Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2017 April ; 2017: 654–658. doi:10.1109/ISBI.2017.7950605.

MAHALANOBIS DISTANCE FOR CLASS AVERAGING OF CRYO-EM IMAGES

Tejal Bhamre¹, Zhizhen Zhao², and Amit Singer³

¹Dept. of Physics, Princeton University

²Dept. of Electrical Engineering, UIUC

³Dept. of Mathematics, Princeton University*

Abstract

Single particle reconstruction (SPR) from cryo-electron microscopy (EM) is a technique in which the 3D structure of a molecule needs to be determined from its contrast transfer function (CTF) affected, noisy 2D projection images taken at unknown viewing directions. One of the main challenges in cryo-EM is the typically low signal to noise ratio (SNR) of the acquired images. 2D classification of images, followed by class averaging, improves the SNR of the resulting averages, and is used for selecting particles from micrographs and for inspecting the particle images. We introduce a new affinity measure, akin to the Mahalanobis distance, to compare cryo-EM images belonging to different defocus groups. The new similarity measure is employed to detect similar images, thereby leading to an improved algorithm for class averaging. We evaluate the performance of the proposed class averaging procedure on synthetic datasets, obtaining state of the art classification.

Index Terms

Cryo-electron microscopy; single particle reconstruction; particle picking; class averaging; Mahalanobis distance; denoising; CTF

1. INTRODUCTION

SPR from cryo-EM is a rapidly advancing technique in structural biology to determine the 3D structures of macromolecular complexes in their native state [5, 7], without the need for crystallization. First, the sample, consisting of randomly oriented, nearly identical copies of a macromolecule, is frozen in a thin ice layer. An electron microscope is used to acquire top view images of the sample, in the form of a large image called a ‘micrograph’, from which individual particle images are picked semi-automatically. After preprocessing the selected raw particle images, the images are clustered. The images within each class are averaged to enjoy a higher SNR than the individual images. This step is known as “class averaging”. To minimize radiation damage, cryo-EM imaging must be constrained to low electron doses,

*This work was partially supported by Award Number R01GM090200 from the NIGMS, FA9550-12-1-0317 from AFOSR, Simons Foundation Investigator Award and Simons Collaborations on Algorithms and Geometry, and the Moore Foundation Data-Driven Discovery Investigator Award.

which results in a very low SNR in the acquired 2D projection images. Class averaging is thus a crucial step in the SPR pipeline: class averages are used for a preliminary inspection of the dataset, to eliminate outliers, and in semi-automated particle picking [16]. Typically, a user manually picks particles from a small number of micrographs. These are used to compute class averages, which are further used as templates to pick particles from all micrographs. Class averages are also used in subsequent stages of the SPR pipeline, such as 3D abinitio modeling.

The two popular approaches for 2D class averaging [12, 13, 15, 22, 10, 11] in cryo-EM are multivariate statistical analysis (MSA)[22] with multi-reference alignment (MRA) [3] and iterative reference-free alignment using K-means clustering [13]. Popular cryo-EM packages like RELION, XMIPP, EMAN2, SPIDER, SPARX, IMAGIC [23, 17, 21, 9, 16, 6] use some of these methods for class averaging. RELION uses a maximum likelihood classification procedure. A faster and more accurate approach for 2D class averaging based on rotationally invariant representation was introduced in [26] and is implemented in the cryo-EM software package ASPIRE (<http://spr.math.princeton.edu/>).

Recently in [1], it was shown that preliminary inspection of the underlying clean images and outlier detection can be performed at an earlier stage, by better denoising the acquired images using an algorithm called Covariance Wiener Filtering (CWF). In CWF, the covariance matrix of the underlying clean projection images is estimated from their noisy, CTF-affected measurements. The covariance is then used in the classical Wiener deconvolution framework to obtain denoised images.

