

Published in final edited form as:

SIAM J Imaging Sci. 2015 January 22; 8(1): 126–185. doi:10.1137/130935434.

Covariance Matrix Estimation for the Cryo-EM Heterogeneity Problem*

E. Katsevich[†], A. Katsevich[‡], and A. Singer[§]

[†] Department of Mathematics, Princeton University, Princeton, NJ 08544

[‡] Department of Mathematics, University of Central Florida, Orlando, FL 32816

[§] Department of Mathematics and PACM, Princeton University, Princeton, NJ 08544-1000

Abstract

In cryo-electron microscopy (cryo-EM), a microscope generates a top view of a sample of randomly oriented copies of a molecule. The problem of single particle reconstruction (SPR) from cryo-EM is to use the resulting set of noisy two-dimensional projection images taken at unknown directions to reconstruct the three-dimensional (3D) structure of the molecule. In some situations, the molecule under examination exhibits structural variability, which poses a fundamental challenge in SPR. The heterogeneity problem is the task of mapping the space of conformational states of a molecule. It has been previously suggested that the leading eigenvectors of the covariance matrix of the 3D molecules can be used to solve the heterogeneity problem. Estimating the covariance matrix is challenging, since only projections of the molecules are observed, but not the molecules themselves. In this paper, we formulate a general problem of covariance estimation from noisy projections of samples. This problem has intimate connections with matrix completion problems and high-dimensional principal component analysis. We propose an estimator and prove its consistency. When there are finitely many heterogeneity classes, the spectrum of the estimated covariance matrix reveals the number of classes. The estimator can be found as the solution to a certain linear system. In the cryo-EM case, the linear operator to be inverted, which we term the projection covariance transform, is an important object in covariance estimation for tomographic problems involving structural variation. Inverting it involves applying a filter akin to the ramp filter in tomography. We design a basis in which this linear operator is sparse and thus can be tractably inverted despite its large size. We demonstrate via numerical experiments on synthetic datasets the robustness of our algorithm to high levels of noise.

Keywords

cryo-electron microscopy; X-ray transform; inverse problems; structural variability; classification; heterogeneity; covariance matrix estimation; principal component analysis; high-dimensional statistics; Fourier projection slice theorem; spherical harmonics

*Received by the editors September 3, 2013; accepted for publication (in revised form) September 22, 2014; published electronically January 22, 2015.

© 2015 Society for Industrial and Applied Mathematics. Copyright © by SIAM
(ekatsevi@princeton.edu). (alexander.katsevich@ucf.edu). (amits@math.princeton.edu).

1. Introduction

1.1. Covariance matrix estimation from projected data

Covariance matrix estimation is a fundamental task in statistics. Statisticians have long grappled with the problem of estimating this statistic when the samples are only partially observed. In this paper, we consider this problem in the general setting where “partial observations” are arbitrary linear projections of the samples onto a lower-dimensional space.

Problem 1.1—Let \mathbf{X} be a random vector on C^p , with $E[\mathbf{X}] = \mu_0$ and $\text{Var}(\mathbf{X}) = \Sigma_0$ ($\text{Var}[\mathbf{X}]$ denotes the covariance matrix of \mathbf{X}). Suppose also that \mathbf{P} is a random $q \times p$ matrix with complex entries, and \mathbf{E} is a random vector in C^q with $E[\mathbf{E}] = 0$ and $\text{Var}[\mathbf{E}] = \sigma^2 \mathbf{I}_q$. Finally, let \mathbf{I} denote the random vector in C^q given by

$$\mathbf{I} = \mathbf{P}\mathbf{X} + \mathbf{E}. \quad (1.1)$$

Assume now that \mathbf{X} , \mathbf{P} , and \mathbf{E} are independent. Estimate μ_0 and Σ_0 given observations $\mathbf{I}_1, \dots, \mathbf{I}_n$ and P_1, \dots, P_n of \mathbf{I} and \mathbf{P} , respectively.

Here, and throughout this paper, we write random quantities in boldface to distinguish them from deterministic quantities. We use regular font (e.g., X) for vectors and matrices, calligraphic font (e.g., \mathcal{X}) for functions, and script font for function spaces (e.g., \mathcal{B}). We denote true parameter values with a subscript of zero (e.g., μ_0), estimated parameter values with a subscript of n (e.g., μ_n), and generic variables with no subscript (e.g., μ).

Problem 1.1 is quite general, and has many practical applications as special cases. The main application this paper addresses is the heterogeneity problem in single particle reconstruction (SPR) from cryo-electron microscopy (cryo-EM). SPR from cryo-EM is an inverse problem where the goal is to reconstruct a three-dimensional (3D) molecular structure from a set of its two-dimensional (2D) projections from random directions [12]. The heterogeneity problem deals with the situation in which the molecule to be reconstructed can exist in several structural classes. In the language of Problem 1.1, \mathbf{X} represents a discretization of the molecule (random due to heterogeneity), P_s the 3D-to-2D projection matrices, and I_s the noisy projection images. The goal of this paper is to estimate the covariance matrix associated with the variability of the molecule. If there is a small, finite number (C) of classes, then Σ_0 has low rank ($C - 1$). This ties the heterogeneity problem to principal component analysis (PCA) [40]. If Σ_0 has eigenvectors V_1, \dots, V_p (called principal components) corresponding to eigenvalues $\lambda_1 \dots \lambda_p$, then PCA states that V_i accounts for a variance of λ_i in the data. In modern applications, the dimensionality p is often large, while \mathbf{X} typically has much fewer intrinsic degrees of freedom [11]. The heterogeneity problem is an example of such a scenario; for this problem, we demonstrate later that the top principal components can be used in conjunction with the images to reconstruct each of the C classes.

