Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2015 April ; 2015: 200–204. doi:10.1109/ISBI.2015.7163849.

COVARIANCE ESTIMATION USING CONJUGATE GRADIENT FOR 3D CLASSIFICATION IN CRYO-EM

Joakim Andén^{*}, Eugene Katsevich[†], and Amit Singer^{*}

^{*} Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ

[†] Department of Statistics, Stanford University, Stanford, CA

Abstract

Classifying structural variability in noisy projections of biological macromolecules is a central problem in Cryo-EM. In this work, we build on a previous method for estimating the covariance matrix of the three-dimensional structure present in the molecules being imaged. Our proposed method allows for incorporation of contrast transfer function and non-uniform distribution of viewing angles, making it more suitable for real-world data. We evaluate its performance on a synthetic dataset and an experimental dataset obtained by imaging a 70S ribosome complex.

Keywords

Cryo-EM; 3D reconstruction; single particle reconstruction; heterogeneity; structural variability; classification; covariance; conjugate gradient

1. INTRODUCTION

A variety of techniques exist to estimate the structure of biological macromolecules: X-ray crystallography, nuclear magnetic resonance (NMR) spectroscopy, and cryo-electron microscopy (Cryo-EM). While X-ray crystallography provides the best resolution, it requires crystallization – a challenging task for biological molecules. NMR spectroscopy is limited in only being suitable for small molecules (< 50 kDa). Cryo-EM does not require crystallization of the molecules and recent advances in detector technology enable structure determination at near-atomic resolution for sizes greater than 200 kDa [1, 2]. However, biological molecules are prone to radiation damage, so the electron dose is limited ($5\text{--}20$ $e^-/\text{\AA}^2$), resulting in images with poor signal-to-noise ratio (SNR). Dealing with high noise levels is therefore crucial to any Cryo-EM analysis.

In this work, we consider single particle reconstruction (SPR) Cryo-EM. Here, we assume that identical copies of a single molecule are rapidly frozen in a thin layer of vitreous ice and kept frozen during imaging. Each copy then has a random orientation and position. To reconstruct the molecular structure from these noisy projections, we first estimate the Euler angles describing each orientation. Classical tomographic inversion methods are then applied to obtain a three-dimensional voxel structure, or “volume,” representing the molecule. Using this volume, Euler angles are re-estimated, and the process is repeated until convergence. This scheme, known as iterative refinement, can be implemented using a variety of algorithms at each step [3, 4].

Most such algorithms assume that all molecules imaged have the same structure. However, this assumption is often invalid since many molecules exist in a multitude of states. Modeling of multiple molecular states in a Cryo-EM dataset is known as the heterogeneity problem and has attracted much attention in recent years.

One proposed solution is to model the distribution of images using a projected mixture model [5]. This approach has proven successful, but is computationally very expensive and requires the number of states to be known in advance. Other methods rely on the fact that Fourier transforms of projections arising from the same volume coincide on a single line, known as the common line. Images can then be clustered by measuring correlations between lines in Fourier space [6, 7].

Often, molecular states will differ only locally, so estimation of Euler angles can be performed by fitting a single-state model. These are sufficiently accurate for a first reconstruction. Once images have been clustered according to different molecular states, angles are re-estimated during iterative refinement on each cluster. In the following, we assume that the angles are known and focus on the task of classification.

The work in this paper draws on a concept introduced by Penczek et al., who determine the volumes of different molecular states by estimating their covariance [8]. Given a set of C volumes, their covariance matrix will be dominated by $C - 1$ eigenvectors, or “eigen volumes.” Each projected image is then approximated as a combination of projected eigen volumes. Clustering the resulting coordinates into C classes now yields a good classification of the images. To estimate the covariance matrix, the authors propose a bootstrapping approach in which multiple subsets of the dataset are used to reconstruct multiple volumes whose covariance is then calculated. Unfortunately, this heuristic method offers no theoretical guarantees.

Katsevich et al. have proposed an estimator for the volume covariance matrix that remedies this problem [9]. This estimator has several useful properties: it converges to the population covariance matrix as the number of images goes to infinity, does not assume a particular distribution of molecular states, and does not require knowing the number of classes C . Indeed, C can be estimated from the spectrum of the covariance matrix.

