteration (long dashed line), and final iteration (solid).

Fig. 8.
Total error (solid line) and large error (dashed) percentages by noise variance.