Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes

Amit Moscovich^{1*}, Amit Halevi^{1*}, Joakim Andén² and Amit Singer^{1,3}

¹ Program in Applied & Computational Mathematics, Princeton University, Princeton, NJ

² Center for Computational Mathematics, Flatiron Institute, New York, NY

³ Department of Mathematics, Princeton University, Princeton, NJ

* Equal contribution.

E-mail: amit@moscovich.org, ahalevi@princeton.edu, janden@flatironinstitute.org and amits@math.princeton.edu

Abstract. Single-particle electron cryomicroscopy is an essential tool for high-resolution 3D reconstruction of proteins and other biological macromolecules. An important challenge in cryo-EM is the reconstruction of non-rigid molecules with parts that move and deform. Traditional reconstruction methods fail in these cases, resulting in smeared reconstructions of the moving parts. This poses a major obstacle for structural biologists, who need high-resolution reconstructions of entire macromolecules, moving parts included. To address this challenge, we present a new method for the reconstruction of macromolecules exhibiting continuous heterogeneity. The proposed method uses projection images from multiple viewing directions to construct a graph Laplacian through which the manifold of three-dimensional conformations is analyzed. The 3D molecular structures are then expanded in a basis of Laplacian eigenvectors, using a novel generalized tomographic reconstruction algorithm to compute the expansion coefficients. These coefficients, which we name *spectral volumes*, provide a high-resolution visualization of the molecular dynamics. We provide a theoretical analysis and evaluate the method empirically on several simulated data sets.

Keywords: single particle electron cryomicroscopy, heterogeneity, tomographic reconstruction, molecular conformation space, manifold learning, Laplacian eigenmaps, diffusion maps

Submitted to: Inverse Problems

1. Introduction

The function of biological macromolecules is determined not only by their chemical composition but also by their 3D configuration. Hence, accurately estimating these configurations is of great importance to the field of structural biology. Macromolecules may deform their structure, resulting in a continuum of possible configurations, known as *conformations*. It is not always possible to isolate different conformations and study each separately. As a result, practitioners often image a heterogeneous sample which is then "purified" computationally.

While X-ray crystallography and nuclear magnetic resonance (NMR) spectroscopy have been very successful in imaging homogeneous molecular structures to high resolution, they rely on aggregate measurements over an entire sample and are therefore ill-suited for imaging heterogeneous molecules. Single-particle electron cryomicroscopy (cryo-EM), on the other hand, produces a separate image for each individual molecule, opening the door to accurate determination of structural variability.

The cryo-EM process consists of rapidly freezing a solution containing the molecular sample and then imaging it using a transmission electron microscope. The electron detector captures a movie where each frame records the electron counts for every pixel. To reduce the effects of ionization damage—which destroys the fine structure of the molecules—the electron dose is kept low, resulting in exceptionally noisy images. See bottom row in Figure 1 for examples. Since each particle is randomly oriented with respect to the imaging plane, the resulting image contains projections of molecules from many random viewing directions. Each projection image is typically modeled as the line integral of the particle's electric potential along the viewing direction, followed by convolution with a point spread function and the addition of noise [1, 2]. The goal of cryo-EM reconstruction is to invert this process and obtain an estimate of the molecular volume from its projection images. In recent years, better sample preparation techniques and improved detectors have led to reconstructions at a near-atomic resolution [3, 4, 5, 6].

The standard assumption in 3D reconstruction by cryo-EM is that of a homogeneous sample with no structural variability. Many methods for homogeneous 3D reconstruction have been proposed [1, 7, 8, 9, 10]. The prevalent methods are based on a Bayesian approach [11] which starts from some initial guess for the volume and then performs expectation-maximization (EM), alternating between estimating a latent distribution of viewing directions for every image and estimating the volume given these distributions. As discussed above, however, the homogeneous assumption does not hold in general. Resolving molecular structures with variability is known as the *heterogeneity problem* in single-particle cryo-EM. Two types of heterogeneity are typically considered: discrete and continuous.

1.1. Discrete heterogeneity

Discrete heterogeneity models the structural variability of the particles using a small number of distinct molecular structures. This can be used to model heterogeneous



Figure 1. The potassium ion channel used to simulate a heterogeneous molecular ensemble. Note the C_4 rotational symmetry. (left) surface plot of the 3D density of a single molecule. We generated two synthetic datasets: ChannelSpin where the top red part is randomly rotated around the z axis (the molecule's axis of symmetry), and ChannelStretch where the bottom blue part is stretched along the x-y plane; (right) two different conformations from ChannelSpin projected along three orientations, from left to right: side view, top view, and oblique view. The top row contains clean projections whereas the bottom row contains corresponding CTF-filtered projections with noise added.

molecular structures that spend most of their time in one of several fixed states. One such molecule is ATP synthase, an enzyme that acts as a molecular stepper motor [12].

Several software packages support reconstruction with discrete heterogeneity, also known as 3D classification, including RELION [13], cryoSPARC [14], FREALIGN/cisTEM [15, 16], and EMAN2 [17]. To recover K distinct volumes, these methods assign, for each particle image, a latent distribution over the set $\{1, \ldots, K\}$. This is incorporated into the EM algorithm, similar to the latent distribution over the viewing directions.

1.2. Continuous heterogeneity

In this model, the molecular volumes in the sample vary continuously, but in a manner that is subject to many constraints due to molecular bonds. If the number of degrees of freedom associated with the flexible motion is small then the space of molecular volumes forms a low-dimensional manifold (up to thermal vibrations). Figure 1 shows a simple molecular model with continuous heterogeneity that we use in our simulations. Here, the continuous motion is the free rotation of the top part around the vertical axis. In this case, the manifold of molecular volumes is diffeomorphic to the unit circle S^1 .

One approach for analyzing structural heterogeneity is to perform principal component analysis (PCA) of the 3D molecular structures represented as densities on an $N \times N \times N$ voxel grid. This idea goes back to [18] and was further developed by [19, 20, 21, 22, 23]. These methods estimate the $N^3 \times N^3$ covariance matrix of the 3D volumes and compute its leading eigenvectors, known as *eigenvolumes*. One variant relies on a consistent least-squares estimator for the covariance [24, 25, 26]. These methods may capture continuous heterogeneity—as illustrated by [26]—but are ill-suited for high-resolution reconstruction, as we discuss in Section 3.3. A notable exception is

the method proposed in [27] that attempts to directly compute the leading eigenvectors, at high resolution, without estimating the entire covariance matrix.

A different approach is taken in [28, 29, 30, 31] and is based on diffusion maps, a non-linear dimensionality reduction method that is well-suited for recovering lowdimensional manifold structure [32, 33]. These methods start by clustering the projection images by their viewing direction and computing a low-dimensional embedding for each cluster. All of these different embeddings are then aligned, yielding a global embedding of conformations. Sets of close points in the global embedding may then be used to reconstruct a 3D volume corresponding to a particular conformation. This approach faces two important challenges: first, unsupervised global registration of the embeddings is by itself a very challenging problem [34, 35]; second, each individual embedding uses only a small subset of images from a particular viewing direction, which may be insufficient for accurate manifold recovery.

The RELION software package has also been recently extended to include multi-body refinement [36]. This method takes a segmentation of a 3D molecular reconstruction and attempts to refine each part separately from a static base model, with independent viewing directions and shift parameters for each part. Multi-body refinement, however, is limited to rigid variability and may fail to accurately reconstruct the interface between moving parts.

Other methods have been proposed based on normal mode analysis of the molecular structure reconstruction [37, 38]. However, the underlying harmonic oscillator model used in these methods may be too simple to describe sophisticated continuous variability such as structural deformations. See [39] for a survey of methods for studying continuous heterogeneity using cryo-EM.

1.3. Our contribution

We present a new method for recovering continuous variability based on manifold learning. In contrast to the viewing-direction specific manifold estimates of [28], our method directly approximates the global manifold of conformations from all projection images, regardless of their viewing direction.

Throughout this paper, we identify molecular volumes with their electric potential sampled on a 3D voxel grid of dimension N^3 . Under the continuous heterogeneity model, a single molecule corresponds to embedded submanifold of \mathbb{R}^{N^3} . This manifold is the range of a smooth function that maps a set of conformation parameters to a volume. A standard technique for approximating smooth functions on manifolds is by series expansion in Laplacian eigenfunctions, a technique that generalizes the familiar Fourier series expansion in Euclidean space. However, to apply this technique we need to have the Laplacian eigenfunctions. This is a "chicken and egg" problem: The computation of the Laplacian eigenfunctions requires the distribution of 3D volumes, which is the very thing we would like to estimate. To resolve this problem, we use the covariance-based approach [26] to obtain low-resolution estimates of the 3D volumes. These reconstructions are then



Figure 2. High-level diagram of our method, illustrated on the ChannelSpin dataset. (i) Classical single-particle reconstruction, to obtain estimates of the CTF and viewing directions. (ii) Covariance estimation of the 3D density. The eigenvectors of the covariance matrix are then used to form a low-resolution 3D reconstruction from each projection image. (iii) Using the low-resolution reconstructions we build an affinity graph and compute its Laplacian eigenvectors. (iv) We expand the unknown volumes in a basis of Laplacian eigenvectors and perform tomographic reconstruction. The result is r spectral volumes (left, overlaid on the mean image) which define a high-resolution reconstruction for each projection image (right).

used to form an empirical graph Laplacian whose r eigenvectors with lowest eigenvalues $\hat{\phi}^{(0)}, \ldots, \hat{\phi}^{(r-1)} \in \mathbb{R}^n$ are used in lieu of the unknown Laplacian eigenfunctions. Then we compute a set of expansion coefficient vectors $\hat{\alpha}^{(0)}, \ldots, \hat{\alpha}^{(r-1)} \in \mathbb{R}^{N^3}$, which we refer to as *spectral volumes*. Together, they define a high-resolution 3D reconstruction for each projection image:

$$\hat{\mathbf{x}}_s := \sum_{\ell=0}^{r-1} \hat{\boldsymbol{\alpha}}^{(\ell)} \hat{\phi}_s^{(\ell)},\tag{1}$$

where $\hat{\beta}_s \in \mathbb{R}^q$ is a vector that parameterizes the low-resolution estimate of the volume \mathbf{x}_s . To compute the expansion coefficients $\{\hat{\alpha}^{(\ell)}\}_{\ell=0}^{r-1}$, we formulate a novel generalized tomographic reconstruction problem posed as a 3D deconvolution, similar to [40]. The convolution kernel is computed efficiently using a non-uniform fast Fourier transform (NUFFT) [41, 42] and the solution is computed using the conjugate gradient method, leveraging the fast Fourier transform (FFT) for the application of the convolution. These computational details are key for scaling up to high-resolution. See Figure 2 for a diagram of the main steps which constitute our method.

Remark 1. The eigenvectors of the Laplacian can be used not only for function representation but also for non-linear dimensionality reduction (e.g. [43, 32]). In our case, they define an embedding of the low-resolution reconstructions that is useful for visualizing the underlying manifold of conformations. In Figure 5 we show the two-dimensional embedding of the ChannelSpin dataset using the second and third eigenvectors. Despite the high noise levels, the underlying circular manifold of motions is recovered.

Remark 2. The spectral volumes have the same dimensionality as the high-resolution volumes that we reconstruct. They may therefore be visualized as 3D molecular volumes, albeit with negative values as well as positive. These visualizations can provide insight regarding the range of motions of the molecule. See Figure 6 for examples and Section 5 for an asymptotic analysis of the spectral volumes.

Section 2 defines the forward model and formulates the inverse problem for continuous heterogeneity in cryo-EM. We describe our method in Section 3, including the generalized tomographic reconstruction from noisy projection images. Section 4 outlines the algorithms used and their computational complexity. In Section 5 we prove the convergence of the spectral volumes and high-resolution reconstructions under the manifold assumption. Finally, we present results on synthetic datasets in Section 6.

2. Problem formulation

We begin by describing the forward model for cryo-EM and then define the inverse problem that we wish to solve, first by considering the simpler case without heterogeneity and then by generalizing to the case of continuous heterogeneity. **Table 1.** List of symbols. Scalars are denoted by italics, vectors by boldface letters, matrices by non-italicized capitals, estimators are decorated with a hat.