There are two main contributions of this paper. First, we introduce a new similarity measure, related to the Mahalanobis distance [8], to compute the similarity of pairs of cryo-EM images. Second, we use the proposed Mahalanobis distance to improve the class averaging algorithm described in [26]. We first obtain for each image a list of S other images suspected as nearest neighbors using the rotation invariant representation (see Sec. 2 for details), and then rank these suspects using the Mahalanobis distance (see Fig. 1). The top K nearest neighbors, where $K < S$, given by this procedure are finally aligned and averaged to produce class averages. We test the new algorithm on a synthetic dataset at various noise levels and observe an improvement in the number of nearest neighbors correctly detected.

2. BACKGROUND

2.1. Image Formation Model

Under the linear, weak phase approximation (see [4, Chapter 2]), the image formation model in cryo-EM is given by

$$y_i = a_i \star x_i + n_i \quad (1)$$

where $\star$ denotes the convolution operation, y_i is the noisy projection image in real space, x_i is the underlying clean projection image in real space, a_i is the point spread function of the microscope, and n_i is additive Gaussian noise that corrupts the image. In the Fourier domain,

images are multiplied with the Fourier transform of the point spread function, called the CTF, and eqn.(1) can be rewritten as

$$Y_i = A_i X_i + N_i \quad (2)$$

where Y_i , X_i and N_i are the Fourier transforms of y_i , x_i and n_i respectively. The CTF is approximately given by (see [4, Chapter 3])

$$CTF(\hat{k}; \Delta \hat{z}^2) = e^{-B \hat{k}^2} \sin \left[-\pi \Delta \hat{z} \hat{k}^2 + \frac{\pi}{2} \hat{k}^4 \right] \quad (3)$$

where $\Delta \hat{z} = \frac{\Delta z}{[C_s \lambda]^{\frac{1}{2}}}$ is the “generalized defocus” and $\hat{k} = [C_s \lambda]^{\frac{1}{4}} k$ is the “generalized spatial frequency”, and B is the B-factor for the Gaussian envelope function. CTF’s corresponding to different defocus values have different zero crossings (see Fig. 2). Note that the CTF inverts the sign of the image’s Fourier coefficients when it is negative, and completely suppresses information at its zero crossings.

2.2. Rotationally Invariant Class Averaging

The procedure for class averaging, described in [26], was demonstrated to be both faster and more accurate than other existing class averaging procedures. It consists of three main steps. First, principal component analysis (PCA) of CTF-corrected phase flipped images is computed. We refer to this step as steerable PCA, because the procedure takes into account that the 2D covariance matrix commutes with in-plane rotations. Second, the bispectrum of the expansion coefficients in the reduced steerable basis is computed. The bispectrum is a rotationally invariant representation of images, but is typically of very high dimensionality. It is projected onto a lower dimensional subspace using a fast, randomized PCA algorithm [14]. One way to compare images after this step is using the normalized cross correlation. At low SNR, it is difficult to identify true nearest neighbors from the cross correlation. Therefore, Vector Diffusion Maps (VDM) [15] was used to further improve the initial classification by taking into account the consistency of alignment transformations among nearest neighbor suspects.

2.3. Covariance Wiener Filtering (CWF)

CWF was proposed in [1] as an algorithm to (i) estimate the CTF-corrected covariance matrix of the underlying clean images (since phase flipping is not an optimal correction) and (ii) using the estimated covariance to solve the associated de-convolution problem in eqn. 2 to obtain denoised images, that are estimates of X_i for each i in eqn. 2. The first step involves estimating the mean image of the dataset, μ , denoted $\hat{\mu}$, followed by solving a least squares problem to estimate the covariance Σ , denoted $\hat{\Sigma}$. Under the assumption of additive white Gaussian noise with variance σ^2 , the estimate of the underlying clean image X_i is given by

$$\hat{X}_i = (I - H_i A_i) \hat{\mu} + H_i Y_i \quad (4)$$

$$\text{where } H_i = \sum A_i^T (A_i \sum A_i^T + \sigma^2 I)^{-1}$$

3. ANISOTROPIC AFFINITY

The Mahalanobis distance in statistics [8] is a generalized, unitless and scale invariant similarity measure that takes correlations in the dataset into account. It is popularly used for anomaly detection and clustering [24, 25].