Another class of applications closely related to Problem 1.1 is missing data problems in statistics. In these problems, X_1, \dots, X_n are samples of a random vector \mathbf{X} . The statistics of this random vector must be estimated in a situation where certain entries of the samples X_s

are missing [31]. This amounts to choosing P_s to be *coordinate-selection operators*, operators which output a certain subset of the entries of a vector. An important problem in this category is PCA with missing data, which is the task of finding the top principal components when some data are missing. Closely related to this is the noisy low rank matrix completion problem [9]. In this problem, only a subset of the entries of a low rank matrix A are known (possibly with some error), and the task is to fill in the missing entries. If we let X_s be the columns of A , then the observed variables in each column are $P_s X_s + E_s$, where P_s acts on X_s by selecting a subset of its coordinates, and E_s is noise. Note that the matrix completion problem involves filling in the missing entries of X_s , while Problem 1.1 requires us only to find the covariance matrix of these columns. However, the two problems are closely related. For example, if the columns are distributed normally, then the missing entries can be found as their expectations conditioned on the known variables [51]. Alternatively, we can find the missing entries by choosing the linear combinations of the principal components that best fit the known matrix entries. A well-known application of matrix completion is in the field of recommender systems (also known as collaborative filtering). In this application, users rate the products they have consumed, and the task is to determine what new products they would rate highly. We obtain this problem by interpreting $A_{i,j}$ as the j th user's rating of product i . In recommender systems, it is assumed that only a few underlying factors determine users' preferences. Hence, the data matrix A should have low rank. A high profile example of recommender systems is the Netflix prize problem [6].

In both of these classes of problems, Σ_0 is large but should have low rank. Despite this, note that Problem 1.1 does not have a low rank assumption. Nevertheless, as our numerical results demonstrate, the spectrum of our (unregularized) covariance matrix estimator reveals low rank structure when it is present in the data. Additionally, the framework we develop in this paper naturally allows for regularization.

Having introduced Problem 1.1 and its applications, let us delve more deeply into one particular application: SPR from cryo-EM.

1.2. Cryo-electron microscopy

Electron microscopy is an important tool for structural biologists, as it allows them to determine complex 3D macromolecular structures. A general technique in electron microscopy is called SPR. In the basic setup of SPR, the data collected are 2D projection images of ideally assumed identical, but randomly oriented, copies of a macromolecule. In particular, one specimen preparation technique used in SPR is called cryo-EM, in which the sample of molecules is rapidly frozen in a thin ice layer [12, 63]. The electron microscope provides a top view of the molecules in the form of a large image called a micrograph. The projections of the individual particles can be picked out from the micrograph, resulting in a set of projection images. Mathematically, we can describe the imaging process as follows. Let $X : \mathbb{R}^3 \rightarrow \mathbb{R}$ represent the Coulomb potential induced by the unknown molecule. We scale the problem to be dimension-free in such a way that most of the "mass" of X lies within the unit ball $B \subset \mathbb{R}^3$ (since we later model X to be bandlimited, we cannot quite assume it is supported in B). To each copy of this molecule corresponds a rotation $R \in$

$SO(3)$, which describes its orientation in the ice layer. The idealized forward projection operator $P = P(R) : L^1(\mathbb{R}^3) \rightarrow L^1(\mathbb{R}^2)$ applied by the microscope is the X-ray transform

$$(PX)(x, y) = \int_{\mathbb{R}} X(R^T r) dz, \quad (1.2)$$

where $r = (x, y, z)^T$. Hence, P first rotates X by R , and then integrates along vertical lines to obtain the projection image. The microscope yields the image PX , discretized onto an $N \times N$ Cartesian grid, where each pixel is also corrupted by additive noise. Let there be $q \approx \pi N^2$ pixels contained in the inscribed disc of an $N \times N$ grid (the remaining pixels contain little or no signal because X is concentrated in B). If $S : L^1(\mathbb{R}^2) \rightarrow \mathbb{R}^q$ is a discretization operator, then the microscope produces images I given by

$$I = SPX + E \quad (1.3)$$

with $E \sim N(0, \sigma^2 I_q)$, where for the purposes of this paper we assume additive white Gaussian noise. The microscope has an additional blurring effect on the images, a phenomenon we will discuss shortly, but will leave out of our model. Given a set of images I_1, \dots, I_n , the cryo-EM problem is to estimate the orientations R_1, \dots, R_n of the underlying volumes and reconstruct X . Note that throughout this paper, we will use “cryo-EM” and “cryo-EM problem” as shorthand for the SPR problem from cryo-EM images; we also use “volume” as a synonym for “3D structure.”