Unfortunately, calculating this estimator involves the inversion of a high-dimensional linear operator, making direct calculation intractable for typical problems. To solve this, the authors replace the operator by a sparse, block-diagonal approximation that can be more easily inverted. However, this is only valid for a uniform distribution of viewing angles and does not incorporate the contrast transfer function (CTF) of the microscope, which is necessary for real-world data.

In this paper, we instead invert the original linear operator using the conjugate gradient (CG) method. The operator can be decomposed as a sum of sparse operators, and so applying it is computationally cheap. As a result, the CG inversion has an overall computational complexity of $O(nN_{\text{res}}^7)$, where n is the number of images and N_{res} is the effective resolution of the model. This approach also has the advantage of enabling a non-uniform

distribution of viewing angles and allows us to incorporate the effect of the CTF, as we demonstrate through classification on both simulated and experimental datasets.

2. CRYO-EM IMAGING MODEL

In this paper, we shall represent the molecular structure using its Coulomb potential function in three dimensions, defined by some $\mathcal{X} \in \mathbf{L}^1(\mathbb{R}^3)$. The imaging process includes a low-pass filtering, so we cannot expect to represent volumes accurately above a certain frequency. More importantly, since our goal is classification, we only need enough resolution to distinguish one molecular state from another. To reduce computational cost, we therefore restrict $\mathcal{X}$ to some finite-dimensional subspace $\mathcal{V}$ of $\mathbf{L}^1(\mathbb{R}^3)$ where the frequency content is concentrated in a ball of radius $N_{\text{res}}\pi/2$, yielding an effective resolution of N_{res} .

We can represent a particular viewing direction as an axis of integration and an in-plane rotation, that is an element R of $\text{SO}(3)$, the group of orientation-preserving rotations in $\mathbb{R}^3$. The projection of $\mathcal{X}$ corresponding to the rotation R is then given by

$$\mathcal{P}\mathcal{X}(x, y) = \int_{\mathbb{R}} \mathcal{X}(R^T r) dz, \quad (1)$$

where $r = (x, y, z)^T$.

An electron microscope never captures the actual projection PX . Instead, it registers a projection filtered by a CTF $H(\omega)$ which depends on microscope optics and the wavelength of the electron beam used [3].

Let us define the D -dimensional Fourier transform of a function $f \in \mathbf{L}^1(\mathbb{R}^D)$ (here D is typically 2 or 3) as

$$\hat{f}(\omega) = \int_{\mathbb{R}^D} f(x) e^{-i\omega^T x} dx \quad (2)$$

for any $\omega \in \mathbb{R}^D$. The Fourier transform of the CTF-filtered projection is then

$$H(\omega) \cdot \widehat{\mathcal{P}\mathcal{X}}(\omega). \quad (3)$$

Instead of applying the CTF to the filtered image, we can apply it to the volume prior to projection. The Fourier slice theorem [10] tells us that

$$\widehat{\mathcal{P}\mathcal{X}}(\omega_1, \omega_2) = \hat{\mathcal{X}}(R^T(\omega_1, \omega_2, 0)^T). \quad (4)$$

The CTF is radially symmetric, so it can be extended symmetrically to $\mathbb{R}^3$. We thus have $H(R^T \omega) = H(\omega)$, and so

$$H(\omega_1, \omega_2) \cdot \widehat{\mathcal{P}\mathcal{X}}(\omega_1, \omega_2) = (H \cdot \hat{\mathcal{X}})(R^T(\omega_1, \omega_2, 0)^T). \quad (5)$$

Letting $\mathcal{T}$ denote spatial filtering by $H(\omega)$, the CTF-filtered projection is then PTX .

CTFs have several zero-crossings, so reconstruction is not possible from a dataset obtained from a single CTF. Experiments thus use a number of microscope configurations, resulting in different CTFs covering the entire frequency spectrum.

Finally, the image is registered on a discrete grid of size N -by- N . As mentioned previously, however, we only consider volumes of effective resolution N_{res} . Consequently, we restrict our images to those in a finite-dimensional space $\mathcal{I}$ with frequency content centered in a ball of radius $N_{\text{res}}\pi/2$. The sampling operator mapping $L^1(\mathbb{R}^2)$ to $\mathcal{I}$ is denoted $\mathcal{S}$.

Putting everything together, the image $\mathcal{I}$ obtained from $\mathcal{X}$ through convolution with $\mathcal{T}$, projection by $\mathcal{P}$ and sampling by $\mathcal{S}$ is given by

$$\mathcal{I} = SPTX = MX, \quad (6)$$

where we have introduced the imaging operator $\mathcal{M} = SPT$.