Our goal is to define a similarity measure to compare how close any two cryo-EM images are, given the CTF-affected, noisy observations for a pair of images from possibly different defocus groups, say Y_i and Y_j in eq. 2. CTF correction is a challenging problem due to the numerous zero crossings of the CTF. A popular, albeit, only partial correction of CTF is phase flipping, which involves simply inverting the sign of the Fourier coefficients. This corrects for the phase inversion caused by the CTF, but does not perform amplitude correction. Since phase flipping is suboptimal as a method for CTF correction, computing nearest neighbors using the Euclidean distance between features constructed from phase flipped, de-noised images can suffer from incorrectly identified neighbors. One simple approach would be to use the Euclidean distance between the CWF denoised images, as a measure of similarity. However, the optimality criterion for obtaining CWF denoised images is different from that for identifying nearest neighbors. Also, after CWF denoising, noise is no longer white, so the Euclidean distance is not an optimal measure of affinity.

In our statistical model, the underlying clean images $X_1, X_2, \dots, X_n \in \mathbb{C}^d$ (where n is the total number of images and d is the total number of pixels in each image) are assumed to be independent, identically distributed (i.i.d.) samples drawn from a Gaussian distribution. Further, we assume that the noise in our model is additive white Gaussian noise

$$X_i \sim \mathcal{N}(\mu, \Sigma) \quad N_i \sim \mathcal{N}(0, \sigma^2 I_d) \quad (5)$$

We note that while the assumption of a Gaussian distribution does not hold in practice, it facilitates the derivation of the new measure. The justification of the new measure is its superiority over the existing class averaging algorithm, as we demonstrate in Sec. 5.

The Gaussian assumption on signal and noise (5) and the image formation model (2) imply that Y_i is also Gaussian

$$Y_i \sim \mathcal{N}(A_i \mu, A_i \sum A_i^T + \sigma^2 I_d), \quad \text{for } i=1, \dots, n. \quad (6)$$

The joint distribution of (X_i, Y_i) is given by

$$\begin{bmatrix} X_i \\ Y_i \end{bmatrix} = \begin{bmatrix} I & 0 \\ A_i & I \end{bmatrix} \times \begin{bmatrix} X_i \\ N_i \end{bmatrix} \quad (7)$$

$$\sim \mathcal{N} \left[\begin{bmatrix} \mu \\ A_i \mu \end{bmatrix}, \begin{bmatrix} \sum & \sum A_i^T \\ A_i \sum & A_i \sum A_i^T + \sigma^2 I \end{bmatrix} \right] \quad (8)$$

The conditional distribution of X_i given Y_i is also Gaussian

$$X_i | Y_i = y_i \sim \mathcal{N}(\alpha_i, L_i) \quad (9)$$

where

$$\begin{aligned} \alpha_i &= \mu + \sum A_i^T (A_i \sum A_i^T + \sigma^2 I)^{-1} (y_i - A_i \mu) \\ L_i &= \sum - \sum A_i^T (A_i \sum A_i^T + \sigma^2 I)^{-1} A_i \sum. \end{aligned} \quad (10)$$

So

$$X_i - X_j | Y_i = y_i, Y_j = y_j \sim \mathcal{N}(\alpha_i - \alpha_j, L_i + L_j) \quad (11)$$