The cryo-EM problem is challenging for several reasons. Unlike most other imaging modalities of computerized tomography, the rotations R_s are unknown, so we must estimate them before reconstructing X . This challenge is one of the major hurdles to reconstruction in cryo-EM. Since the images are not perfectly centered, they also contain in-plane translations, which must be estimated as well. The main challenge in rotation estimation is that the projection images are corrupted by extreme levels of noise. This problem arises because only low electron doses can scan the molecule without destroying it. To an extent, this problem is mitigated by the fact that cryo-EM datasets often have tens or even hundreds of thousands of images, which makes the reconstruction process more robust. Another issue with transmission electron microscopy in general is that technically, the detector only registers the magnitude of the electron wave exiting the specimen. Zernike realized in the 1940s that the phase information could also be recovered if the images were taken out of focus [60]. While enabling measurement of the full output of the microscope, this out-of-focus imaging technique produces images representing the convolution of the true image with a point spread function (PSF). The Fourier transform of the PSF is called the contrast transfer function (CTF). Thus the true images are multiplied by the CTF in the Fourier domain to produce the output images. Hence, the P_s operators in practice also include the blurring effect of a CTF. This results in a loss of information at the zero crossings of the (Fourier-domain) CTF and at high frequencies [12]. In order to compensate for the former effect, images are taken with several different defocus values, whose corresponding CTFs have different zero crossings.

The field of cryo-EM has recently seen a drastic improvement in detector technology. New direct electron detector cameras have been developed, which, according to a recent article in *Science*, have “unprecedented speed and sensitivity” [24]. This technology has enabled SPR from cryo-EM to succeed on smaller molecules (up to size ~150 kDa) and achieve higher resolutions (up to 3Å) than before. Such high resolution allows tracing of the polypeptide chain and identification of residues in protein molecules [28, 3, 15, 34, 68]. Recently, single particle methods have provided high resolution structures of the TRPV1 ion channel [30] and of the large subunit of the yeast mitochondrial ribosome [1]. While X-ray crystallography is still the imaging method of choice for small molecules, cryo-EM now holds the promise of reconstructing larger, biomedically relevant molecules not amenable to crystallization.

The most common method for solving the basic cryo-EM problem is guessing an initial structure and then performing an iterative refinement procedure, where iterations alternate between (1) estimating the rotations of the experimental images by matching them with projections of the current 3D model and (2) tomographic inversion producing a new 3D model based on the experimental images and their estimated rotations [12, 61, 44]. There are no convergence guarantees for this iterative scheme, and the initial guess can incur bias in the reconstruction. An alternative is to estimate the rotations and reconstruct an accurate initial structure directly from the data. Such an *ab initio* structure is a much better initialization for the iterative refinement procedure. This strategy helps avoid bias and reduce the number of refinement iterations necessary to converge [70]. In the *ab initio* framework, rotations can be estimated by one of several techniques (see, e.g., [55, 64] and references therein).

1.3. Heterogeneity problem

As presented above, a key assumption in the cryo-EM problem is that the sample consists of (rotated versions of) identical molecules. However, in many datasets this assumption does not hold. Some molecules of interest exist in more than one conformational state. For example, a subunit of the molecule might be present or absent, have a few different arrangements, or be able to move continuously from one position to another. These structural variations are of great interest to biologists, as they provide insight into the functioning of the molecule. Unfortunately, standard cryo-EM methods do not account for heterogeneous samples. New techniques must be developed to map the space of molecules in the sample, rather than just reconstruct a single volume. This task is called the *heterogeneity problem*. A common case of heterogeneity is when the molecule has a finite number of dominant conformational classes. In this discrete case, the goal is to provide biologists with 3D reconstructions of all these structural states. While cases of continuous heterogeneity are possible, in this paper we mainly focus on the discrete heterogeneity scenario.

While we do not investigate the 3D rotation estimation problem in the heterogeneous case, we conjecture that this problem can be solved without developing sophisticated new tools. Consider, for example, the case when the heterogeneity is small, i.e., the volumes X_1, \dots, X_n can be rotationally aligned so they are all close to their mean (in some norm). For example,

this property holds when the heterogeneity is localized (e.g., as in Figure 1). In this case, one might expect that by first assuming homogeneity, existing rotation estimation methods would yield accurate results. Even if the heterogeneity is large, an iterative scheme can be devised to alternately estimate the rotations and conformations until convergence (though this convergence is local, at best). Thus, in this publication, we assume that the 3D rotations R_s (and in-plane translations) have already been estimated.

With the discrete heterogeneity and known rotations assumptions, we can formulate the heterogeneity problem as follows.

Problem 1.2 (heterogeneity problem)—Suppose a heterogeneous molecule can take on one of C different states: $X^1, \dots, X^C \in B$, where B is a finite-dimensional space of bandlimited functions (see section 3.2). Let $\Omega = \{1, 2, \dots, C\}$ be a sample space, and p_1, \dots, p_C probabilities (summing to one) so that the molecule assumes state c with probability p_c . Represent the molecule as a random field $\mathbf{X} : \Omega \times \mathbb{R}^3 \rightarrow \mathbb{R}$, with

$$\mathbb{P}[X = X^c] = p_c, \quad c = 1, \dots, C. \quad (1.4)$$

Let \mathbf{R} be a random rotation with some distribution over $SO(3)$, and define the corresponding random projection $\mathbf{P} = \mathbf{P}(\mathbf{R})$ (see (1.2)). Finally, $\mathbf{E} \sim N(0, \sigma^2 \mathbf{I}_q)$. Assume that \mathbf{X} , \mathbf{R} , \mathbf{E} are independent. A random image of a particle is obtained via

$$\mathbf{I} = \mathbf{S}\mathbf{P}\mathbf{X} + \mathbf{E}, \quad (1.5)$$

where $\mathbf{S} : L^1(\mathbb{R}^2) \rightarrow \mathbb{R}^q$ is a discretization operator. Given observations $\mathbf{I}^1, \dots, \mathbf{I}_n$ and $\mathbf{R}_1, \dots, \mathbf{R}_n$ of \mathbf{I} and \mathbf{R} , respectively, estimate the number of classes C , the structures X^c , and the probabilities p_c .