Since both $\mathcal{V}$ and $\mathcal{I}$ are of finite dimension, we can represent them using finite bases. Let $\dim \mathcal{V} = p$ and $\dim \mathcal{I} = q$. We can then represent $\mathcal{X}$ and $\mathcal{I}$ as vectors X and I in $\mathbb{R}^p$ and $\mathbb{R}^q$, respectively. Similarly, $\mathcal{P}$ and $\mathcal{T}$ have matrix representations P and T , respectively, obtained by least-squares approximation. Taken together, we have the imaging matrix M . Note that $\mathcal{S}$ is no longer present since P and T already project onto a finite-dimensional space.

3. VOLUME COVARIANCE

3.1. Covariance estimator

To model the variability of volumes in the dataset, let $\mathbf{X}_s$ for $s = 1, \dots, n$ be a collection of independent and identically distributed discrete random variables in $\mathbb{R}^p$, each taking the value X_c with probability p_c for $c = 1, \dots, C$. These random variables have mean $\mu_0 = \mathbb{E}[\mathbf{X}_s]$ and covariance matrix

$$\Sigma_0 = \text{Var}[\mathbf{X}_s] = \mathbb{E}[(\mathbf{X}_s - \mathbb{E}[\mathbf{X}_s])(\mathbf{X}_s - \mathbb{E}[\mathbf{X}_s])^H], \quad (7)$$

where u^H is the conjugate transpose of the vector u . Since $\mathbf{X}_s$ is a discrete random variable with C states, Σ has rank $C - 1$.

To estimate μ_0 and Σ_0 , we consider the statistics of the projected images. Specifically, we define the random variables

$$\mathbf{I}_s = M_s \mathbf{X}_s + \mathbf{E}_s, \quad (8)$$

where $\mathbf{E}_s$ are independent and identically distributed zero-mean random noise vectors, independent of M_s and $\mathbf{X}_s$, with $\text{Var}[\mathbf{E}_s] = \sigma^2 I_q$. The expected value of $\mathbf{I}_s$ is

$$\mathbb{E}[\mathbf{I}_s] = M_s \mu_0, \quad (9)$$

while its covariance is given by

$$\text{Var}[\mathbf{I}_s] = M_s \Sigma_0 M_s^H + \sigma^2 I_q, \quad (10)$$

where M_s^H is the conjugate transpose of the imaging operator M_s and I_q is the q -by- q identity matrix.

Let us consider the realizations I_s of $\mathbf{I}_s$ for $s = 1, \dots, n$. Following (9) and (10), we define the following estimators for μ_0 and Σ_0 :

$$\mu_n = \arg \min_{\mu} \frac{1}{n} \sum_{s=1}^n \|I_s - M_s \mu\|^2, \quad (11)$$

$$\Sigma_n = \arg \min_{\Sigma} \frac{1}{n} \sum_{s=1}^n \left\| (I_s - M_s \mu_n) (I_s - M_s \mu_n)^H - (M_s \Sigma M_s^H + \sigma^2 I_q) \right\|_F^2, \quad (12)$$

where $\|\cdot\|_F$ is the Frobenius matrix norm.

Differentiating and setting to zero in (11), we get

$$A_n \mu_n = b_n, \quad (13)$$

where A_n and b_n are given by

$$A_n = \frac{1}{n} \sum_{s=1}^n M_s^H M_s, \quad b_n = \frac{1}{n} \sum_{s=1}^n M_s^H I_s \quad (14)$$

Applying the same process to (12), we obtain

$$L_n(\Sigma_n) = B_n, \quad (15)$$

where $L_n: \mathbb{C}^{p \times p} \rightarrow \mathbb{C}^{p \times p}$ is the linear operator defined by

$$L_n(\Sigma) = \frac{1}{n} \sum_{s=1}^n M_s^H M_s \Sigma M_s^H M_s \quad (16)$$

and

$$B_n = \frac{1}{n} \sum_{s=1}^n M_s^H (I_s - M_s \mu_n) (I_s - M_s \mu_n)^H M_s - \sigma^2 \frac{1}{n} \sum_{s=1}^n M_s^H M_s. \quad (17)$$

Calculating μ_n and Σ_n thus amounts solving (13) and (15). Since $\mathcal{I}$ contains images of effective resolution N_{res} , $q = O(N_{\text{res}}^2)$. Likewise, $p = O(N_{\text{res}}^3)$. The matrix A_n is p -by- p , can

therefore be naïvely inverted with a complexity $O(N_{\text{res}}^9)$. However, if we were to compute the matrix representation of L_n , this would be a p^2 -by- p^2 matrix, and its inversion would take $O(N_{\text{res}}^{18})$. So while we may be able to calculate A_n^{-1} , inverting L_n poses a much greater challenge.