Let $X_i - X_j = x_{ij}$ and $\alpha_i - \alpha_j = \alpha_{ij}$. Then, for small ε , the probability that the ℓ_p distance between X_i and X_j is smaller than ε is

$$\begin{aligned} \Pr(\|X_{ij}\|_p < \varepsilon | Y_i = y_i, Y_j = y_j) &= \frac{1}{(2\pi)^{\frac{d}{2}} |L_i + L_j|^{\frac{1}{2}}} \times \int_{B_p(0, \varepsilon)} \exp\left\{-\frac{1}{2}(x_{ij} - \alpha_{ij})^T (L_i + L_j)^{-1} (x_{ij} - \alpha_{ij})\right\} dx_{ij} \\ & \quad (12) \end{aligned}$$

$$= \frac{\varepsilon^d \text{Vol}(B_p(0, 1))}{(2\pi)^{\frac{d}{2}} |L_i + L_j|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2} \alpha_{ij}^T (L_i + L_j)^{-1} \alpha_{ij}\right\} + \mathcal{O}(\varepsilon^{d+1}) \quad (13)$$

where $B_p(0, 1)$ is the ℓ_p ball of radius 1 in $\mathbb{R}^d$ centered at the origin (which follows from Taylor expanding the integrand). The probability of $\|X_{ij}\|_p < \epsilon$ given the noisy images y_i and y_j is a measure of the likelihood for the underlying clean images x_i and x_j to originate from the same viewing direction. We define our similarity measure after taking the logarithm on both sides of eqn. 13), dropping out the constants independent of i and j , and substituting back α_{ij} :

$$-\frac{1}{2}\log(|L_i+L_j|) - \frac{1}{2}(\alpha_i - \alpha_j)^T (L_i + L_j)^{-1} (\alpha_i - \alpha_j) \quad (14)$$

Notice the resemblance of the second term in eq. (14) to the classical Mahalanobis distance [8]. This term takes into account the anisotropic nature of the covariance matrix by appropriately normalizing each dimension when computing the distance between two points. Note that this distance is different for different pairs of points since it depends on $L_i + L_j$, unlike the Euclidean distance and the classical Mahalanobis distance.

4. ALGORITHM FOR IMPROVED CLASS AVERAGING USING MAHALANOBIS DISTANCE

We propose an improved class averaging algorithm that incorporates the affinity measure (14). The quantities α_i , L_i are computed for each image and defocus group respectively (in practice Σ is replaced by its estimate $\hat{\Sigma}$), using CWF [1]. The estimated covariance using CWF is block diagonal in the Fourier Bessel basis. In practice, we use α_i , L_i projected onto the subspace spanned by the principal components (for each angular frequency block). We obtain an initial list of S nearest neighbors for each image using the Initial Classification algorithm in [26]. Then, for the list of nearest neighbors corresponding to each image, the affinity (14) is computed and used to pick the closest K nearest neighbors, where $K < S$.

5. NUMERICAL EXPERIMENTS

We test the improved class averaging algorithm on a synthetic dataset that consists of projection images generated from the volume of *P. falciparum* 80S ribosome bound to E-tRNA, available on the Electron Microscopy Data Bank (EMDB) as EMDB 6454. The algorithm was implemented in the UNIX environment, on a machine with total RAM of 1.5 TB, running at 2.3 GHz, and with 60 cores. For the results described here, we used 10,000 projection images of size 65×65 that were affected by the various CTF's and additive white Gaussian noise at various noise levels, in particular, we show here results for 4 values of the SNR. The images were divided into 20 defocus groups. Initial classification was first used to select $S = 50$ nearest neighbors for each image. After rotationally aligning the suspected neighbors, the Mahalanobis distance was computed between each image and the 50 aligned suspects. We then pick the closest $K = 10$ neighbors for each image (in practice, the choice of K depends on the SNR and the number of images). For comparison, we compute 10 nearest neighbors for each image using only Initial Classification (with or without using the optional VDM step). Table 1 shows the number of pairs of nearest neighbor images detected with each method at various SNR's, that have correlation > 0.9 between the original clean