Note that $\mathbf{S}\mathbf{P}|_B$ is a (random) linear operator between finite-dimensional spaces, and so it has a matrix version $\mathbf{P} : \mathbb{R}^p \rightarrow \mathbb{R}^q$, where $p = \dim B$. If we let \mathbf{X} be the random vector on \mathbb{R}^p obtained by expanding \mathbf{X} in the basis for B , then we recover the equation $\mathbf{I} = \mathbf{P}\mathbf{X} + \mathbf{E}$ from Problem 1.1. Thus, the main factors distinguishing Problem 1.2 from Problem 1.1 are that the former assumes a specific form for \mathbf{P} and posits a discrete distribution on \mathbf{X} . As we discuss in section 4, Problem 1.2 can be solved by first estimating the covariance matrix as in Problem 1.1, finding coordinates for each image with respect to the top eigenvectors of this matrix, and then applying a standard clustering procedure to these coordinates.

One of the main difficulties of the heterogeneity problem is that, compared to usual SPR, we must deal with an even lower effective signal-to-noise ratio (SNR). Indeed, the signal we seek to reconstruct is the variation of the molecules around their mean, as opposed to the mean volume itself. We propose a precise definition of SNR in the context of the heterogeneity problem in section 7.1. Another difficulty is the indirect nature of our problem. Although the heterogeneity problem is an instance of a clustering problem, it differs from usual such problems in that we do not have access to the objects we are trying to cluster—

only projections of these objects onto a lower-dimensional space are available. This makes it challenging to apply any standard clustering technique directly.

The heterogeneity problem is considered one of the most important problems in cryo-EM. In his 2013 Solvay public lecture on cryo-EM, Dr. Joachim Frank emphasized the importance of “the ability to obtain an entire inventory of coexisting states of a macromolecule from a single sample” [13]. Speaking of approaches to the heterogeneity problem in a review article, Frank discussed “the potential these new technologies will have in exploring functionally relevant states of molecular machines” [14]. It is stressed there that much room for improvement remains; current methods cannot automatically identify the number of conformational states and have trouble distinguishing between similar conformations.

1.4. Previous work

Much work related to Problems 1.1 and 1.2 has already been done. There is a rich statistical literature on the covariance estimation problem in the presence of missing data, a special case of Problem 1.1. In addition, work on the *low rank matrix sensing problem* (a generalization of matrix completion) is also closely related to Problem 1.1. Regarding Problem 1.2, several approaches to the heterogeneity problem have been proposed in the cryo-EM literature.

1.4.1. Work related to Problem 1.1—Many approaches to covariance matrix estimation from missing data have been proposed in the statistics literature [31]. The simplest approach to dealing with missing data is to ignore the samples with any unobserved variables. Another simple approach is called *available case analysis*, in which the statistics are constructed using all the available values. For example, the (i, j) entry of the covariance matrix is constructed using all samples for which the i th and j th coordinates are simultaneously observed. These techniques work best under certain assumptions on the pattern of missing entries, and more sophisticated techniques are preferred [31]. One of the most established such approaches is *maximum likelihood estimation* (MLE). This involves positing a probability distribution on \mathbf{X} (e.g., multivariate normal) and then maximizing the likelihood of the observed partial data with respect to the parameters of the model. Such an approach to fitting models from partial observations was known as early as the 1930s, when Wilks used it for the case of a bivariate normal distribution [66]. Wilks proposed to maximize the likelihood using a gradient-based optimization approach. In 1977, Dempster, Laird, and Rubin introduced the expectation-maximization (EM) algorithm [10] to solve maximum likelihood problems. The EM algorithm is one of the most popular methods for solving missing data problems in statistics. Also, there is a class of approaches to missing data problems called *imputation*, in which the missing values are filled either by averaging the available values or through more sophisticated regression-based techniques. Finally, see [32, 33] for other approaches to related problems.

Closely related to covariance estimation from missing data is the problem of PCA with missing data. In this problem, the task is to find the leading principal components, and not necessarily the entire covariance matrix. Not surprisingly, EM-type algorithms are popular for this problem as well. These algorithms often search directly for the low rank factors. See

[18] for a survey of approaches to PCA with missing data. Closely related to PCA with missing data is the low rank matrix completion problem. Many of the statistical methods discussed above are also applicable to matrix completion. In particular, EM algorithms to solve this problem are popular, e.g., [51, 27].

Another more general problem setup related to Problem 1.1 is the *low rank matrix sensing problem*, which generalizes the low rank matrix completion problem. Let $A \in \mathbb{R}^{p \times n}$ be an unknown rank- k matrix, and let $M : \mathbb{R}^{p \times n} \rightarrow \mathbb{R}^d$ be a linear map, called the *sensing matrix*. We would like to find A , but we only have access to the (possibly noisy) data $M(A)$. Hence, the low rank matrix sensing problem can be formulated as follows [19]:

$$\text{minimize } \|M(A) - b\| \quad \text{s.t.} \quad \text{rank}(A) \leq k. \quad (1.6)$$

Note that when Σ_0 is low rank, Problem 1.1 is a special case of the low rank matrix sensing problem. Indeed, consider putting the unknown vectors X_1, \dots, X_n together as the columns of a matrix A . The rank of this matrix is the number of degrees of freedom in \mathbf{X} (in the cryo-EM problem, this relates to the number of heterogeneity classes of the molecule). The linear projections P_1, \dots, P_n can be combined into one sensing matrix M acting on A . In this way, our problem falls into the realm of matrix sensing.