3.2. Inversion of L_n

Since direct inversion of the matrix of L_n is not an option, we turn to other methods of solving (15). If L_n can be calculated fast, the conjugate gradient method provides a viable approach for estimating Σ_n .

By choosing appropriate spaces of $\mathcal{V}$ and $\mathcal{J}$ and equipping these with well-behaved bases, P_s can be expressed as a block-diagonal matrix consisting of $O(N_{\text{res}})$ blocks of size $O(N_{\text{res}})$ -by- $O(N_{\text{res}}^2)$. The application of the CTF, T , is also represented by a block-diagonal matrix with diagonal blocks in this basis. Details on the construction of these bases can be found in Katsevich et al. [9]. All matrix multiplications are therefore done in blocks, reducing computational complexity.

The CTF matrices T_s in L_n result in certain frequencies being amplified and others attenuated. Because the noise in our images I_s is white, the stability of the inversion ω thus depends on $|H_s(\omega)|$. We would therefore like to only invert L_n when $|H_s(\omega)|$ is large on average, corresponding to the dominant eigenvalues of L_n . Since the conjugate gradient method first constructs the larger eigenvalues of the inverse before moving on to smaller eigenvalues, the desired result is obtained by limiting the number of iterations, yielding an implicit regularization [11]. The overall complexity is therefore that of applying L_n , which can be shown to be $O(nN_{\text{res}}^7)$.

3.3. Classification

As mentioned previously, Σ_0 has $C - 1$ non-zero eigenvalues and the eigenvectors, together with μ_0 , define an affine space containing all the volumes. Due to noise, this is not the case for Σ_n , although it converges to Σ_0 as n increases. In numerical experiments, we find that for large n , Σ_n will contain $C - 1$ dominant eigenvalues and the associated eigenvectors approximate the eigenvectors of Σ_0 .

Assembling the dominant $C-1$ eigenvectors into a matrix U_n , we can associate with each image I_s a coordinate vector a_s such that $\|(\mu_n + U_n a_s) - I_s\|^2$ is minimized. If I_s is a projection of the volume X_c , $\mu_n + U_n a_s$ should be close to X_c . As a result, the a_s cluster according to molecular state.

This lets us classify the images according to their molecular structure. Applying a clustering algorithm to the a_s vectors, the images generated by the a given volume will be found in the same cluster. We use a Gaussian mixture model (GMM) trained using the expectation-maximization (EM) algorithm [12]. Once images are associated with a particular molecular structure, standard tomographic inversion techniques can be applied to recover that structure.

4. NUMERICAL EXPERIMENTS

4.1. Synthetic data

To evaluate the above method, we apply it to a synthetic dataset consisting of $n = 10000$ images generated from $C = 2$ volumes, projected along random viewing directions, and filtered by one of seven CTFs. Each image is sampled on a 65-by-65 grid and a Gaussian noise of variance σ^2 is added. In this section, we have a heterogeneous SNR $\text{SNR}_{\text{het}} = 0.005$ (for a discussion of SNR_{het} , see [9]). Sample projections are shown in Figure 1(a,b).

The algorithm is run with an effective resolution of $N_{\text{res}} = 17$ for 10 iterations. The total running time is 50 min on a 3 GHz quad-core CPU with 4 GB of memory. The eigenvalues of Σ_n are shown in Figure 2(a). One eigenvalue is separated from the rest by a spectral gap of 8.4, representing the heterogeneity in the dataset. Calculating the coordinates α_s of the images I_s , we obtain two well-separated distributions. Clustering the coordinates recovers the original classes with 99.8% accuracy.

4.2. Ribosome 70S complex

We also apply the method an experimental dataset consisting of $n = 10000$ images sampled on a 130-by-130 grid. Sample images are shown in Figure 1(c,d). The projections correspond to different molecular states of a 70S ribosomal complex from E. Coli generously provided by J. Frank's group at Columbia University [13]. To estimate orientations, the RELION software was run with one class [14].