images, indicating that they are indeed neighbors. We note an improvement in the number of true nearest neighbors detected by the improved classification algorithm using the Mahalanobis distance. Figure 3 shows the estimated probability density function of the angular distance between nearest neighbor images, using 1) Initial Classification only 2) Improved classification using the Mahalanobis distance by repeating this experiment at four different SNR's. Figure 4a shows the results of Initial Classification and the improved class averaging algorithm on this synthetic dataset. We compare the quality of the class averages from [26] and this paper with $K = 10$. Averaging over a large number of nearest neighbors reduces the noise variance. However, it also blurs the underlying clean signal, since the neighbors are not exactly from the same viewing direction. Therefore, it is crucial to correctly identify only the top few nearest neighbors and average them in order to sufficiently reduce the noise without blurring the features too much. We see in Figure 4b that noise reduces in the class averages when K increases. The procedure in [26] took 168 seconds, while the improved classification using the anisotropic affinity took 860 seconds for the experiment described here.

6. DISCUSSION

We introduced a new similarity measure to compare CTF-affected cryo-EM images belonging to different defocus groups. The anisotropic affinity derived in this paper is similar to the one that appears in [18, 20] but also includes an additional logarithmic term. We provided a new probabilistic interpretation for this anisotropic affinity. The affinity can also be used as a similarity measure for any manifold learning procedure [20, 18] such as diffusion maps [19, 2], with or without missing data, and extended to other imaging modalities where images are affected by different point spread functions or blurring kernels.

References

1. Bhamre T, Zhang T, Singer A. Denoising and covariance estimation of single particle cryo-EM images. *Journal of Structural Biology*. 2016; 195(1):72–81. [PubMed: 27129418]
2. Coifman R, Lafon S. Diffusion maps. *Applied and Computational Harmonic Analysis*. 2006; 21(1): 5–30.
3. Dube P, Tavares P, Lurz R, van Heel M. Bacteriophage SPP1 portal protein: a DNA pump with 13-fold symmetry. *EMBO J*. 1993; 15:1303–1309.
4. Frank, J. Electron microscopy of macromolecular assemblies. In: Frank, J., editor. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies*. Academic Press; Burlington: 1996. p. 12-53.
5. Frank, J. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies : Visualization of Biological Molecules in Their Native State: Visualization of Biological Molecules in Their Native State*. Oxford University Press; USA: 2006.
6. Kimanius D, Forsberg BO, Scheres S, Lindahl E. Accelerated cryo-EM structure determination with parallelisation using GPUs in RELION-2. *bioRxiv*. 2016
7. Kühlbrandt W. The resolution revolution. *Science*. 2014; 343(6178):1443–1444. [PubMed: 24675944]
8. Mahalanobis, PC. On the generalised distance in statistics. *Proceedings National Institute of Science; India*. Apr. 1936; p. 49-55.
9. Marabini R, Masegosa IM, Martín MS, Marco S, Fernández J, de la Fraga L, Vaquerizo C, Carazo J. Xmipp: An image processing package for electron microscopy. *Journal of Structural Biology*. 1996; 116(1):237–240. [PubMed: 8812978]