\mathbf{Name}	Domain	Description
n	\mathbb{N}	Number of images and underlying molecular volumes
s	$1, \ldots, n$	Index to molecular image/volume
N	\mathbb{N}	Image/volume size
\check{N}	\mathbb{N}	Downsampled image/volume size
\mathbf{x}, \mathbf{x}_s	\mathbb{R}^{N^3}	Molecular volume
$\hat{\mathbf{x}}_s$	\mathbb{R}^{N^3}	Our high-resolution molecular volume estimate
u	$\{1,\ldots,N\}^3$	Voxel index
\mathbf{y}, \mathbf{y}_s	\mathbb{R}^{N^2}	Molecular image
$\mathbf{h},\mathbf{h}_{s}$	\mathbb{R}^{N^2}	Contrast transfer function (CTF)
R, R_s	SO(3)	3D viewing orientation
\mathbf{P},\mathbf{P}_s	$\mathbb{R}^{N^2 \times N^3}$	Imaging matrix (rotation, projection, and CTF)
\mathcal{F}_d		The d -dimensional discrete Fourier transform
μ	\mathbb{R}^{N^3} or $\mathbb{R}^{\check{N}^3}$	Mean volume (high-res or low-res)
Σ	$\mathbb{R}^{\check{N}^3 imes \check{N}^3}$	Covariance matrix of molecular volumes
q	\mathbb{N}	Number of PCA eigenvolumes
$\hat{\mathrm{V}}_q$	$\mathbb{R}^{\check{N}^3 imes q}$	Eigenvolumes of the estimated covariance matrix
$oldsymbol{eta}(\mathbf{x}), oldsymbol{eta}_s$	\mathbb{R}^{q}	PCA coordinates of a molecular volume
В	$\subseteq \mathbb{R}^q$	The domain of PCA coordinates
$ u(\mathbf{B})$		Measure of volumes in PCA coordinate representation
W	$\mathbb{R}^{n imes n}$	Edge weights matrix
L	$\mathbb{R}^{n \times n}$	Graph Laplacian matrix
\mathcal{M}	$\subset \mathbb{R}^{N^3}$	Riemannian submanifold of molecular volumes
$\phi^{(\ell)}$	$\mathbf{B} \to \mathbb{R}$	Laplace–Beltrami eigenfunction of the $\ell^{\rm th}$ smallest eigenvalue
$\hat{oldsymbol{\phi}}^{(\ell)}$	\mathbb{R}^n	Laplacian eigenvector of the ℓ^{th} smallest eigenvalue
r	\mathbb{N}	Number of spectral volumes
Κ	$\mathbb{R}^{rN^3 imes rN^3}$	Matrix of weighted projection-backprojections
b	\mathbb{R}^{rN^3}	Concatenation of weighted back-projection images
$oldsymbol{lpha}^{(\ell)}$	\mathbb{R}^{N^3}	Spectral volumes

2.1. Forward model

A sample of many identical molecules is prepared in a solution and then rapidly frozen, forming a thin sheet of vitreous ice which is then imaged using a transmission electron microscope (TEM). The resulting image is a measurement of the electrostatic potential of this thin sheet, integrated along the direction perpendicular to the imaging plane. The individual molecules, known as "particles" in the cryo-EM literature, are all captured in different orientations.

For every molecule in a particular 3D conformation, there is a corresponding realvalued electrostatic density map which we simply refer to as the *volume* and discretize it on an $N \times N \times N$ grid of voxels. We now describe the data generation model. First, the volumes $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are drawn i.i.d. from some distribution on \mathbb{R}^{N^3} which describes the structural variability of the molecule. Then, linear imaging operators $P_1, \ldots, P_n \in \mathbb{R}^{N^2 \times N^3}$ are drawn i.i.d from some distribution. These operators are the composition of a volume rotation operator R_s followed by a tomographic projection onto the imaging plane and finally convolved with a point spread function that is determined by the electron microscope configuration and position of the sample. The individual particle images $\mathbf{y}_1, \ldots, \mathbf{y}_n \in \mathbb{R}^{N^2}$ are formed by

$$\mathbf{y}_s = \mathbf{P}_s \mathbf{x}_s + \boldsymbol{\varepsilon}_s \quad \forall s = 1, 2, \dots, n, \tag{2}$$

where $\boldsymbol{\varepsilon}_s$ are noise terms. For simplicity, we assume that $\boldsymbol{\varepsilon}_s \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{N \times N})$ and ignore possible in-plane translations often used to account for inaccurate positioning of the particle image centers.

We consider the volumes $\mathbf{x}_s \in \mathbb{R}^{N^3}$ as functions $\mathbf{x}_s : M_N^3 \to \mathbb{R}$, where $M_N := [-1, -1 + 2/N, \dots, 1 - 2/N]$, is the grid for even values of N (a similar grid may be defined for odd N). Similarly, the images $\mathbf{y}_s \in \mathbb{R}^{N^2}$ are functions $\mathbf{y}_s : M_N^2 \to \mathbb{R}$.

To define the imaging operators P_s we must define the tomographic projection operation. One approach to this is in terms of line integrals perpendicular to the projection plane, but since the volumes lie on a discrete grid of measure zero, one must incorporate an interpolation scheme. An alternative is to express tomographic projection in the Fourier domain. Let **s** be a *d*-dimensional signal on M_N^d , its discrete Fourier transform (DFT) is given by

$$(\mathcal{F}_d \mathbf{s})(\mathbf{k}) := \sum_{\mathbf{u} \in M_N^d} e^{-2\pi i \langle \mathbf{k}, \mathbf{u} \rangle} \mathbf{s}[\mathbf{u}] \quad \forall \mathbf{k} \in \mathbb{R}^d$$
(3)

where **k** is a wave vector that corresponds to a particular directional frequency. By the Fourier slice theorem, a tomographic projection along the z axis in the spatial domain is equivalent to a restriction to the x-y plane in the Fourier domain [44]. We use this fact to express the projection image $P_s \mathbf{x}_s$ in the Fourier domain as follows:

$$(\mathcal{F}_{2}\mathbf{P}_{s}\mathbf{x}_{s})([k_{1},k_{2}]^{\mathrm{T}}) = (\mathcal{F}_{3}\mathbf{x}_{s})(R_{s}^{-1}[k_{1},k_{2},0]^{\mathrm{T}}) \cdot (\mathcal{F}_{2}\mathbf{h}_{s})([k_{1},k_{2}]^{\mathrm{T}}).$$
(4)

where $[k_1, k_2]$ is a wave vector in the resulting 2D projection image, $R_s \in \mathbb{R}^{3\times 3}$ is the rotation of particle number s and \mathbf{h}_s is the point-spread function whose Fourier transform $\mathcal{F}_2\mathbf{h}_s$ is known as the *contrast transfer function* (CTF). See Section 2 of [26] for more details on the forward model.

2.2. Inverse problem

Homogeneous case. The traditional inverse problem in single-particle cryo-EM assumes that all of the molecular volumes in the sample are identical. Thus, the forward model (2) simplifies to

$$\mathbf{y}_s = \mathbf{P}_s \boldsymbol{\mu} + \boldsymbol{\varepsilon}_s \quad \forall s = 1, 2, \dots, n, \tag{5}$$

where μ is a mean volume. Suppose the orientations and CTFs are known so that we have the imaging operators P_1, \ldots, P_n . Then for a white Gaussian noise model the

maximum-likelihood estimate of μ is given by the solution of the least-squares problem,

$$\hat{\boldsymbol{\mu}} = \underset{\boldsymbol{\mu} \in \mathbb{R}^{N^3}}{\operatorname{arg\,min}} \sum_{s=1}^n \|\mathbf{y}_s - \mathbf{P}_s \boldsymbol{\mu}\|^2.$$
(6)

This problem and regularized variants of it are not well-posed in general, with the condition number depending on the distribution of the viewing angles, the CTFs, and the desired resolution of the reconstruction. Nevertheless, high accuracy solutions are routinely obtained using cryo-EM software packages. [13, 14, 16, 17].

Continuous heterogeneity. Our main goal when analyzing a heterogeneous sample is to estimate the density of volumes $\mathbf{x} \in \mathbb{R}^{N^3}$ associated with a given molecule. We approach this problem by performing reconstructions of the individual volumes $\mathbf{x}_1, \ldots, \mathbf{x}_n$. Clearly, estimating nN^3 voxel values from merely nN^2 noisy measurements is an ill-posed problem and much harder than the homogeneous problem, where only a single volume of N^3 voxels needs to be estimated. In this paper we make two main assumptions: The first is that the molecular volumes in the sample lie near a low-dimensional manifold. This model is natural since many heterogeneous macromolecules only have a few degrees of freedom that describe their range of motions [28, 29, 30, 31]. Varying these degrees of freedom traces out a smooth, low-dimensional manifold $\mathcal{M} \subset \mathbb{R}^{N^3}$. The second assumption is that the imaging operators P_s can be accurately estimated using standard cryo-EM reconstruction tools. This is the case when the molecule contains a large fixed component and a smaller heterogeneous part. A good indication that this is indeed the case for a particular dataset is when the reconstruction of the mean volume has a high resolution in some regions and lower resolution in others.

In the next section, we explain how we combine these assumptions with spectral techniques for function approximation on low-dimensional spaces to reconstruct all of the volumes in a heterogeneous molecular sample.

3. Methods

In this section, we describe our spectral approach to the reconstruction of molecular samples with continuous heterogeneity. Our approach is based on the representation and approximation of molecular volumes using an orthogonal basis expansion of eigenfunctions. By expanding the molecular volumes in this basis and imposing the projection constraints we obtain a generalized spectral formulation of the cryo-EM reconstruction problem.

3.1. Manifold spectral representation

Our method builds on the output of a low-resolution reconstruction method [26] that we describe in Section 3.3. In this method, each reconstructed volume is a linear combination of q PCA eigenvolumes, hence it defines some mapping $(\mathbf{y}_s, \mathbf{P}_s) \mapsto \boldsymbol{\beta}_s$ where $\boldsymbol{\beta}_s \in \mathbb{R}^q$ is the vector of eigenvolume coefficients corresponding to a low-dimensional representation of \mathbf{x}_s . In what follows, we ignore potential ambiguities due to the projection and consider

the low-resolution reconstruction as a linear dimensionality reduction of the underlying volume $\mathbf{x}_s \mapsto \boldsymbol{\beta}_s$. Since we assumed the underlying manifold of volumes is *d*-dimensional, then if d < q the image of this mapping is some compact domain $\mathbf{B} \subseteq \mathbb{R}^q$ that is a *d*-dimensional immersed manifold.

In what follows we consider the approximation of smooth functions on general domains **B** via eigenfunctions of the Laplacian operator. We briefly review some relevant facts [45]. The Laplacian has a set of real eigenfunctions $\phi^{(\ell)} : \mathbf{B} \to \mathbb{R}$ that form a complete orthonormal basis of $L^2(\mathbf{B})$ with corresponding non-negative eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \ldots \rightarrow \infty$. The smoothness of $\phi^{(\ell)}$ is controlled by λ_ℓ , which corresponds to the spatial frequency of $\phi^{(\ell)}$. Consequently, the eigenfunctions with lowest eigenvalues form a natural basis for approximating smooth functions on **B**. In fact, this basis is optimal for the approximation of smooth functions for approximation and regression over arbitrary domains is a generalization of the classical approach for signal representation by Fourier series expansion [47].

Let us therefore consider the basis formed by the first r eigenfunctions $\phi^{(0)}, \ldots, \phi^{(r-1)}$. Fix a voxel $\mathbf{u} \in N^3$ and consider its associated restriction function $\mathbf{x}[\mathbf{u}]$. We may approximate this function using low-frequency eigenfunctions

$$\mathbf{x}[\mathbf{u}] \approx \sum_{\ell=0}^{r-1} \alpha_{\mathbf{u}}^{(\ell)} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x})), \tag{7}$$

where $\beta(\mathbf{x}) \in \mathbf{B}$ is the image of \mathbf{x} in PCA coordinates. This can be written more succinctly by aggregating the coefficients for all voxels into a single volume, yielding

$$\mathbf{x} \approx \sum_{\ell=0}^{r-1} \boldsymbol{\alpha}^{(\ell)} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x})), \quad \forall \boldsymbol{\beta} \in \mathbf{B}.$$
 (8)

We call the coefficient vectors $\boldsymbol{\alpha}^{(0)}, \ldots, \boldsymbol{\alpha}^{(r-1)} \in \mathbb{R}^{N^3}$ spectral volumes. Note that the above construction does not rely on a voxel-wise representation of the volumes as the same type of expansion can be done for volumes represented in any spatial basis.