Running the algorithm with $N_{\text{res}} = 17$ for 10 iterations, we obtain Σ_n with the spectrum shown in Figure 3(a). The total running time was 46 min. Here, one dominant eigenvalue at $\lambda_1 = 0.072$ is well-separated from the bulk of the spectrum by a gap of 1.75. We therefore conclude that the dataset contains two classes. Calculating the coordinate α_1 , we obtain the bimodal distribution in 3(b). We then cluster α_1 and send each class to RELION for reconstruction, which yields the cross-sections shown in Figure 3(c) and (d). The volumes are differentiated in the central cavity and in the rotation of the upper part. Compared to the labeling provided with the dataset, our clustering achieves an accuracy of 87%.

5. CONCLUSION

Replacing the sparse approximation of Katsevich et al. with an iterative approach, we obtain a more flexible method for covariance estimation capable of tackling experimental datasets. Compared to other algorithms, it is less computationally intensive and provides an estimate of the number of classes, simplifying the classification of molecular structure.

Acknowledgments

This research was partially supported by Award Number R01GM090200 from the NIGMS and Award Number LTR DTD 06-05-2012 from the Simons Foundation.

REFERENCES

1. Bai X, McMullan G, Scheres SHW. How Cryo-EM is revolutionizing structural biology. Trends in Biochemical Sciences. 2014 In Press.

2. Kühlbrandt W. The resolution revolution. *Science*. 2014; 343:1443–1444. [PubMed: 24675944]
3. Frank, J. *Three-Dimensional Electron Microscopy Of Macromolecular Assemblies: Visualization Of Biological Molecules In Their Native*. Oxford University Press; USA: 2006.
4. van Heel M, Gowen B, Matadeen R, Orlova E, Finn R, Pape T, Cohen D, Stark H, Schmidt R, Schatz M, et al. Single-particle electron cryo-microscopy: towards atomic resolution. *Quarterly reviews of biophysics*. 2000; 33(04):307–369. [PubMed: 11233408]
5. Sigworth F, Doerschuk P, Carazo J, Scheres S. Chapter Ten: An introduction to maximum-likelihood methods in Cryo-EM. *Methods in enzymology*. 2010; 482:263–294. [PubMed: 20888965]
6. Herman G, Kalinowski M. Classification of heterogeneous electron microscopic projections into homogeneous subsets. *Ultramicroscopy*. 2008; 108:327–338. [PubMed: 17574340]
7. Shatsky M, Hall R, Nogales E, Malik J, Brenner S. Automated multi-model reconstruction from single-particle electron microscopy data. *Journal of Structural Biology*. 2010; 170:98–108. [PubMed: 20085819]
8. Penczek P, Kimmel M, Spahn C. Identifying conformational states of macromolecules by eigen-analysis of resampled Cryo-EM images. *Structure*. 2011; 19(11):1582–1590. [PubMed: 22078558]
9. Katsevich E, Katsevich A, Singer A. Covariance matrix estimation for the cryo-em heterogeneity problem. *SIAM Journal on Imaging Sciences*. 2015; 8(1):126–185. [PubMed: 25699132]
10. Natterer, F. *The mathematics of computerized tomography*. Springer; 1986.
11. Hanke, M. *Conjugate gradient type methods for ill-posed problems*. Vol. 327. CRC Press; 1995.
12. Dempster A, Laird N, Rubin D. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1977:1–38.
13. Liao, H.; Frank, J. *Biomedical Imaging: From Nano to Macro*, 2010 IEEE International Symposium on. IEEE; 2010. Classification by bootstrapping in single particle methods; p. 169–172.
14. Scheres S. RELION: implementation of a bayesian approach to Cryo-EM structure determination. *Journal of structural biology*. 2012; 180(3):519–530. [PubMed: 23000701]

Fig. 1.
Sample projection images from the synthetic dataset (a,b) and experimental images of the 70S ribosome (c,d).

Fig. 2.

(a) The eigenvalue histogram of Σ_n obtained from synthesized data. (b) The histogram of the coordinate α_1 .

Fig. 3.

(a) The eigenvalue histogram for Σ_n obtained from experimental images of the 70S ribosome complex. (b) The histogram of the coordinate α_1 . (c,d) Cross-sections of estimated volumes.