10. Park W, Chirikjian G. An assembly automation approach to alignment of noncircular projections in electron microscopy. *IEEE Transactions on Automation Science and Engineering*. 2014; 11(3): 668–679.
11. Park W, Midgett CR, Madden DR, Chirikjian GS. A stochastic kinematic model of class averaging in single-particle electron microscopy. *I J Robotics Res*. 2011; 30:730–754.
12. Penczek PA, Radermacher M, Frank J. Three-dimensional reconstruction of single particles embedded in ice. *Ultramicroscopy*. 1992; 40:33–53. [PubMed: 1580010]
13. Penczek PA, Zhu J, Frank J. A common-lines based method for determining orientations for $N > 3$ particle projections simultaneously. *Ultramicroscopy*. 1996; 63(3–4):205–218. [PubMed: 8921628]
14. Rokhlin V, Szlam A, Tygert M. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications*. 2010; 31(3):1100–1124.
15. Schatz M, van Heel M. Invariant classification of molecular views in electron micrographs. *Ultramicroscopy*. 1990; 32:255–264. [PubMed: 2382314]
16. Scheres SHW. RELION: Implementation of a bayesian approach to cryo-em structure determination. *Journal of Structural Biology*. 2012; 180(3):519–530. [PubMed: 23000701]
17. Shaikh TR, Gao H, Baxter WT, FJA, Boisset N, Leith A, Frank J. SPIDER image processing for single-particle reconstruction of biological macromolecules from electron micrographs. *Nature Protocols*. 2008; 3(12):1941–1974. [PubMed: 19180078]
18. Singer A, Coifman RR. Non-linear independent component analysis with diffusion maps. *Applied and Computational Harmonic Analysis*. 2008; 25(2):226–239.
19. Singer A, Wu HT. Vector diffusion maps and the connection Laplacian. *Communications on Pure and Applied Mathematics*. 2012; 65(8):1067–1144.
20. Talmon R, Coifman RR. Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proceedings of the National Academy of Sciences*. 2013; 110(31):12535–12540.
21. Tang G, Peng L, Baldwin PR, Mann DS, Jiang W, Rees I, Ludtke SJ. EMAN2: An extensible image processing suite for electron microscopy. *Journal of Structural Biology*. 2007; 157:38–46. Software tools for macromolecular microscopy. [PubMed: 16859925]
22. van Heel M, Frank J. Use of multivariate statistics in analysing the images of biological macromolecules. *Ultramicroscopy*. 1981; 6(2):187–194. [PubMed: 7268930]
23. van Heel M, Harauz G, Orlova EV, Schmidt R, Schatz M. A new generation of the IMAGIC image processing system. *Journal of Structural Biology*. 1996; 116(1):17–24. [PubMed: 8742718]
24. Xiang S, Nie F, Zhang C. Learning a Mahalanobis distance metric for data clustering and classification. *Pattern Recognition*. 2008; 41(12):3600–3612.
25. Zhao X, Li Y, Zhao Q. Mahalanobis distance based on fuzzy clustering algorithm for image segmentation. *Digit Signal Process*. Aug; 2015 43(C):8–16.
26. Zhao Z, Singer A. Rotationally invariant image representation for viewing direction classification in cryo-EM. *Journal of Structural Biology*. 2014; 186(1):153–166. [PubMed: 24631969]

Fig. 1.
Pipeline of algorithm

Fig. 2.

CTF's for different values of the defocus. CTF parameters used are: the amplitude contrast $\alpha = 0.07$, the electron wavelength $\lambda = 2.51 \text{ pm}$, the spherical aberration constant $C_s = 2.0$, the B-factor $B = 10$, the defocus = $1 \mu\text{m}$, $1.3 \mu\text{m}$, and $1.6 \mu\text{m}$, and the pixel size is $2.82 \text{ \AA}$. See eq. (3)

Fig. 3.

The estimated probability density function of the angular distance (in degrees) between images classified into the same class by 1) Initial Classification and 2) Improved Classification using the anisotropic affinity at different SNR's.

Fig. 4.

Results of class averaging of a synthetic dataset of 10000 projection images of size 65×65 , affected by CTF and $\text{SNR} = 1 \Rightarrow 40$. (a) We show class averages with Initial Classification in the second row, and with the improved algorithm using the anisotropic affinity in the third row. We use $K = 10$ and $S = 50$. (b) Class averages for one image in the dataset with the improved algorithm using the anisotropic affinity, for $S = 50$, and using $K = 10, 20, 30$.

Table 1

Number of nearest neighbors with correlation > 0.9 , using 10,000 images, $K = 10$ and $S = 50$.

SNR	This work	[26] (VDM)	[26] (No VDM)
1/40	58373	56896	49560
1/60	34965	32113	29219
1/100	17262	14431	13706