The eigenfunctions are unknown, so we employ a widely used technique from the field of manifold learning, replacing them with estimates given by eigenvectors of a data-driven graph Laplacian. More specifically, we build a weighted undirected graph, where the vertices correspond to the projection images $\mathbf{y}_1, \ldots, \mathbf{y}_n$ and the edge weights are estimates of the affinity between the underlying molecular conformations. In our case, the affinities are computed from the low-resolution reconstruction coordinate $\hat{\beta}_s$ described in Section 3.3. We then form the symmetric normalized graph Laplacian and compute its r eigenvectors with the lowest eigenvalues,

$$\hat{\boldsymbol{\phi}}^{(0)},\ldots,\hat{\boldsymbol{\phi}}^{(r-1)}\in\mathbb{R}^n.$$

See Section 4.1 for the specific algorithms used for forming the graph and computing these eigenvectors. As we explain in Section 5.3, we may assume that these estimates converge to the eigenfunctions in the sense that

$$\hat{\phi}_s^{(\ell)} \approx \frac{1}{\sqrt{n}} \phi^{(\ell)}(\boldsymbol{\beta}_s) \quad \forall s = 1, 2, \dots, n,$$
(9)

where the \sqrt{n} factor is needed for proper normalization, so that

$$\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)} \right)^{2} = 1.$$
 (10)

We can now write a data-driven variant of the spectral expansion in (8),

$$\mathbf{x}_s \approx \sqrt{n} \sum_{\ell=0}^{r-1} \boldsymbol{\alpha}^{(\ell)} \hat{\phi}_s^{(\ell)} \quad \forall s = 1, 2, \dots, n.$$
(11)

In the next section we explain how we estimate the coefficients of this expansion.

3.2. Generalized tomographic reconstruction

We assume that the molecular orientations can be accurately estimated using standard methods for homogeneous cryo-EM reconstruction [22, 48], so that the projection operators P_s are estimated to high accuracy. By applying the imaging matrix P_s to both sides of (11) and plugging in the forward model (2), we obtain

$$\mathbf{y}_s \approx \sqrt{n} \sum_{\ell=0}^{r-1} \left(\mathbf{P}_s \boldsymbol{\alpha}^{(\ell)} \right) \hat{\phi}_s^{(\ell)} \quad \forall s = 1, 2, \dots, n.$$
(12)

We seek spectral volumes that minimize the squared error

$$\left(\hat{\boldsymbol{\alpha}}^{(0)},\ldots,\hat{\boldsymbol{\alpha}}^{(r-1)}\right) := \operatorname*{arg\,min}_{\left\{\boldsymbol{\alpha}^{(0)},\ldots,\boldsymbol{\alpha}^{(r-1)}\right\}} \sum_{s=1}^{n} \left\| \mathbf{y}_{s} - \sqrt{n} \sum_{\ell=0}^{r-1} \left(\mathbf{P}_{s} \boldsymbol{\alpha}^{(\ell)}\right) \hat{\phi}_{s}^{(\ell)} \right\|^{2}.$$
(13)

The minimizer can be calculated efficiently by forming the normal equations and solving them using the conjugate gradient method. Using this method also makes it easy to add regularization terms. See Section 4.2 for more details on the numerical solution of this minimization problem. Note that in contrast to the low-resolution PCA eigenvolumes, the spectral volumes are at the full resolution L.

Our high-resolution reconstructions of the molecular volumes are now given by

$$\hat{\mathbf{x}}_s = \sqrt{n} \sum_{\ell=0}^{r-1} \hat{\phi}_s^{(\ell)} \hat{\boldsymbol{\alpha}}^{(\ell)} \quad \forall s = 1, 2, \dots, n.$$
(14)

This estimator generalizes the least-squares estimator (6) for a single mean volume to multiple volumes $\hat{\alpha}^{(0)}, \ldots, \hat{\alpha}^{(r-1)}$ whose contribution to the reconstructed volumes is determined by the weights $\hat{\phi}^{(0)}, \ldots, \hat{\phi}^{(r-1)}$.

3.3. Low-resolution reconstruction

While the approach outlined above provides a recipe for computing the eigenvectors $\hat{\phi}^{(0)}, \ldots, \hat{\phi}^{(r-1)}$ and using them to obtain high-resolution volume estimates, a crucial ingredient is missing still: the graph weights W_{ij} . We would like them to approximate an affinity of the underlying molecular volumes.

Several approaches have been proposed for computing affinities between projection images of heterogeneous ensembles. One of the earliest was to compute affinities using a common-line distance [49]. This is done without estimating the relative orientations and so this procedure finds the best common-line correspondence out of all candidate common lines, resulting in very noisy affinity estimates. To reduce the noise one can first estimate the orientations of the projection images and then compute the common line distance based on the relative orientation. This was proposed in [50], however, the resulting affinity measure is still very noisy, so the authors first performed 2D class averaging within each set of projection images from the same orientation. However, this may average different conformations together.

We define the affinity W_{ij} using the L^2 distance between low-resolution reconstructions, obtained using the covariance estimation method [26]. This approach achieves robustness to noise without averaging different conformations together. We briefly describe their method now.

The first step is to estimate the mean $\mu = \mathbb{E}[\mathbf{x}]$ of the distribution of molecular volumes. This is done by taking the derivative of Equation (6) with respect to μ and setting it equal to zero. This yields the normal equations

$$\frac{1}{n} \left(\sum_{s=1}^{n} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s} \right) \hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{s=1}^{n} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{y}_{s}.$$
(15)

This formulation corresponds to the maximum-likelihood estimator of $\mathbb{E}[\mathbf{x}]$ in the setting of Gaussian white noise. As a consequence, $\hat{\boldsymbol{\mu}}$ is a consistent estimator [24]. A similar estimator for the covariance matrix $\text{Cov}[\mathbf{x}] := \mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^{\text{T}}]$ is given by

$$\hat{\Sigma} = \underset{\Sigma \in \mathbb{R}^{N^3 \times N^3}}{\operatorname{arg\,min}} \sum_{s=1}^{n} \left\| \left(\mathbf{P}_s \Sigma \mathbf{P}_s^{\mathrm{T}} + \sigma^2 \mathbf{I}_{N^2} \right) - (\mathbf{y}_s - \mathbf{P}_s \hat{\boldsymbol{\mu}}) (\mathbf{y}_s - \mathbf{P}_s \hat{\boldsymbol{\mu}})^{\mathrm{T}} \right\|_{\mathrm{F}}^2.$$
(16)

While not a maximum-likelihood estimator, it is consistent under mild conditions [24]. Computing its normal equations yields a linear system in $O(N^6)$ variables. Fortunately, this linear system can be reformulated as a deconvolution problem in six dimensions. Precalculating the convolution kernel requires $O(N^6 \log N + nN^4)$ operations, but it can then be applied with complexity $O(N^6 \log N)$. The equations can now be solved using the preconditioned conjugate gradient method. Empirically, it takes around 50 iterations to converge [26].

While more efficient than a naive approach, the algorithm outlined above still scales poorly in image size N. As a result, this covariance estimation method is not currently practical for N > 25. Furthermore, from a simple dimensionality argument, to estimate the $O(N^6)$ elements of $Cov[\mathbf{x}]$ from n images of size $N \times N$, we need at least $n = O(N^6/N^2) = O(N^4)$ images. So to apply the algorithm to experimental data, we must first downsample the images from $N \times N$ to $\check{N} \times \check{N}$. It is possible to gain insight on the structural variability using this approach, but the resulting reconstructions are of low-resolution.

After obtaining the mean and covariance estimates $\hat{\mu}$ and $\hat{\Sigma}$, the volumes $\mathbf{x}_1, \ldots, \mathbf{x}_n$ can be reconstructed by the PCA method introduced in [22]. First, the *q* eigenvectors, or eigenvolumes, of $\hat{\Sigma}$ are extracted and arranged as columns in a $\check{N}^3 \times q$ matrix \hat{V}_q . They represent the principal directions of molecular volume variability in $\mathbb{R}^{\check{N}^3}$. Together with the estimated mean, they define an affine q-dimensional subspace of $\mathbb{R}^{\check{N}^3}$ of the form $\hat{\mu} + \hat{V}_q \beta$, where $\beta \in \mathbf{B} \subseteq \mathbb{R}^q$ is a coordinates vector. Each image \mathbf{y}_s may then be associated with a volume in the affine subspace through [26]

$$\hat{\boldsymbol{\beta}}_{s} := \operatorname*{arg\,min}_{\boldsymbol{\beta} \in \mathbb{R}^{q}} \frac{1}{\sigma^{2}} \left\| \mathbf{y}_{s} - \mathbf{P}_{s} \left(\hat{\boldsymbol{\mu}} + \hat{\mathbf{V}}_{q} \boldsymbol{\beta} \right) \right\|^{2} + \left\| \boldsymbol{\Lambda}_{q}^{-1/2} \boldsymbol{\beta} \right\|^{2}, \tag{17}$$

where $\Lambda_q = \hat{V}_q^T \hat{\Sigma} \hat{V}_q$ is the diagonal matrix of the leading q eigenvalues of $\hat{\Sigma}$. The above estimator is the maximum a posteriori (MAP) estimator of the coordinates of \mathbf{x}_s for Gaussian distributions of \mathbf{x}_s and $\boldsymbol{\varepsilon}_s$. It is also equal to the Wiener filter estimator and the linear minimum mean squared error estimator of the coordinates [51, 52].

Given the solutions to (17), we have a low-resolution estimate of each volume \mathbf{x}_s given by $\hat{\boldsymbol{\mu}} + \hat{V}_q \hat{\boldsymbol{\beta}}_s$. We assume that the manifold structure of \mathcal{M} is not destroyed by the mapping of projection images to coordinate vectors in \mathbb{R}^q , hence that it is possible to invert this process and associate a unique molecular conformation with every low-dimensional reconstruction. If the intrinsic dimensionality of the conformation space is low and the volumes vary smoothly along this space then the inverse map $\mathbf{B} \to \mathbb{R}^{N^3}$ can be approximated by a small number of spectral volumes.

4. Algorithms and computational complexity

In this section, we provide the technical details of our reconstruction method. In Section 4.1 we describe the precise methods used to form the graph Laplacian and compute its eigenvectors, and in Section 4.2 we describe the deconvolution-based solution of the generalized tomographic reconstruction problem (13).

4.1. Graph computations

To compute the PCA eigenvolumes, we begin by downsampling the input images to size $\check{N} \times \check{N}$, where \check{N} is typically about 16. These images are then fed into the mean and covariance estimation pipeline described in [26]. It has computational complexity $O(n\check{N}^4 + \sqrt{\kappa'}\check{N}^6\log\check{N})$. The condition number κ' is of the order of 100. The top q eigenvectors of the estimated covariance $\hat{\Sigma}$ are computed and the q-dimensional coordinates $\hat{\beta}_s$ of each image are obtained via (17). This step has computational complexity $O(q\check{N}^3\log\check{N} + nq^2\check{N}^2)$, following the algorithm described in [26]. A weighted undirected graph is then constructed with vertices that correspond to the images $\mathbf{y}_1, \ldots, \mathbf{y}_n$ and edge weights calculated from the PCA coordinates $\hat{\boldsymbol{\beta}}_1, \ldots, \hat{\boldsymbol{\beta}}_n$. We tested two kinds of weight matrices:

- (i) Gaussian kernel weights $W_{ij} = e^{-\|\hat{\beta}_i \hat{\beta}_j\|^2 / 2\sigma^2}$.
- (ii) Binary symmetric KNN matrices, whereby $W_{ij} = 1$ if and only if $\hat{\beta}_i$ is one of the k nearest neighbors of $\hat{\beta}_i$ or vice versa, and $W_{ij} = 0$ otherwise.

In our preliminary experiments we obtained similar results with both choices. For our final results, we chose to use the symmetric KNN graph since it is sparse, which reduces the memory and computational costs. For the Laplacian matrix, we use the symmetric normalized graph Laplacian

$$L := D^{-1/2} (D - W) D^{-1/2} = I - D^{-1/2} W D^{-1/2},$$
(18)

where D is a diagonal matrix that satisfies $D_{ii} = \sum_{j} W_{ij}$. The symmetry of L permits the use of specialized algorithms for eigenvector calculation and guarantees that the resulting eigenvectors are orthogonal. See the tutorial by [53] for other common choices of weight and Laplacian matrices.

We build the KNN weights matrix W using MATLAB's knnsearch function which for low dimensions is based on a KDTree [54]. The running time of this part is $O(qn \log n)$ where q is the dimension of the PCA coordinates $\hat{\beta}_s$ used in the low-resolution reconstruction. We then form the Laplacian matrix L and compute its r eigenvectors $\hat{\phi}^{(0)}, \ldots, \hat{\phi}^{(r-1)}$ with lowest eigenvalues using MATLAB's eigs function. This function implements the Krylov–Schur algorithm [55]. The matrices W and L are stored as sparse matrices of average degree O(k), hence their memory usage is O(nk). There exist newer methods of computing eigenvectors, such as the algebraic multigrid preconditioner used by the megaman manifold learning package [56, 57]. We did not incorporate such methods in the current work, as the eigenvector calculation step was not a bottleneck in our implementation.