One of the first algorithms for matrix sensing was inspired by the compressed sensing theory [46]. This approach uses a matrix version of l_1 regularization called nuclear norm regularization. The nuclear norm is the sum of the singular values of a matrix, and is a convex proxy for its rank. Another approach to this problem is alternating minimization, which decomposes A into a product of the form UV^T and iteratively alternates between optimizing with respect to U and V . The first proof of convergence for this approach was given in [19]. Both the nuclear norm and alternating minimization approaches to the low rank matrix sensing problem require a *restricted isometry property* on M for theoretical guarantees.

While the aforementioned algorithms are widely used, we believe they have limitations as well. EM algorithms require postulating a distribution over the data and are susceptible to getting trapped in local optima. Regarding the former point, Problem 1.1 avoids any assumptions on the distribution of \mathbf{X} , so our estimator should have the same property. Matrix sensing algorithms (especially alternating minimization) often assume that the rank is known in advance. However, there is no satisfactory statistical theory for choosing the rank. By contrast, the estimator we propose for Problem 1.1 allows automatic rank estimation.

1.4.2. Work related to Problem 1.2—Several approaches to the heterogeneity problem have been proposed. Here we give a brief overview of some of these approaches.

One approach is based on the notion of common lines. By the Fourier projection slice theorem (see Theorem 3.1), the Fourier transforms of any two projection images of an object will coincide on a line through the origin, called a common line. The idea of Shatsky et al. [52] was to use common lines as a measure of how likely it is that two projection images correspond to the same conformational class. Specifically, given two projection images and

their corresponding rotations, we can take their Fourier transforms and correlate them on their common line. From there, a weighted graph of the images is constructed, with edges weighted based on this common line measure. Then spectral clustering is applied to this graph to classify the images. An earlier common lines approach to the heterogeneity problem is described in [16].

Another approach is based on MLE. It involves positing a probability distribution over the space of underlying volumes, and then maximizing the likelihood of the images with respect to the parameters of the distribution. For example, Wang et al. [65] model the heterogeneous molecules as a mixture of Gaussians and employ the EM algorithm to find the parameters. A challenge with MLE approaches is that the resulting objective functions are nonconvex and have a complicated structure. For more discussion of the theory and practice of maximum likelihood methods, see [53] and [50], respectively. Also see [49] for a description of a software package which uses maximum likelihood to solve the heterogeneity problem.

A third approach to the heterogeneity problem is to use the covariance matrix of the set of original molecules. Penczek, Kimmel, and Spahn outline a bootstrapping approach in [43] (see also [41, 42, 67, 29]). In this approach, one repeatedly takes random subsets of the projection images and reconstructs 3D volumes from these samples. Then, one can perform PCA on this set of reconstructed volumes, which yields a few dominant “eigenvolumes.” Penczek, Kimmel, and Spahn propose to then produce mean-subtracted images by subtracting projections of the mean volume from the images. The next step is to project each of the dominant eigenvolumes in the directions of the images, and then obtain a set of coordinates for each image based on its similarity with each of the eigenvolume projections. Finally, using these coordinates, this resampling approach proceeds by applying a standard clustering algorithm such as K -means to classify the images into classes.

While existing methods for the heterogeneity problem have their success stories, each suffers from its own shortcomings: the common line approach does not exploit all the available information in the images, the maximum likelihood approach requires explicit a priori distributions and is susceptible to local optima, and the bootstrapping approach based on covariance matrix estimation is a heuristic sampling method that lacks in theoretical guarantees.

Note that the above overview of the literature on the heterogeneity problem is not comprehensive. For example, very recently, an approach to the heterogeneity problem based on normal mode analysis was proposed [20].

1.5. Our contribution

In this paper, we propose and analyze a covariance matrix estimator Σ_n to solve the general statistical problem (Problem 1.1), and then apply this estimator to the heterogeneity problem (Problem 1.2).

Our covariance matrix estimator has several desirable properties. First, we prove that the estimator is consistent as $n \rightarrow \infty$ for fixed p, q . Second, our estimator does not require a prior distribution on the data, unlike MLE methods. Third, when the data have low intrinsic

dimension, our method does not require knowing the rank of Σ_0 in advance. The rank can be estimated from the spectrum of the estimated covariance matrix. This sets our method apart from alternating minimization algorithms that search for the low rank matrix factors themselves. Fourth, our estimator is given in closed form and its computation requires only a single linear inversion.

To implement our covariance matrix estimator in the cryo-EM case, we must invert a high-dimensional matrix L_n (see definition (2.8)). The size of this matrix is so large that typically it cannot even be stored on a computer; thus, inverting L_n is the greatest practical challenge we face. We consider two possibilities of addressing this challenge. In the primary approach we consider, we replace L_n by its limiting operator L , which does not depend on the rotations R_s and is a good approximation of L_n as long as these rotations are distributed uniformly enough. We then carefully construct new bases for images and volumes to make L a sparse, block diagonal matrix. While L has dimensions on the order of $N_{res}^6 \times N_{res}^6$, this matrix has only $O(N_{res}^9)$ total nonzero entries in the bases we construct, where N_{res} is the grid size corresponding to the target resolution. These innovations lead to a practical algorithm to estimate the covariance matrix in the heterogeneity problem. The second approach we consider is an iterative inversion of L_n , which has a low storage requirement and avoids the requirement of uniformly distributed rotations. We compare the complexities of these two methods, and find that each has its strengths and weaknesses.