4.2. Spectral volume estimation

Recall that the spectral volumes are defined in (13) as minimizers of the generalized tomographic reconstruction equation, (13). To find this minimum, we compute the gradient with respect to $\{\boldsymbol{\alpha}^{(\ell)}\}_{\ell=0}^{r-1}$ and set it to zero, obtaining the normal equations

$$\frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{y}_{s} = \sum_{m=0}^{r-1} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \hat{\phi}_{s}^{(m)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s} \boldsymbol{\alpha}^{(m)} \quad \forall \ell = 0, 1, \dots, r-1.$$
(19)

We can rewrite the equation in vector notation by defining the vectors $\mathbf{b}^{(0)}, \ldots, \mathbf{b}^{(r-1)} \in \mathbb{R}^{N^3}$ to be weighted backprojected images

$$\mathbf{b}^{(\ell)} = \frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{y}_{s} , \qquad (20)$$

and $K \in \mathbb{R}^{rN^3 \times rN^3}$ to be an $r \times r$ block matrix, with blocks of size $L^3 \times L^3$. Each block is a weighted sum of projection-backprojection matrices, with its (ℓ, m) block given by

$$\mathbf{K}^{(\ell,m)} = \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \hat{\phi}_{s}^{(m)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s}.$$
(21)

By defining the vector $\mathbf{b} \in \mathbb{R}^{rN^3}$ to be the concatenation of $\mathbf{b}^{(0)}, \ldots, \mathbf{b}^{(r-1)}$ and $\boldsymbol{\alpha} \in \mathbb{R}^{rN^3}$ to be the concatenation of $\boldsymbol{\alpha}^{(0)}, \ldots, \boldsymbol{\alpha}^{(r-1)}$ we can rewrite (19) as

$$\mathbf{b} = \mathbf{K}\boldsymbol{\alpha}.\tag{22}$$

Since K is of size $N^3r \times N^3r$, it would be very expensive to directly solve this equation using standard direct inversion algorithms such as those based on LU or Cholesky decomposition, since this would require $O(N^9r^3)$ operations. Instead, we draw on the work of [40] and note that applying $P_s^T P_s$ to a volume is equivalent to convolving the volume with a kernel calculated from R_s and \mathbf{h}_s . The backprojected images $P_s^T \mathbf{y}_s$ may be calculated from R_s , \mathbf{h}_s , and \mathbf{y}_s through a similar method. A complication arises from the fact that the points $R_s^{-1}[k_1, k_2, 0]^T$ in (4) do not lie on a regular grid, hence to evaluate the expression $(\mathcal{F}_3\mathbf{x}_s)(R_s^{-1}[k_1, k_2, 0]^T)$ we need to compute Fourier amplitudes on a non-regular grid which cannot be achieved through the standard FFT. Instead, we use the FINUFFT non-uniform fast Fourier transform software package [58]. It has computational complexity $O(N^3 \log N + S)$ where S is the number of points at which the transform is computed. Here, $S = N^2 n$, as both \mathbf{h}_s and \mathbf{y}_s are of size $N \times N$, and we consider n instances of projection images. We must calculate $K^{(\ell,m)}$ for each of the r^2 (ℓ, m) -pairs, and $\mathbf{b}^{(\ell)}$ for each ℓ . Thus, the total time to calculate K is $O(r^2N^3 \log N + r^2nN^2)$ and to calculate \mathbf{b} is $O(rN^3 \log N + rnN^2)$.

Each step of the conjugate gradient method involves applying the forward operator as well as performing several vector dot products and additions. Applying the forward operator is done using r^2 FFT operations of size $N \times N \times N$, which has total complexity $O(r^2N^3 \log N)$. The complexity of the conjugate gradient method is thus $O(\sqrt{\kappa}r^2N^3 \log N)$, where κ is the condition number of K, since the conjugate gradient method converges in $O(\sqrt{\kappa})$ steps [59, 60]. The total complexity, including calculating K and **b** and solving the normal equations, is therefore $O(r^2nN^2 + \sqrt{\kappa}r^2N^3 \log N)$. For our synthetic data sets **ChannelSpin** and **ChannelStretch** with a uniform distribution of viewing directions and $n \gg N$ and using r = 16 spectral volumes, we have found that κ is of the order of 10–30. See Section 6.3 for empirical runtimes on these data sets.

Remark 3. The running time may be reduced by computing an approximation to K. In the proof of Theorem 1 we show that $K^{(\ell,m)} \to \delta_{\ell,m} \mathbb{E}[P^T P]$ in probability. We can thus approximate K by setting the off-diagonal blocks to zero and setting the diagonal blocks to the empirical estimate of $\mathbb{E}[P^T P]$

$$\mathbf{K}^{(\ell,\ell)} = \frac{1}{n} \sum_{s=1}^{n} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s}.$$
(23)

With this approximation, the time to approximate K is $O(N^3 \log N + nN^2)$, it is now dominated by the computation of **b**. The time to multiply vectors by K is $O(rN^3 \log N)$, so the total runtime drops by a factor of r to $O(rnN^2 + \sqrt{\kappa}rN^3 \log N)$.

5. Theory

In this section, we analyze the solution to the generalized tomographic reconstruction as defined in (13), starting with a simplified special case.

5.1. Warmup: Spectral volumes without projections

We first analyze the spectral volumes solution in an easy setting where the imaging operators P_1, \ldots, P_n are all equal to the identity matrix. That is, we have direct, albeit

noisy, measurements $\mathbf{z}_s = \mathbf{x}_s + \boldsymbol{\varepsilon}_s$ without projections and point spread function. This case is directly applicable for reconstructing a manifold of 2D images, as we demonstrate in Section 6.1. In this setting, the spectral volumes $\hat{\boldsymbol{\alpha}}^{(0)}, \ldots, \hat{\boldsymbol{\alpha}}^{(r-1)}$ minimize

$$\sum_{s=1}^{n} \left\| \mathbf{z}_{s} - \sqrt{n} \sum_{\ell=0}^{r-1} \hat{\phi}_{s}^{(\ell)} \boldsymbol{\alpha}^{(\ell)} \right\|^{2}.$$
(24)

In this sum, each voxel **u** can be considered separately, giving

$$\hat{\alpha}^{(\ell)}[\mathbf{u}] = \underset{\alpha^{(\ell)}[\mathbf{u}]}{\operatorname{arg\,min}} \sum_{s=1}^{n} \left| z_s[\mathbf{u}] - \sqrt{n} \sum_{\ell=0}^{r-1} \hat{\phi}_s^{(\ell)} \alpha^{(\ell)}[\mathbf{u}] \right|^2.$$
(25)

For a symmetric graph Laplacian L, the eigenvectors $\hat{\phi}_0, \ldots, \hat{\phi}_{r-1}$ form an orthonormal set. Hence the coefficient $\hat{\alpha}^{(\ell)}[\mathbf{u}]$ is given by an orthonormal projection of $z[\mathbf{u}]$ onto $\hat{\phi}^{(\ell)}$

$$\hat{\alpha}^{(\ell)}[\mathbf{u}] = \frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} z_{s}[\mathbf{u}] = \frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} (x_{s}[\mathbf{u}] + \varepsilon_{s}[\mathbf{u}]),$$
(26)

or, in vector form,

$$\hat{\boldsymbol{\alpha}}^{(\ell)} = \frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \left(\mathbf{x}_{s} + \boldsymbol{\varepsilon}_{s} \right) = \frac{1}{\sqrt{n}} \sum_{s=1}^{n} \hat{\phi}_{s}^{(\ell)} \mathbf{x}_{s} + \mathcal{N}\left(0, \frac{\sigma^{2}}{n} \mathbf{I}_{N^{2}} \right).$$
(27)

The last equality stems from the fact that the noise terms satisfy $\varepsilon_s \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{N^2})$. Consequently, the spectral volumes in this simplified model are, up to a noise term, orthogonal projections of the true volumes $\mathbf{x}_1, \ldots, \mathbf{x}_n$ onto the basis of Laplacian eigenvectors. In the next subsection, we show that this is also the case when tomographic projections are incorporated into the model.

5.2. Spectral volumes with projections

We now consider the full forward model with non-trivial imaging operators P_1, \ldots, P_n . First note that in our model, the images $\mathbf{y}_1, \ldots, \mathbf{y}_n$ and the imaging operators are random vectors, therefore the Laplacian eigenvectors $\hat{\phi}^{(0)} \ldots \hat{\phi}^{(r-1)} \in \mathbb{R}^n$ are also random vectors. For our analysis, we make the following two assumptions:

Assumption 1. Let \mathbf{y}_s be an image, drawn according to the forward model (2). Then its Laplacian eigenvector coordinates $\hat{\phi}_s^{(0)}, \ldots, \hat{\phi}_s^{(r-1)}$ are independent of \mathbf{P}_s .

In other words, the Laplacian eigenmap (or diffusion map) coordinates are independent of the viewing direction and CTF of the particle. We can justify this assumption by assuming that the covariance-based method of [26] performs accurate low-resolution reconstruction, regardless of the viewing angle.

Assumption 2. For any r > 0, the following sum converges in probability:

$$\max_{\ell \in \{0,\dots,r-1\}} \sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)} \right)^{4} \to 0.$$
(28)

That is, for any $\epsilon > 0$, the probability that $\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)} \right)^{4} > \epsilon$ tends to zero as $n \to \infty$.

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 17

Remark 4. From the normalization constraint $\sum_{s=1}^{n} (\hat{\phi}_{s}^{(\ell)})^{2} = 1$, unless the energy of $\hat{\phi}_{s}^{(\ell)}$ is highly concentrated, we expect to have $\hat{\phi}_{s}^{(\ell)} = O(1/\sqrt{n})$ and thus $\sum_{s} (\hat{\phi}_{s}^{(\ell)})^{4} = O(1/n)$. In the next subsection, we discuss convergence of the Laplacian eigenvectors to limiting eigenfunctions. If this spectral convergence holds then (28) follows.

Before stating our main result, we recall big-O in probability notation for stochastic boundedness: a sequence of random variables $\{X_n\}_{n=1}^{\infty}$ satisfies $X_n = O_P(f(n))$ if for every $\epsilon > 0$, there is some bound M_{ϵ} such that $\Pr[|X_n|/f(n) > M_{\epsilon}] < \epsilon$. We now state our main result which characterizes the estimated spectral volumes $\hat{\boldsymbol{\alpha}}^{(0)}, \ldots, \hat{\boldsymbol{\alpha}}^{(r-1)}$ up to a stochastically bounded error.

Theorem 1. (Spectral volume convergence) Let $\hat{\phi}^{(\ell)}$ be an eigenvector of the symmetric graph Laplacian described in Section 4.1. Under Assumptions 1 and 2, the spectral volumes as defined in (13) satisfy

$$\hat{\boldsymbol{\alpha}}^{(\ell)} = \mathbb{E}\left[\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{x}_{s}\right] + O_{\mathrm{P}}\left(\frac{1}{\sqrt{n}}\right),\tag{29}$$

where the expectation is taken with respect to the random draw of projection images as described in Section 2.1.

The proof is in Appendix A.

5.3. Convergence of the reconstructed volumes

Consider the graph Laplacian eigenvectors $\{\hat{\phi}^{(\ell)}\}_{\ell=0}^{n-1}$ computed from the low-resolution reconstruction coordinates $\hat{\beta}_1, \ldots, \hat{\beta}_n \in \mathbf{B}$. Several variants of the discrete graph Laplacian are known to converge to a continuous linear operator on **B**. This convergence is not only pointwise, but also spectral, meaning that the eigenvectors of the graph Laplacian converge to the eigenfunctions $\phi^{(\ell)}$ of this operator [61, 62, 63]. In particular cases, the limiting operator is the continuous Laplacian, but more generally it is a weighted Laplacian operator, or Fokker–Planck operator, which has an additional drift term towards, or away from, high-density regions [33, 64, 65]. For our theoretical analysis we only need spectral convergence towards some set of eigenfunctions, not necessarily the Laplacian eigenfunctions. We formulate this requirement in the following assumption. For simplicity, we ignore possible eigenvalue multiplicities.

Assumption 3. The eigenvectors $\hat{\phi}^{(0)}, \hat{\phi}^{(1)}, \ldots \in \mathbb{R}^n$ of the graph Laplacian, ordered by their eigenvalues, converge to a set of functions $\phi^{(0)}, \phi^{(1)}, \ldots : \mathbf{B} \to \mathbb{R}$, in the sense that

$$\sup_{s=1,\dots,n} |\sqrt{n}\hat{\phi}_s^{(\ell)} - \phi^{(\ell)}(\hat{\boldsymbol{\beta}}_s)| \to 0$$
(30)

where this convergence is in probability. Furthermore, the limiting functions form an orthonormal set with respect to the measure $\nu(\mathbf{B})$,

$$\langle \phi^{(\ell)}, \phi^{(m)} \rangle = \int_{\mathbf{B}} \phi^{(\ell)}(\boldsymbol{\beta}) \phi^{(m)}(\boldsymbol{\beta}) \mathrm{d}\nu(\boldsymbol{\beta}) = \delta_{\ell,m}.$$
(31)

Note that Assumption 2 follows from this assumption.

Remark 5. The \sqrt{n} term in (30) is necessary for the eigenvector normalization, since

$$\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)} \right)^{2} \to \sum_{s=1}^{n} \left(\frac{1}{\sqrt{n}} \phi^{(\ell)}(\hat{\boldsymbol{\beta}}_{s}) \right)^{2} \to \int_{\mathbf{B}} \left(\phi^{(\ell)}(\boldsymbol{\beta}) \right)^{2} \mathrm{d}\nu(\boldsymbol{\beta}) = 1.$$
(32)

By Propositions 2.1 and 2.2 of [24], the low-resolution mean and covariance estimates are consistent. However, unlike these aggregate quantities, the PCA coordinates $\hat{\beta}_s$ are computed from a single image, so they must contain an irreducible error term due to the finite noise level. We codify this in the following assumption.

Assumption 4. The estimated PCA coordinates are correct up to some stochastically bounded noise term,

$$\hat{\boldsymbol{\beta}}_s = \boldsymbol{\beta}(\mathbf{x}_s) + O_P(1). \tag{33}$$

We now show that, up to noise, the spectral volumes are merely voxel-wise orthogonal projections of the true volumes $\mathbf{x}_1, \ldots, \mathbf{x}_n$ onto a basis of eigenfunctions.

Corollary 1. Under Assumptions 1, 3 and 4 it follows from Theorem 1 that

$$\hat{\boldsymbol{\alpha}}^{(\ell)} = \mathbb{E}[\phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}) + O_P(1))\mathbf{x}] + O_P\left(\frac{1}{\sqrt{n}}\right)$$
(34)

where the expectation is with respect to the distribution of $\mathbf{x} \in \mathcal{M}$.

So far we have treated the convergence of the spectral volumes. We now turn to the convergence of the high-resolution reconstructions. As discussed in Section 3.1, we assume that the manifold of the molecular volumes can be well approximated by a small number of eigenfunctions. We define this notion precisely in the following assumption.

Assumption 5. There is some non-negative decreasing function h(r) that satisfies $h(r) \rightarrow 0$ such that h(r) bounds the approximation of \mathcal{M} by r spectral volumes:

$$\left\|\mathbf{x} - \sum_{\ell=0}^{r-1} \boldsymbol{\alpha}^{(\ell)} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}))\right\| = O(h(r)) \qquad \forall \mathbf{x} \in \mathcal{M}.$$
(35)

We can prove that the true volumes are recovered up to noise.

Theorem 2. Consider a sample from a manifold that conforms to (35), then it follows from Assumptions 1, 3, 4 and 5 that as $n \to \infty$ we have

$$\hat{\mathbf{x}}_{s} = \mathbf{x}_{s} + \sum_{\ell=0}^{r-1} O_{P}(C_{\ell}) \boldsymbol{\alpha}^{(\ell)} + O_{P}(h(r)) \left(\sum_{\ell=0}^{r-1} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_{s})) + O_{P}(C_{\ell}) \right), \quad (36)$$

where C_{ℓ} is an upper bound on the first derivative of $\phi^{(\ell)}$.

The proof of this theorem is in Appendix A. Note that the first error term contains an irreducible error from the finite level of noise in the PCA coordinate assignment.

6. Results

In this section, we apply our method to several synthetic datasets with a low-dimensional conformation space. We first consider clean images of a clock face with a rotating hand, then more realistic datasets of molecular volumes with one- and two-dimensional motions.

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 19



Figure 3. (left column) Clean 2D clock faces; (middle column) noisy clock faces, used as inputs to the reconstruction algorithm; (right column) corresponding reconstructions using r = 31 spectral volumes.



Figure 4. Spectral volumes $\hat{\alpha}^{(0)}, \ldots, \hat{\alpha}^{(15)}$ of the 2D clock. (left) First spectral volume $\alpha^{(0)}$ which converges to the mean; (right) Other spectral volumes ordered vertically in pairs of the same eigenvalue. Eigenvalues increase from left to right.

6.1. Clock dataset

We begin with a toy model of a 2D clock face with a single moving hand. Since the objects we wish to reconstruct are images rather than volumes, no projections are involved. This is the setting studied in Section 5.1. The simulated dataset is comprised of $n = 10^4$ noisy images $\mathbf{z}_1, \ldots, \mathbf{z}_n \in \mathbb{R}^{N \times N}$ with N = 128, where each image shows the clock hand at a random angle with additive Gaussian noise. The affinity matrix was chosen to be $W_{ij} = e^{-\|\mathbf{z}_i - \mathbf{z}_j\|^2/N^2\sigma^2}$, where σ is the standard deviation of the Gaussian noise. We then constructed a normalized graph Laplacian, extracted its eigenvectors, and computed r = 31 spectral volumes by solving the least-squares problem of (24). Figure 3 shows representative input images and their corresponding reconstructions and Figure 4 shows the estimated spectral volumes (images in this case). We need a large value of r to get good reconstructions, since the clock hand has sharp discontinuities.

To analyze this example we note that the clock dataset has the manifold geometry of the unit circle S^1 . Ignoring an arbitrary phase offset, the set of real eigenfunctions of the Laplace–Beltrami operator on S^1 are $\phi^{(0)}(\theta) = 1/2\pi$ and for all integer $\ell \geq 1$, Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 20

$$\phi^{(2\ell-1)}(\theta) = \sqrt{1/\pi} \sin(\ell\theta) \text{ and } \phi^{(2\ell)}(\theta) = \sqrt{1/\pi} \cos(\ell\theta). \text{ It follows from (27) that}$$
$$\hat{\boldsymbol{\alpha}}^{(\ell)} \to \mathbb{E}[\phi^{(\ell)}(\mathbf{z})\mathbf{z}] = \int_{\mathbf{z}\sim\mathcal{M}} \phi^{(\ell)}(\mathbf{z})\mathbf{z} \mathrm{d}\mathbf{z}. \tag{37}$$

Let \mathbf{z}_{θ} denote the image with the clock hand at angle θ , we may rewrite the above as

$$\hat{\boldsymbol{\alpha}}^{(\ell)} \to \int_0^{2\pi} \phi^{(\ell)}(\theta) \mathbf{z}_{\theta} \mathrm{d}\theta.$$
(38)

Rather than fixing a pixel and rotating the clock hand, we may fix the clock hand and rotate the pixel. For pixels [x, y] inside the disk of the clock face,

$$\hat{\boldsymbol{\alpha}}^{(\ell)}[x,y] \to \int_0^{2\pi} \phi^{(\ell)}(\theta) \mathbf{z}_{\theta=0}[R_{-\theta}[x,y]^T] \mathrm{d}\theta.$$
(39)

We conclude that in the case of simple rotation heterogeneity, a pixel of the ℓ^{th} spectral volume in the rotating domain converges to the ℓ^{th} Fourier coefficient of the function $f(\theta) = \mathbf{z}_0[R_{-\theta}[x, y]]$. Put differently, the coefficients $\hat{\boldsymbol{\alpha}}^{(0)}[x, y], \hat{\boldsymbol{\alpha}}^{(1)}[x, y], \ldots$ converge to the Fourier coefficients of polar representation of the rotating domain.

We tested a similar clock dataset in 3D, using the same clock hand shape, this time with tomographic projections. The results we obtained are similar to the results on the 2D clock dataset, in accordance with Corollary 1. See Appendix B for details.

6.2. Simulated ion channel

We created two synthetic datasets based on a voltage-gated potassium channel, shown in Figure 1. The first dataset ChannelSpin demonstrates a rotational motion of the top part about the z axis. The second dataset ChannelStretch demonstrates a nonrigid stretching of the bottom part. Specifically, each slice was displaced along the x-y plane by an amount that is proportional to the squared distance of the slice from the center of the molecule. We used a spatial resolution of N = 108 and generated n = 10,000 volumes for each dataset. The angles of rotation in the ChannelSpin dataset were drawn uniformly, which gives a conformational manifold diffeomorphic to the circle. For the ChannelStretch dataset, we drew x and y displacement values uniformly and independently between -16and +16 pixels, which gives a conformational manifold diffeomorphic to the square. In both datasets, we projected the molecules along random orientations, drawn uniformly from SO(3). We then applied a simulated point spread function with a defocus value chosen uniformly at random from 15.0, 16.7, 18.3, 20.0, 21.7, 23.3, or 25.0 pixels. Finally, we added white Gaussian noise, with a variance chosen such that the total energy of the noise was 30 times that of the total energy of each clean image. No additional in-plane shift was applied. See Figure 1 for example images.

Using the projection images, we ran the covariance estimation method with q = 4 components to build the adjacency matrix for the ChannelSpin dataset and and q = 8 for the ChannelStretch dataset. We then reconstructed the volumes using r = 1, ..., 15 spectral volumes. For both the covariance and spectral volume estimation procedures, we used the true orientations of the projection images.



Figure 5. Laplacian eigenmaps embedding of the ion channel datasets. (left) scatter plot of 1000 samples from the ChannelSpin dataset, showing the first two nontrivial Laplacian eigenvector coordinates. (right) 3D scatter plot of 2000 samples from the ChannelStretch dataset, showing the first three nontrivial eigenvectors. This dataset forms a saddle over a 2D square. The blue component of the color is given by the position along the line $\hat{\phi}_s^{(1)} = \hat{\phi}_s^{(2)}$, whereas the red component is given by the position along the line $\hat{\phi}_s^{(1)} = -\hat{\phi}_s^{(2)}$.

Examining the Laplacian eigenmaps embedding Figure 5 shows the embeddings of a random sample from the ChannelSpin and ChannelStretch datasets. The embedding of ChannelSpin clearly shows a circle whereas the embedding of the ChannelStretch dataset shows a 2-dimensional square in the $\hat{\phi}_s^{(1)} - \hat{\phi}_s^{(2)}$ plane that is shaped like a saddle. Both of these results are in accordance with the underlying geometry of the original motions.

Examining the spectral volumes Figure 6 shows the first few spectral volumes. For the ChannelSpin dataset, as expected $\hat{\alpha}^{(0)}$ captures the mean over all rotations. Higher order spectral volumes have increasing angular frequency, capturing more and more detail. Note that $\hat{\alpha}^{(1)}$ and $\hat{\alpha}^{(2)}$ are a quarter period out of phase, while $\hat{\alpha}^{(3)}$ has twice the angular frequency. Due to the C_4 symmetry of the ion channel, the lowest frequency of the ChannelSpin dataset has a period of 90 degrees.

For the ChannelStretch dataset, we see that $\hat{\boldsymbol{\alpha}}^{(0)}$ captures the fixed part of the molecule with high resolution and shows a "smeared" bottom portion. The first and second nontrivial spectral volumes each have a low spatial frequency along the x and y axes. Higher spectral volumes show higher spatial frequencies. $\hat{\boldsymbol{\alpha}}^{(3)}$ shows a mix of the directions of $\hat{\boldsymbol{\alpha}}^{(1)}$ and $\hat{\boldsymbol{\alpha}}^{(2)}$ whereas $\hat{\boldsymbol{\alpha}}^{(4)}$ is similar to $\hat{\boldsymbol{\alpha}}^{(1)}$ but with a double spatial frequency. Recall that by Corollary 1 up to noise the spectral volumes are $\mathbb{E}[\phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}))\mathbf{x}]$. In this case the eigenfunctions of the Laplacian on the square are the 2D discrete cosine transform basis functions, which are up to scale $\phi^{(n_x,n_y)} = \cos(n_x x) \cos(n_y y)$ with eigenvalue $\propto n_x^2 + n_y^2$, See [45, Section 3.1]. This agrees with our empirical observations.

Reconstruction accuracy Figure 7 shows reconstructions with increasing numbers of spectral volumes alongside the original simulated volume. Note that the reconstructions

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 22



Figure 6. Spectral volumes computed from the ChannelSpin dataset (top) and ChannelStretch dataset (bottom). The zeroth spectral volume $\hat{\alpha}^{(0)}$ (grey) is shown on the left. The next four figures, from left to right, show higher order spectral volumes, superimposed over $\hat{\alpha}^{(0)}$. Red and blue represent negative and positive values of the higher-order spectral volume, respectively. The two data sets are viewed from different angles.



Figure 7. Reconstructed volumes from ChannelSpin (top row) and ChannelStretch (bottom row), using $r \in \{1, 3, 5, 15\}$ spectral volumes. Also shown are the low-resolution reconstructions of the covariance-based method described in Section 3.3.

for the ChannelSpin volume are of higher quality than for the ChannelStretch volume. This is expected, since the manifold of conformations of ChannelSpin is one-dimensional whereas ChannelStretch has two-dimensional motion. Hence more samples are needed to get a dense cover of the conformational manifold.