The limiting operator L is a fundamental object in tomographic problems involving variability, and we call it the *projection covariance transform*. The projection covariance transform relates the covariance matrix of the imaged object to data that can be acquired from the projection images. Standard weighted back-projection tomographic reconstruction algorithms involve application of the ramp filter to the data [38], and we find that the inversion of L entails applying a similar filter, which we call the *triangular area filter*. The triangular area filter has many of the same properties as the ramp filter, but reflects the slightly more intricate geometry of the covariance estimation problem. The projection covariance transform is an interesting mathematical object in its own right, and we begin studying it in this paper.

Finally, we numerically validate the proposed algorithm (the first algorithm discussed above). We demonstrate this method's robustness to noise on synthetic datasets by obtaining a meaningful reconstruction of the covariance matrix and molecular volumes even at low SNR levels. Excluding precomputations (which can be done once and for all), reconstructions for 10000 projection images of size 65×65 pixels takes fewer than five minutes on a standard laptop computer.

The paper is organized as follows. In section 2, we construct an estimator for Problem 1.1, state theoretical results about this estimator, and connect our problem to high-dimensional PCA. In section 3, we specialize the covariance estimator to the heterogeneity problem and investigate its geometry. In section 4, we discuss how to reconstruct the conformations once we have estimated the mean and covariance matrix. In section 5, we discuss computational aspects of the problem and construct a basis in which L is block diagonal and sparse. In

section 6, we explore the complexity of the proposed approach. In section 7, we present numerical results for the heterogeneity problem. We conclude with a discussion of future research directions in section 8. Appendices A, B, and C contain calculations and proofs.

2. An estimator for Problem 1.1

2.1. Constructing an estimator

We define estimators μ_n and Σ_n through a general optimization framework based on the model (1.1). As a first step, let us calculate the first- and second-order statistics of \mathbf{I} , conditioned on the observed matrix P_s for each s . Using the assumptions in Problem 1.1, we find that

$$\mathbb{E}[\mathbf{I}|P=P_s] = \mathbb{E}[PX+E|P=P_s] = \mathbb{E}[P|P=P_s] \mathbb{E}[X] = P_s \mu_0 \quad (2.1)$$

and

$$\text{Var}[\mathbf{I}|P=P_s] = \text{Var}[PX|P=P_s] + \text{Var}[E] = P_s \Sigma_0 P_s^H + \sigma^2 \mathbf{I}_p. \quad (2.2)$$

Note that P_s^H denotes the conjugate transpose of P_s .

Based on (2.1) and (2.2), we devise least-squares optimization problems for μ_n and Σ_n :

$$\mu_n = \underset{\mu}{\text{argmin}} \frac{1}{n} \sum_{s=1}^n \|\mathbf{I}_s - P_s \mu\|^2, \quad (2.3)$$

$$\Sigma_n = \underset{\Sigma}{\text{argmin}} \frac{1}{n} \sum_{s=1}^n \left\| (\mathbf{I}_s - P_s \mu_n) (\mathbf{I}_s - P_s \mu_n)^H - (P_s \Sigma P_s^H + \sigma^2 \mathbf{I}_p) \right\|_F^2. \quad (2.4)$$

Here we use the Frobenius norm, which is defined by $\|A\|_F^2 = \sum_{i,j} |A_{i,j}|^2$

Note that these optimization problems do not encode any prior knowledge about μ_0 or Σ_0 . Since Σ_0 is a covariance matrix, it must be positive semidefinite (PSD). As discussed above, in many applications Σ_0 is also low rank. The estimator Σ_n need not satisfy either of these properties. Thus, regularization of (2.4) is an option worth exploring. Nevertheless, here we only consider the unregularized estimator Σ_n . Note that in most practical problems, we only are interested in the leading eigenvectors of Σ_n , and if these are estimated accurately, then it does not matter if Σ_n is not PSD or low rank. Our numerical experiments show that in practice, the top eigenvectors of Σ_n are indeed good estimates of the true principal components for high enough SNR.

Note that we first solve (2.3) for μ_n , and then use this result in (2.4). This makes these optimization problems quadratic in the elements of μ and Σ , and hence they can be solved by setting the derivatives with respect to μ and Σ to zero. This leads to the following equations for μ_n and Σ_n (see Appendix A for the derivative calculations):

$$\frac{1}{n} \left(\sum_{s=1}^n P_s^H P_s \right) \mu_n = \frac{1}{n} \sum_{s=1}^n P_s^H I_s =: b_n, \quad (2.5)$$

$$\frac{1}{n} \sum_{s=1}^n P_s^H P_s \Sigma_n P_s^H P_s = \frac{1}{n} \sum_{s=1}^n P_s^H (I_s - P_s \mu_n) (I_s - P_s \mu_n)^H P_s - \sigma^2 \frac{1}{n} \sum_{s=1}^n P_s^H P_s =: B_n. \quad (2.6)$$

When $p = q$ and $\mathbf{P} = \mathbf{I}_p$, μ_n and Σ_n reduce to the sample mean and sample covariance matrix. When \mathbf{P} is a coordinate-selection operator (recall the discussion following the statement of Problem 1.1), (2.5) estimates the mean by averaging all the available observations for each coordinate, and (2.6) estimates each entry of the covariance matrix by averaging over all samples for which both coordinates are observed. These are exactly the available-case estimators discussed in [31, section 3.4].