To quantify the accuracy of our reconstructions, we use the *Fourier shell correlation* (FSC), which is the standard evaluation criterion in the cryo-EM literature [66]. Given

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 23



Figure 8. FSC curves for r = 2, ..., 16 from bottom to top, comparing reconstructed volumes to originals. Each curve is the FSC of $\mathbf{x}_s - \boldsymbol{\mu}$ and $\sum_{\ell=1}^r \hat{\phi}_s^{(\ell)} \hat{\alpha}^{(\ell)}$, averaged over s = 1, ..., n. (left) ChannelSpin; (right) ChannelStretch. The bottom x-axis denotes spatial frequency, and the top axis the corresponding wavelength (so that the rightmost position is the Nyquist frequency and a wavelength of 2 pixels.

two volumes $\mathbf{x}_1, \mathbf{x}_2$, the FSC takes their Fourier transforms and computes the correlation between each frequency shell. Because we wish to estimate the quality of reconstructing the variable part of the molecule, instead of reporting the FSC between the reconstructed volumes and the original volumes, we report the mean-subtracted FSC

$$\operatorname{FSC}\left(\mathbf{x}_{s}-\boldsymbol{\mu},\sum_{\ell=1}^{r-1}\hat{\phi}_{s}^{(\ell)}\hat{\boldsymbol{\alpha}}^{(\ell)}\right),\tag{40}$$

Figure 8 shows the results for each simulated dataset. As expected, the reconstruction quality increases with the number of spectral volumes. The reconstruction of the high frequencies is less accurate than that of the low frequencies. In both cases, as r increases, the FSC curves converge, indicating an upper bound on accuracy after which more spectral volumes yield diminishing returns. This bound is determined by the number of available images, and their noise levels.

6.3. Runtime

Table 2 details the running time of our method for the ChannelSpin simulation with N = 108, r = 16, and q = 8. The method is implemented in MATLAB 2017b and runs on 16 cores of a 2.3 GHz Intel Xeon CPU; memory usage was about 60 GB.

7. Conclusion

Today, rigid macromolecules are routinely reconstructed to near-atomic resolution using standard cryo-EM software tools. However, the high-resolution reconstruction of

Procedure	Running time (sec)
Calculation of $\hat{\mu}$	624.6
Calculation of $\hat{\Sigma}$	5044.7
Calculation $\hat{\mathbf{V}}_q$	0.8
Calculation of $\{\hat{oldsymbol{eta}}_s\}$	2084.8
Calculation of $\{\hat{\phi}_s\}$	531.5
Calculation of K	12378.0
Calculation of ${\bf b}$	4014.1
Estimation of $\{\hat{\boldsymbol{lpha}}^{(\ell)}\}$	1769.9

 Table 2. Runtimes for the main steps of our algorithm

molecular samples with continuous heterogeneity remains one of the grand challenges of the field. This work describes a new method which addresses this challenge. It combines spectral graph theory with recent techniques for covariance-based low-resolution reconstruction. Our procedure computes conformation-dependent Laplacian eigenmap coordinates and then generates a set of spectral volumes that characterize the variability of the molecule under study. Together these define a high-resolution 3D reconstruction for every projection image.

In the context of machine learning, our method combines and extends two classical methods: (i) the low-resolution covariance-based reconstruction, which we use to form the affinity graph, may be viewed as a generalization of PCA, as it finds the principal volumes in the space of molecular conformations. Unlike PCA, the input is projection images rather than full observations. (ii) the construction of an affinity graph from the low-resolution reconstructions and the generalized tomographic reconstruction using the eigenvectors of the graph Laplacian. This can be viewed as an extension of standard approaches for nonparametric regression, semi-supervised learning and matrix completion on graphs and manifolds (see for example [67, 68, 63, 69, 70]). The key difference is that rather than partially-labeled or noisy observations, we reconstruct a smooth high-dimensional function from noisy tomographic measurement. We note that the combination of PCA and graph Laplacian representations has been used for dimensionality reduction and denoising, for example in [71, 72].

Similar to the hyper-molecules method proposed by [73, 74], our method expands the molecular volumes which generated the projection images using a small set of basis volumes. However, in the hyper-molecules model the basis volumes are obtained from a user-specified manifold. Similarly, the multi-body refinement of RELION 3 [36] requires that the user manually segment the molecule into components that exhibit motion relative to each other. In contrast, our method is data-driven and requires no such user input. It relies only on the assumption that the molecule deforms continuously in a manner that is determined by a small number of parameters.

To conclude, in this paper we have described a method for the reconstruction of molecules with continuous heterogeneity, studied it theoretically and demonstrated the high-resolution reconstruction of synthetic data with one-dimensional and twodimensional motion manifolds. In future work, we will continue to scale up the method and apply it to the analysis of experimental datasets.

Acknowledgments

We would like to thank Ronald R. Coifman and Mark Tygert for interesting discussions. AS, AM, and AH were partially supported by NIGMS Award Number R01GM090200, AFOSR FA9550-17-1-0291, ARO W911NF-17-1-0512, Simons Investigator Award, the Moore Foundation Data-Driven Discovery Investigator Award, and NSF BIGDATA Award IIS-1837992. The Flatiron Institute is a division of the Simons Foundation. 3D molecular graphics were rendered using UCSF Chimera [75].

References

- Joachim Frank. Three-Dimensional Electron Microscopy of Macromolecular Assemblies. Oxford University Press, mar 2006. doi:10.1093/acprof:oso/9780195182187.001.0001
- [2] Miloš Vulović, Raimond B.G. Ravelli, Lucas J. van Vliet, Abraham J. Koster, Ivan Lazić, Uwe Lücken, Hans Rullgård, Ozan Öktem, and Bernd Rieger. Image formation modeling in cryoelectron microscopy. J. Struct. Biol., 183(1):19–32, jul 2013. doi:10.1016/j.jsb.2013.05.008
- [3] Werner Kühlbrandt. The Resolution Revolution. Science, 343(6178):1443–1444, mar 2014. doi:10.1126/science.1251652
- [4] Alexey Amunts et al. Structure of the Yeast Mitochondrial Large Ribosomal Subunit. Science, 343(6178):1485–1489, mar 2014. doi:10.1126/science.1249410
- [5] Maofu Liao, Erhu Cao, David Julius, and Yifan Cheng. Structure of the TRPV1 ion channel determined by electron cryo-microscopy. *Nature*, 504(7478):107–112, 2013. doi:10.1038/nature12822
- [6] Alberto Bartesaghi et al. Atomic Resolution Cryo-EM Structure of β-Galactosidase. Structure, 26(6):848–856.e3, 2018. doi:10.1016/j.str.2018.04.004
- [7] Alex Barnett, Leslie Greengard, Andras Pataki, and Marina Spivak. Rapid Solution of the Cryo-EM Reconstruction Problem by Frequency Marching. SIAM J. Imaging Sci., 10(3):1170–1195, jan 2017. doi:10.1137/16M1097171
- [8] Yifan Cheng, Nikolaus Grigorieff, Pawel A. Penczek, and Thomas Walz. A Primer to Single-Particle Cryo-Electron Microscopy. Cell, 161(3):438–449, 2015. doi:10.1016/j.cell.2015.03.050
- [9] Jacqueline L.S. Milne et al. Cryo-electron microscopy a primer for the non-microscopist. FEBS J., 280(1):28–45, 2012. doi:10.1111/febs.12078
- [10] Kutti R. Vinothkumar and Richard Henderson. Single particle electron cryomicroscopy: trends, issues and future perspective. Q. Rev. Biophys., 49:1–25, jul 2016. doi:10.1017/S0033583516000068
- [11] Sjors H.W. Scheres. A Bayesian View on Cryo-EM Structure Determination. J. Mol. Biol., 415(2):406–418, jan 2012. doi:10.1016/j.jmb.2011.11.010
- [12] Ryohei Yasuda, Hiroyuki Noji, Kazuhiko Kinosita, and Masasuke Yoshida. F1-ATPase Is a Highly Efficient Molecular Motor that Rotates with Discrete 120° Steps. Cell, 93(7):1117–1124, jun 1998. doi:10.1016/S0092-8674(00)81456-7
- [13] Sjors H.W. Scheres. RELION: Implementation of a Bayesian approach to cryo-EM structure determination. J. Struct. Biol., 180(3):519–530, dec 2012. doi:10.1016/j.jsb.2012.09.006
- [14] Ali Punjani, John L. Rubinstein, David J. Fleet, and Marcus A. Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. Nat. Methods, 14(3):290–296, mar 2017. doi:10.1038/nmeth.4169

- [15] Dmitry Lyumkis, Axel F. Brilot, Douglas L. Theobald, and Nikolaus Grigorieff. Likelihood-based classification of cryo-EM images using FREALIGN. J. Struct. Biol., 183(3):377–388, sep 2013. doi:10.1016/j.jsb.2013.07.005
- [16] Timothy Grant, Alexis Rohou, and Nikolaus Grigorieff. cisTEM, user-friendly software for singleparticle image processing. *Elife*, 7(3):377–388, mar 2018. doi:10.7554/eLife.35383
- [17] Guang Tang, Liwei Peng, Philip R. Baldwin, Deepinder S. Mann, Wen Jiang, Ian Rees, and Steven J. Ludtke. EMAN2: An extensible image processing suite for electron microscopy. J. Struct. Biol., 157(1):38–46, jan 2007. doi:10.1016/j.jsb.2006.05.009
- [18] Weiping Liu and Joachim Frank. Estimation of variance distribution in three-dimensional reconstruction I Theory. J. Opt. Soc. Am. A, 12(12):2615, dec 1995. doi:10.1364/JOSAA.12.002615
- [19] Pawel A. Penczek. Variance in three-dimensional reconstructions from projections. In Proc. IEEE Int. Symp. Biomed. Imaging, pages 749–752. IEEE, 2002. doi:10.1109/ISBI.2002.1029366
- [20] Pawel A. Penczek, Chao Yang, Joachim Frank, and Christian M.T. Spahn. Estimation of variance in single-particle reconstruction using the bootstrap technique. J. Struct. Biol., 154(2):168–183, may 2006. doi:10.1016/j.jsb.2006.01.003
- [21] Pawel A. Penczek, Joachim Frank, and Christian M.T. Spahn. A method of focused classification, based on the bootstrap 3D variance analysis, and its application to EF-G-dependent translocation. J. Struct. Biol., 154(2):184–194, may 2006. doi:10.1016/j.jsb.2005.12.013
- [22] Pawel A. Penczek, Marek Kimmel, and Christian M.T. Spahn. Identifying Conformational States of Macromolecules by Eigen-Analysis of Resampled Cryo-EM Images. *Structure*, 19(11):1582–1590, nov 2011. doi:10.1016/j.str.2011.10.003
- [23] Hstau Y. Liao and Joachim Frank. Classification by bootstrapping in single particle methods. In 2010 IEEE Int. Symp. Biomed. Imaging From Nano to Macro, volume 29169, pages 169–172. IEEE, 2010. doi:10.1109/ISBI.2010.5490386
- [24] Eugene Katsevich, Alexander Katsevich, and Amit Singer. Covariance Matrix Estimation for the Cryo-EM Heterogeneity Problem. SIAM J. Imaging Sci., 8(1):126–185, jan 2015. doi:10.1137/130935434
- [25] Joakim Andén, Eugene Katsevich, and Amit Singer. Covariance estimation using conjugate gradient for 3D classification in cryo-EM. In 2015 IEEE 12th Int. Symp. Biomed. Imaging (ISBI), pages 200–204. IEEE, apr 2015. doi:10.1109/ISBI.2015.7163849
- [26] Joakim Andén and Amit Singer. Structural Variability from Noisy Tomographic Projections. SIAM J. Imaging Sci., 11(2):1441–1492, may 2018. doi:10.1137/17M1153509
- [27] Hemant D. Tagare, Alp Kucukelbir, Fred J. Sigworth, Hongwei Wang, and Murali Rao. Directly reconstructing principal components of heterogeneous particles from cryo-EM images. J. Struct. Biol., 191(2):245–262, aug 2015. doi:10.1016/j.jsb.2015.05.007
- [28] Ali Dashti, et al. Trajectories of the ribosome as a Brownian nanomachine. Proc. Natl. Acad. Sci., 111(49):17492–17497, dec 2014. doi:10.1073/pnas.1419276111
- [29] Peter Schwander, Russell Fung, and Abbas Ourmazd. Conformations of macromolecules and their complexes from heterogeneous datasets. *Philos. Trans. R. Soc. B Biol. Sci.*, 369(1647):1–8, jun 2014. doi:10.1098/rstb.2013.0567
- [30] Joachim Frank and Abbas Ourmazd. Continuous changes in structure mapped by manifold embedding of single-particle data in cryo-EM. *Methods*, 100:61–67, may 2016. doi:10.1016/j.ymeth.2016.02.007
- [31] Ali Dashti, Danya Ben Hail, Ghoncheh Mashayekhi, Peter Schwander, Amedée des Georges, Joachim Frank, and Abbas Ourmazd. Functional Pathways of Biomolecules Retrieved from Single-particle Snapshots. Technical report, 2018. doi:10.1101/291922
- [32] Ronald R. Coifman, Stéphane Lafon, Ann B. Lee, Mauro Maggioni, Boaz Nadler, Fred Warner, and Steven W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proc. Natl. Acad. Sci.*, 102(21):7426–7431, may 2005. doi:10.1073/pnas.0500334102
- [33] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. Appl. Comput. Harmon. Anal., 21(1):5–30,

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 27

jul 2006. doi:10.1016/j.acha.2006.04.006

- [34] Chang Wang and Sridhar Mahadevan. Manifold alignment without correspondence. Int. Jt. Conf. Artif. Intell. (IJCAI), 1:1273–1278, 2009.
- [35] Zhen Cui, Hong Chang, Shiguang Shan, and Xilin Chen. Generalized Unsupervised Manifold Alignment. In Neural Inf. Process. Syst. (NIPS), pages 2429–2437, 2014.
- [36] Takanori Nakane, Dari Kimanius, Erik Lindahl, and Sjors H.W. Scheres. Characterisation of molecular motions in cryo-EM single-particle data by multi-body refinement in RELION. *eLife*, 7:1–18, jun 2018. doi:10.7554/eLife.36861
- [37] Qiyu Jin, Carlos Oscar S. Sorzano, José Miguel de la Rosa-Trevín, José Román Bilbao-Castro, Rafael Núñez-Ramírez, Oscar Llorca, Florence Tama, and Slavica Jonić. Iterative Elastic 3D-to-2D Alignment Method Using Normal Modes for Studying Structural Dynamics of Large Macromolecular Complexes. Structure, 22(3):496–506, mar 2014. doi:10.1016/j.str.2014.01.004
- [38] Sandra Schilbach, Merle Hantsche, Dmitry Tegunov, Christian Dienemann, Cristoph Wigge, Henning Urlaub, and Patrick Cramer. Structures of transcription pre-initiation complex with TFIIH and Mediator. Nature, 551(7679):204–209, nov 2017. doi:10.1038/nature24282
- [39] Carlos Oscar S. Sorzano, et al. Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. Acta Crystallogr. Sect. F Struct. Biol. Commun., 75(1):19–32, jan 2019. doi:10.1107/S2053230X18015108
- [40] Lanhui Wang, Yoel Shkolnisky, and Amit Singer. A Fourier-based Approach for Iterative 3D Reconstruction from Cryo-EM Images. jul 2013. arXiv:1307.5824
- [41] Alok Dutt and Vladimir Rokhlin. Fast Fourier Transforms for Nonequispaced Data. SIAM J. Sci. Comput., 14(6):1368–1393, nov 1993. doi:10.1137/0914081
- [42] Leslie Greengard and June-Yub Lee. Accelerating the Nonuniform Fast Fourier Transform. SIAM Rev., 46(3):443–454, jan 2004. doi:10.1137/S003614450343200X
- [43] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. Neural Comput., 15(6):1373–1396, jun 2003. doi:10.1162/089976603321780317
- [44] Frank Natterer. The Mathematics of Computerized Tomography. Society for Industrial and Applied Mathematics, jan 2001. doi:10.1137/1.9780898719284
- [45] Denis S. Grebenkov and Binh T. Nguyen. Geometrical Structure of Laplacian Eigenfunctions. SIAM Rev., 55(4):601–667, jan 2013. doi:10.1137/120880173
- [46] Yonathan Aflalo, Haim Brezis, and Ron Kimmel. On the Optimality of Shape and Data Representation in the Spectral Domain. SIAM J. Imaging Sci., 8(2):1141–1160, jan 2015. doi:10.1137/140977680
- [47] Wlodzimier Greblicki and Miroslaw Pawlak. Fourier and Hermite series estimates of regression functions. Ann. Inst. Stat. Math., 37(3):443, dec 1985. doi:10.1007/BF02481112
- [48] Hstau Y. Liao, Yaser Hashem, and Joachim Frank. Efficient Estimation of Three-Dimensional Covariance and its Application in the Analysis of Heterogeneous Samples in Cryo-Electron Microscopy. Structure, 23(6):1129–1137, jun 2015. doi:10.1016/j.str.2015.04.004
- [49] Gabor T. Herman and Miroslaw Kalinowski. Classification of heterogeneous electron microscopic projections into homogeneous subsets. Ultramicroscopy, 108(4):327–338, mar 2008. doi:10.1016/j.ultramic.2007.05.005
- [50] Maxim Shatsky, Richard J. Hall, Eva Nogales, Jitendra Malik, and Steven E. Brenner. Automated multi-model reconstruction from single-particle electron microscopy data. J. Struct. Biol., 170(1):98–108, apr 2010. doi:10.1016/j.jsb.2010.01.007
- [51] Stéphane Mallat. A Wavelet Tour of Signal Processing. Elsevier, 2009. doi:10.1016/B978-0-12-374370-1.X0001-8
- [52] Steven M. Kay. Fundamentals of Statistical Signal Processing: Estimation Theory. Prentice-Hall, 1993.
- [53] Ulrike von Luxburg. A tutorial on spectral clustering. Stat. Comput., 17(4):395–416, dec 2007. doi:10.1007/s11222-007-9033-z
- [54] Jerome H. Friedman, Jon Louis Bentley, and Raphael Ari Finkel. An Algorithm for Finding Best

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 28

Matches in Logarithmic Expected Time. ACM Trans. Math. Softw., 3(3):209–226, sep 1977. doi:10.1145/355744.355745

- [55] G.W. Stewart. A Krylov–Schur Algorithm for Large Eigenproblems. SIAM J. Matrix Anal. Appl., 23(3):601–614, jan 2002. doi:10.1137/S0895479800371529
- [56] James McQueen, Marina Meilă, Jacob VanderPlas, and Zhongyue Zhang. Megaman: Scalable Manifold Learning in Python. J. Mach. Learn. Res. (JMLR), 17(148):1–5, 2016.
- [57] Luke N. Olson and Jacob B. Schroder. PyAMG: Algebraic Multigrid Solvers in Python v4.0, 2018.
- [58] Alex H. Barnett, Jeremy F. Magland, and Ludvig af Klinteberg. A parallel non-uniform fast Fourier transform library based on an "exponential of semicircle" kernel. SIAM J. Sci. Comput. Press., pages 1–25, apr 2019.
- [59] Gene H. Golub and Charles F. Van Loan. Matrix Computations. Johns Hopkins University Press, 4th edition, 2013.
- [60] Lloyd N. Trefethen and David Bau. Numerical Linear Algebra. SIAM, Philadelphia, PA 19104-2688, jan 1997.
- [61] Ulrike von Luxburg, Mikhail Belkin, and Olivier Bousquet. Consistency of spectral clustering. Ann. Stat., 36(2):555–586, apr 2008. doi:10.1214/00905360700000640
- [62] Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On Learning with Integral Operators. J. Mach. Learn. Res. (JMLR), 11:905–934, 2010.
- [63] Ann B. Lee and Rafael Izbicki. A spectral series approach to high-dimensional nonparametric regression. *Electron. J. Stat.*, 10(1):423–463, 2016. doi:10.1214/16-EJS1112
- [64] Boaz Nadler, Stephane Lafon, Ronald R. Coifman, and Ioannis G. Kevrekidis. Diffusion Maps, Spectral Clustering and Eigenfunctions of Fokker–Planck operators. In *Neural Inf. Process. Syst.* (*NIPS*), pages 955–962, 2005.
- [65] Daniel Ting, Ling Huang, and Michael Jordan. An Analysis of the Convergence of Graph Laplacians. In Int. Conf. Mach. Learn. (ICML), 2010.
- [66] Richard Henderson, et al. Outcome of the First Electron Microscopy Validation Task Force Meeting. Structure, 20(2):205–214, feb 2012. doi:10.1016/j.str.2011.12.014
- [67] Mikhail Belkin and Partha Niyogi. Semi-Supervised Learning on Riemannian Manifolds. Mach. Learn., 56(1-3):209–239, jul 2004. doi:10.1023/B:MACH.0000033120.25363.1e
- [68] Xueyuan Zhou and Nathan Srebro. Error Analysis of Laplacian Eigenmaps for Semi-supervised Learning. Int. Conf. Artif. Intell. Stat. (AISTATS), 15:901–908, 2011.
- [69] Amit Moscovich, Ariel Jaffe, and Boaz Nadler. Minimax-optimal semi-supervised regression on unknown manifolds. In Int. Conf. Artif. Intell. Stat. (AISTATS), volume 54, pages 933–942, apr 2017.
- [70] Paul Villoutreix, Joakim Andén, Bomyi Lim, Hang Lu, Ioannis G. Kevrekidis, Amit Singer, and Stanislav Y. Shvartsman. Synthesizing developmental trajectories. *PLOS Comput. Biol.*, 13(9):1–15, sep 2017. doi:10.1371/journal.pcbi.1005742
- [71] Amit Singer and Hau-Tieng Wu. Two-Dimensional Tomography from Noisy Projections Taken at Unknown Random Directions. SIAM J. Imaging Sci., 6(1):136–175, jan 2013. doi:10.1137/090764657
- [72] Amit Singer. Spectral independent component analysis. Appl. Comput. Harmon. Anal., 21(1):135– 144, jul 2006. doi:10.1016/j.acha.2006.03.003
- [73] Roy R. Lederman and Amit Singer. Continuously heterogeneous hyper-objects in cryo-EM and 3-D movies of many temporal dimensions. 2017. arXiv:1704.02899
- [74] Roy R. Lederman, Joakim Andén, and Amit Singer. Hyper-Molecules: on the Representation and Recovery of Dynamical Structures, with Application to Flexible Macro-Molecular Structures in Cryo-EM. submitted to Inverse Problems.
- [75] Eric F. Pettersen, Thomas D. Goddard, Conrad C. Huang, Gregory S. Couch, Daniel M. Greenblatt, Elaine C. Meng, and Thomas E. Ferrin. UCSF Chimera—A visualization system for exploratory research and analysis. J. Comput. Chem., 25(13):1605–1612, oct 2004. doi:10.1002/jcc.20084

Appendix A. Proofs

Before proving Theorem 1, we require a technical lemma that bounds the variance of linear combinations of random variables.

Lemma 1. Fix n and let Z_1, \ldots, Z_n be i.i.d. random variables with finite variance. Let W_1, \ldots, W_n be identically distributed, but possibly dependent, random weights. Denote the (unnormalized) sample moments by

$$\sum_{s=1}^{n} W_s = m_1 \qquad \sum_{s=1}^{n} W_s^2 = m_2.$$
(A.1)

If the weights W_1, \ldots, W_n are independent of Z_1, \ldots, Z_n , we have

$$\operatorname{Var}\left(\sum_{s=1}^{n} W_s Z_s\right) \le m_2 \mathbb{E}[Z_1^2]. \tag{A.2}$$

Proof. By definition,

$$\operatorname{Var}\left(\sum_{s=1}^{n} W_s Z_s\right) = \mathbb{E}\left[\left(\sum_{s=1}^{n} W_s Z_s\right)^2\right] - \left(\mathbb{E}\left[\sum_{s=1}^{n} W_s Z_s\right]\right)^2.$$
(A.3)

Since W_s and Z_s are independent, we can rewrite the second term, which yields

$$\left(\mathbb{E}\left[\sum_{s=1}^{n} W_s Z_s\right]\right)^2 = \left(\sum_{s=1}^{n} \mathbb{E}[W_s]\mathbb{E}[Z_s]\right)^2 = m_1^2 \left(\mathbb{E}[Z_1]\right)^2.$$
(A.4)

Similarly, $W_s W_t$ is independent of Z_1, \ldots, Z_n , so we may break up the expectations in the first term of (A.3). We then split the double sum into a sum over index pairs s = t and a sum over $s \neq t$, obtaining

$$\mathbb{E}\left[\left(\sum_{s=1}^{n} W_s Z_s\right)^2\right] = \sum_{s,t=1}^{n} \mathbb{E}[W_s W_t] \mathbb{E}[Z_s Z_t]$$
(A.5)

$$= \sum_{s=1}^{n} \mathbb{E}[W_s^2] \mathbb{E}[Z_s^2] + \sum_{t=1}^{n} \sum_{s \neq t} \mathbb{E}[W_s W_t] \mathbb{E}[Z_s Z_t] \quad (A.6)$$

$$= \mathbb{E}[Z_1^2] \sum_{s=1}^n \mathbb{E}[W_s^2] + (\mathbb{E}[Z_1])^2 \sum_{t=1}^n \sum_{s \neq t} \mathbb{E}[W_s W_t].$$
(A.7)

The second term may be bounded by the constraint $\sum_{s=1}^{n} W_s = m_1$, since

$$\left(\mathbb{E}[Z_1]\right)^2 \mathbb{E}\left[\sum_{t=1}^n \sum_{s \neq t} W_s W_t\right] \le \left(\mathbb{E}[Z_1]\right)^2 \mathbb{E}\left[\left(\sum_{s=1}^n W_s\right)^2\right] = \left(\mathbb{E}[Z_1]\right)^2 m_1^2. \quad (A.8)$$

Putting it all together and incorporating the second moment constraint, we obtain

$$\operatorname{Var}\left(\sum_{s=1}^{n} W_{s} Z_{s}\right) \leq \mathbb{E}[Z_{s}^{2}] \sum_{s=1}^{n} \mathbb{E}[W_{s}^{2}] \leq m_{2} \mathbb{E}[Z_{1}^{2}].$$
(A.9)

Proof of Theorem 1.