Observe that (2.5) requires inversion of the matrix

$$A_n = \frac{1}{n} \sum_{s=1}^n P_s^H P_s, \quad (2.7)$$

and (2.6) requires inversion of the linear operator $L_n : \mathbb{C}^{p \times p} \rightarrow \mathbb{C}^{p \times p}$ defined by

$$L_n(\Sigma) = \frac{1}{n} \sum_{s=1}^n P_s^H P_s \Sigma P_s^H P_s. \quad (2.8)$$

Since the P_s are drawn independently from \mathbf{P} , the law of large numbers implies that (2.9) $A_n \rightarrow A$ and $L_n \rightarrow L$ almost surely,

$$A_n \rightarrow A \quad \text{and} \quad L_n \rightarrow L \quad \text{almost surely}, \quad (2.9)$$

where the convergence is in the operator norm, and

$$A = \mathbb{E} \left[P^H P \right] \quad \text{and} \quad L(\Sigma) = \mathbb{E} \left[P^H P \Sigma P^H P \right]. \quad (2.10)$$

The invertibilities of A and L depend on the distribution of \mathbf{P} . Intuitively, if \mathbf{P} has a nonzero probability of “selecting” any coordinate of its argument, then A will be invertible. If \mathbf{P} has a nonzero probability of “selecting” any pair of coordinates of its argument, then L will be invertible. In this paper, we assume that A and L are invertible. In particular, we will find that in the cryo-EM case, A and L are invertible if, for example, the rotations are sampled uniformly from $SO(3)$. Under this assumption, we will prove that A_n and L_n are invertible with high probability for sufficiently large n . In the case when A_n or L_n are not invertible, we cannot define estimators from the above equations, so we simply set them to zero. Since the RHS quantities b_n and B_n are noisy, it is also not desirable to invert A_n or L_n when these matrices are nearly singular. Hence, we propose the following estimators:

$$\mu_n = \begin{cases} A_n^{-1} b_n & \text{if } \|A_n^{-1}\| \leq 2\|A^{-1}\|, \\ 0 & \text{otherwise;} \end{cases} \quad \Sigma_n = \begin{cases} L_n^{-1}(B_n) & \text{if } \|L_n^{-1}\| \leq 2\|L^{-1}\|, \\ 0 & \text{otherwise} \end{cases} \quad (2.11)$$

The factors of 2 are somewhat arbitrary; any $a > 1$ would do.

Let us make a few observations about A_n and L_n . By inspection, A_n is symmetric and PSD. We claim that L_n satisfies the same properties, with respect to the Hilbert space $\mathbb{C}^{p \times p}$ equipped with the inner product $\langle A, B \rangle = \text{tr}(B^H A)$. Using the property $\text{tr}(AB) = \text{tr}(BA)$, we find that for any Σ_1, Σ_2 ,

$$\begin{aligned} \langle L_n(\Sigma_1), \Sigma_2 \rangle &= \text{tr}(\Sigma_2^H L_n(\Sigma_1)) = \text{tr}\left[\frac{1}{n} \sum_s \Sigma_2^H P_s^H P_s \Sigma_1 P_s^H P_s\right] \\ &= \text{tr}\left[\frac{1}{n} \sum_s P_s^H P_s \Sigma_2^H P_s^H P_s \Sigma_1\right] = \langle \Sigma_1, L(\Sigma_2) \rangle. \end{aligned} \quad (2.12)$$

Thus, L_n is self-adjoint. Next, we claim that L_n is PSD. Indeed,

$$\begin{aligned} \langle L_n(\Sigma), \Sigma \rangle &= \text{tr}(\Sigma^H L_n(\Sigma)) = \text{tr}\left[\frac{1}{n} \sum_s \Sigma^H P_s^H P_s \Sigma P_s^H P_s\right] \\ &= \frac{1}{n} \sum_s \text{tr}\left[\left(P_s \Sigma P_s^H\right)^H \left(P_s \Sigma P_s^H\right)\right] = \sum_s \frac{1}{n} \|P_s \Sigma P_s^H\|_F^2 \geq 0. \end{aligned} \quad (2.13)$$

2.2. Consistency of μ_n and Σ_n

In this section, we state that under mild conditions on $\mathbf{P}, \mathbf{X}, \mathbf{E}$, the estimators μ_n and Σ_n are consistent. Note that here, and throughout this paper, $\|\cdot\|$ will denote the Euclidean norm for vectors and the operator norm for matrices. Also, define

$$\|Y\|_j = \mathbb{E}\left[\|Y - \mathbb{E}[Y]\|^j\right]^{1/j}, \quad (2.14)$$

where Y is a random vector.

Proposition 2.1—*Suppose A (defined in (2.10)) is invertible, that \mathbf{P} is bounded almost surely, and that $\|\mathbf{X}\|_2, \|\mathbf{E}\|_2 < \infty$. Then, for fixed p, q we have*

$$\mathbb{E}\|\mu_n - \mu_0\| = O\left(\frac{1}{\sqrt{n}}\right). \quad (2.15)$$

Hence, under these assumptions, μ_n is consistent.