We will prove the convergence of the solution to (22) by proving that both **b** and K converge in probability as $n \to \infty$. We start by computing the expectation of **b**.

$$\mathbb{E}[\mathbf{b}^{(\ell)}] = \mathbb{E}\left[\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{P}_{s}^{\mathrm{T}}\mathbf{y}_{s}\right]$$
(A.10)

$$= \sum_{s=1}^{n} \frac{1}{\sqrt{n}} \mathbb{E}[\hat{\phi}_{s}^{(\ell)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s} \mathbf{x}_{s} + \hat{\phi}_{s}^{(\ell)} \mathbf{P}_{s}^{\mathrm{T}} \boldsymbol{\varepsilon}_{s}] \qquad (\mathrm{By}\ (2))$$
(A.11)

$$=\sum_{s=1}^{n} \frac{1}{\sqrt{n}} \mathbb{E}[\hat{\phi}_{s}^{(\ell)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s} \mathbf{x}_{s}] + 0 \tag{A.12}$$

$$= \mathbb{E}[\mathbf{P}^{\mathrm{T}}\mathbf{P}]\mathbb{E}\left[\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{x}_{s}\right] \qquad (\text{By Assumption 1}).$$
(A.13)

Now consider the variance of the i^{th} element of the vector $\mathbf{b}^{(\ell)}$,

$$\operatorname{Var}(\mathbf{e}_{i}^{\mathrm{T}}\mathbf{b}^{(\ell)}) = \operatorname{Var}\left(\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{e}_{i}^{\mathrm{T}}\mathrm{P}_{s}^{\mathrm{T}}\mathbf{y}_{s}\right).$$
(A.14)

We apply Lemma 1 with $W_s = \hat{\phi}_s^{(\ell)}$ and $Z_s = \mathbf{e}_i^{\mathrm{T}} \mathbf{P}_s^{\mathrm{T}} \mathbf{y}_s$ to obtain

$$\operatorname{Var}(\mathbf{e}_{i}^{\mathrm{T}}\mathbf{b}^{(\ell)}) < \frac{1}{n}\mathbb{E}[Z_{1}^{2}].$$
(A.15)

We now compute the expectation and variance of the matrix K.

$$\mathbb{E}\left[\mathbf{K}^{(\ell,m)}\right] = \sum_{s=1}^{n} \mathbb{E}\left[\hat{\phi}_{s}^{(\ell)} \hat{\phi}_{s}^{(m)} \mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s}\right]$$
(A.16)

$$= \sum_{s=1}^{n} \mathbb{E}[\hat{\phi}_{s}^{(\ell)} \hat{\phi}_{s}^{(m)}] \mathbb{E}[\mathbf{P}_{s}^{\mathrm{T}} \mathbf{P}_{s}] \qquad (\text{By Assumption 1})$$
(A.17)

$$= \mathbb{E}\left[\sum_{s} \hat{\phi}_{s}^{(\ell)} \hat{\phi}_{s}^{(m)}\right] \mathbb{E}[\mathbf{P}^{\mathrm{T}}\mathbf{P}] = \delta_{\ell,m} \mathbb{E}[\mathbf{P}^{\mathrm{T}}\mathbf{P}].$$
(A.18)

For the variance, we compute the variance of a single entry $\mathbf{K}_{i,j}^{(\ell,m)} = \mathbf{e}_i^{\mathrm{T}} \mathbf{K}^{(\ell,m)} \mathbf{e}_j$. Case 1: $\ell = m$

$$\operatorname{Var}\left(\mathbf{e}_{i}^{\mathrm{T}}\mathrm{K}^{(\ell,\ell)}\mathbf{e}_{j}\right) = \operatorname{Var}\left(\sum_{s=1}^{n}\left(\hat{\phi}_{s}^{(\ell)}\right)^{2}\mathbf{e}_{i}^{\mathrm{T}}\mathrm{P}_{s}^{\mathrm{T}}\mathrm{P}_{s}\mathbf{e}_{j}\right).$$
(A.19)

Let $W_s = (\hat{\phi}_s^{(\ell)})^2$ and $Z_s = \mathbf{e}_i^{\mathrm{T}} \mathbf{P}_s^{\mathrm{T}} \mathbf{P}_s \mathbf{e}_j$. By Assumption 2, $\sum_{s=1}^n W_s^2 \to 0$. Since Z_s has finite variance, we apply Lemma 1 to obtain that as n tends to infinity, $\operatorname{Var}(\mathbf{e}_i^{\mathrm{T}} \mathbf{K}^{(\ell,\ell)} \mathbf{e}_j) \to 0$ in probability. Case 2: $\ell \neq m$

$$\operatorname{Var}\left(\mathbf{e}_{i}^{\mathrm{T}}\mathrm{K}^{(\ell,m)}\mathbf{e}_{j}\right) = \operatorname{Var}\left(\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\hat{\phi}_{s}^{(m)}\mathbf{e}_{i}^{\mathrm{T}}\mathrm{P}_{s}^{\mathrm{T}}\mathrm{P}_{s}\mathbf{e}_{j}\right).$$
(A.20)

Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes 31

By Cauchy–Schwarz and Assumption 2, the following converges in probability:

$$\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)}\right)^{2} \left(\hat{\phi}_{s}^{(m)}\right)^{2} \leq \sqrt{\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(\ell)}\right)^{4}} \sqrt{\sum_{s=1}^{n} \left(\hat{\phi}_{s}^{(m)}\right)^{4}} \to 0. \tag{A.21}$$

Again, we apply Lemma 1. This time with $W_s = \hat{\phi}_s^{(\ell)} \hat{\phi}_s^{(m)}$ and $Z_s = \mathbf{e}_i^{\mathrm{T}} \mathbf{P}_s^{\mathrm{T}} \mathbf{P}_s \mathbf{e}_j$ to obtain that $\operatorname{Var}(\mathbf{e}_i^{\mathrm{T}} \mathrm{K}^{(\ell,m)} \mathbf{e}_i) \to 0$ in probability. To summarize, we proved the following results:

$$\mathbf{b}^{(\ell)} = \mathbb{E}[\mathbf{P}^{\mathrm{T}}\mathbf{P}]\mathbb{E}\left[\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{x}_{s}\right] + O_{\mathrm{P}}\left(\frac{1}{\sqrt{n}}\right)$$
(A.22)

$$\mathbf{K}^{(\ell,m)} \to \delta_{\ell,m} \mathbb{E}[\mathbf{P}^{\mathrm{T}}\mathbf{P}]$$
 In probability. (A.23)

By (22), the vector of spectral volumes satisfies $\boldsymbol{\alpha} = \mathbf{K}^{-1}\mathbf{b}$. Denote $\mathbf{K} = \mathbb{E}[\mathbf{K}] + \Delta \mathbf{K}$ and $\mathbf{b} = \mathbb{E}[\mathbf{b}] + \Delta \mathbf{b}$. Expanding K^{-1} to first order,

$$K^{-1} = \mathbb{E}[K]^{-1} + \mathbb{E}[K]^{-1} \Delta K \mathbb{E}[K]^{-1} + O(\|\Delta K\|^2).$$
(A.24)

Since $\Delta K \to 0$ and $\Delta b = O_P(1/\sqrt{n})$, the spectral volumes satisfy,

$$\boldsymbol{\alpha} = (\mathbb{E}[\mathbf{K}] + \Delta \mathbf{K})^{-1} (\mathbb{E}[\mathbf{b}] + \Delta \mathbf{b}) = \mathbb{E}[\mathbf{K}]^{-1} \mathbb{E}[\mathbf{b}] + O_{\mathbf{P}}(1/\sqrt{n}).$$
(A.25)

Plugging in Equations (A.22) and (A.23), we obtain

$$\hat{\boldsymbol{\alpha}}^{(\ell)} = \mathbb{E}\left[\frac{1}{\sqrt{n}}\sum_{s=1}^{n}\hat{\phi}_{s}^{(\ell)}\mathbf{x}_{s}\right] + O_{\mathrm{P}}\left(\frac{1}{\sqrt{n}}\right).$$
(A.26)

Proof of Theorem 2. It follows from Corollary 1 and Assumption 5 that

$$\hat{\boldsymbol{\alpha}}^{(\ell)} = \mathbb{E}[\phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}) + O_P(1))\mathbf{x}] + O_P(1/\sqrt{n})$$
(A.27)

$$= \mathbb{E}\left[\phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}) + O_P(1))\left(O(h(r)) + \sum_{m=0}^{r-1} \boldsymbol{\alpha}^{(m)} \phi^{(m)}(\boldsymbol{\beta}(\mathbf{x}))\right)\right] + O_P\left(\frac{1}{\sqrt{n}}\right) (A.28)$$
$$= \boldsymbol{\alpha}^{(\ell)} + O_P(h(r)) \qquad (hr \mathbb{E}[\phi^{(\ell)} \phi^{(m)}] = \boldsymbol{\delta}$$

$$= \boldsymbol{\alpha}^{(\ell)} + O_P(h(r)). \qquad (by \ \mathbb{E}[\phi^{(\ell)}\phi^{(m)}] = \delta_{\ell,m})$$
(A.29)

By the definition of the high-resolution reconstructions (14) we now have

$$\hat{\mathbf{x}}_{s} = \sqrt{n} \sum_{\ell} \hat{\phi}_{s}^{(\ell)} \hat{\boldsymbol{\alpha}}^{(\ell)} = \sqrt{n} \sum_{\ell} \hat{\phi}_{s}^{(\ell)} (\boldsymbol{\alpha}^{(\ell)} + O_{P}(h(r)))$$
(A.30)

By Assumption 3 we may rewrite this as

$$\hat{\mathbf{x}}_s = \sum_{\ell} (\phi^{(\ell)}(\hat{\boldsymbol{\beta}}_s) + o_P(1))(\boldsymbol{\alpha}^{(\ell)} + O_P(h(r)))$$
(A.31)

By Assumption 4 we have $\phi^{(\ell)}(\hat{\boldsymbol{\beta}}_s) = \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_s) + O_P(1)) = \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_s)) + O_P(C_\ell)$, where the last equality stems from the fact that $\phi^{(\ell)}$ is a smooth function on a compact domain and therefore its derivatives are bounded. We plug this back into (A.31),

$$\hat{\mathbf{x}}_s = \sum_{\ell} (\phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_s)) + O_P(C_{\ell}))(\boldsymbol{\alpha}^{(\ell)} + O_P(h(r)))$$
(A.32)

$$=\sum_{\ell} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_s))\boldsymbol{\alpha}^{(\ell)} + \sum_{\ell} O_P(C_{\ell})\boldsymbol{\alpha}^{(\ell)}$$
(A.33)

$$+ O_P(h(r)) \sum_{\ell} \phi^{(\ell)}(\boldsymbol{\beta}(\mathbf{x}_s)) + \sum_{\ell} O_P(C_\ell) O_P(h(r)).$$
(A.34)
proof by reusing Assumption 5 on the first term.

We conclude the proof by reusing Assumption 5 on the first term.

Appendix B. 3D Clock



Figure B1. Two poses of the clock model and their reconstructions using r = 7.



Figure B2. 3D clock spectral volumes, shown superimposed on $\alpha^{(0)}$ to aid in context. Grey is $\alpha^{(0)}$, red and blue are negative and positive portions of the higher order spectral volume, respectively.

This appendix presents the 3D analogue to the 2D clock simulation. Here the spatial resolution is N = 256 and the number of volumes is $n = 10^5$. Clock hand angles were drawn uniformly at random from the circle and the viewing orientations were drawn at uniformly from SO(3). No noise was added and no CTF applied in order to test the behavior under ideal conditions. The covariance estimation method was run with q = 8 components in order to generate the adjacency matrix. We then constructed a symmetric normalized graph Laplacian and performed reconstructions with r = 16 spectral volumes. Figure B2 shows the spectral volumes for the 3D clock. As can be seen, the spectral volumes here resemble very closely those from Section 4, limited to the region that the clock hand rotates in. The zeroth spectral volume looks like the mean volume whereas higher order spectral volumes come in pairs of increasing angular frequency.