Proposition 2.2—*Suppose A and L (defined in (2.10)) are invertible, that \mathbf{P} is bounded almost surely, and that there is a polynomial Q for which*

$$\|X\|_j, \|E\|_j \leq Q(j), \quad j \in \mathbb{N}. \quad (2.16)$$

Then, for fixed p, q , we have

$$\mathbb{E}\|\Sigma_n - \Sigma_0\| = O\left(\frac{Q(\log n)}{\sqrt{n}}\right). \quad (2.17)$$

Hence, under these assumptions, Σ_n is consistent.

Remark 2.3—The moment growth condition (2.16) on X and E is not very restrictive. For example, bounded, subgaussian, and subexponential random vectors all satisfy (2.16) with $\deg Q = 1$ (see [62, sections 5.2 and 5.3]).

See Appendix B for the proofs of Propositions (2.1) and (2.2). We mentioned that μ_n and Σ_n are generalizations of available-case estimators. Such estimators are known to be consistent when the data are *missing completely at random* (MCAR). This means that the pattern of missingness is independent of the (observed and unobserved) data. Accordingly, in Problem 1.1, we assume that \mathbf{P} and \mathbf{X} are independent, a generalization of the MCAR condition. The above propositions state that the consistency of μ_n and Σ_n also generalizes to Problem 1.1.

2.3. Connection to high-dimensional PCA

While the previous section focused on the “fixed p , large n ” regime, in practice both p and n are large. Now, we consider the latter regime, which is common in modern high-dimensional statistics. In this regime, we consider the properties of the estimator Σ_n when Σ_0 is low rank, and the task is to find its leading eigenvectors. What is the relationship between the spectra of Σ_n and Σ_0 ? Can the rank of Σ_0 be deduced from that of Σ_n ? To what extent do the leading eigenvectors of Σ_n approximate those of Σ_0 ? In the setting of (1.1) when $\mathbf{P} = \mathbf{I}_p$, the theory of high-dimensional PCA provides insight into such properties of the sample covariance matrix (and thus of Σ_n). In particular, an existing result gives the correlation between the top eigenvectors of Σ_n and Σ_0 for given settings of SNR and p/n . It follows from this result that if the SNR is sufficiently high compared to p/n , then the top eigenvector of Σ_n is a useful approximation of the top eigenvector of Σ_0 . If generalized to the case of nontrivial \mathbf{P} , this result would be a useful guide for using the estimator Σ_n to solve practical problems, such as Problem 1.2. In this section, we first discuss the existing high-dimensional PCA literature, and then raise some open questions about how these results generalize to the case of nontrivial \mathbf{P} .

Given independently and identically distributed (i.i.d.) samples $I_1, \dots, I_n \in \mathbb{R}^p$ from a centered distribution \mathbf{I} with covariance matrix $\tilde{\Sigma}_0$ (called the population covariance matrix), the sample covariance matrix $\tilde{\Sigma}_n$ is defined by

$$\tilde{\Sigma}_n = \frac{1}{n} \sum_{s=1}^n I_s I_s^H. \quad (2.18)$$

We use the new tilde notation because in the context of Problem 1.1, $\tilde{\Sigma}_0$ is the signal-plus-noise covariance matrix, as opposed to the covariance of the signal itself. High-dimensional PCA is the study of the spectrum of $\tilde{\Sigma}_n$ for various distributions of \mathbf{I} in the regime where $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma$.

The first case to consider is $\mathbf{X} = 0$, i.e., $\mathbf{I} = \mathbf{E}$, where $\mathbf{E} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_p)$. In a landmark paper, Marcčenko and Pastur [35] proved that the spectrum of $\tilde{\Sigma}_n$ converges to the *Marcčenko–Pastur (MP) distribution*, which is parameterized by γ and σ^2 :

$$MP(x) = \frac{1}{2\pi\sigma^2} \frac{\sqrt{(\gamma_+ - x)(x - \gamma_-)}}{\gamma x} 1_{[\gamma_-, \gamma_+]}, \quad \gamma_{\pm} = \sigma^2(1 \pm \sqrt{\gamma})^2. \quad (2.19)$$

The above formula assumes $\gamma \leq 1$; a similar formula governs the case $\gamma > 1$. Note that there are much more general statements about classes of \mathbf{I} for which this convergence holds; see, e.g., [54]. See Figure 2(a) for MP distributions with a few different parameter settings.

Johnstone [21] took this analysis a step further and considered the limiting distribution of the largest eigenvalue of $\tilde{\Sigma}_n$. He showed that the distribution of this eigenvalue converges to the Tracy–Widom distribution centered on the right edge of the MP spectrum. In the same paper, Johnstone considered the *spiked covariance model*, in which

$$\mathbf{I} = \mathbf{X} + \mathbf{E}, \quad (2.20)$$

where \mathbf{E} is as before and $\Sigma_0 = \text{Var}[\mathbf{X}] = \text{diag}(\tau_1^2, \dots, \tau_r^2, 0)$, so that the population covariance matrix is $\tilde{\Sigma}_0 = \text{diag}(\tau_1^2 + \sigma^2, \dots, \tau_r^2 + \sigma^2, \sigma^2, \dots, \sigma^2)$. Here, \mathbf{X} is the signal and \mathbf{E} is the noise. In this view, the goal is to accurately recover the top r eigenvectors, as these will determine the subspace on which \mathbf{X} is supported. The question then is the following: for what values of τ_1, \dots, τ_r will the top r eigenvectors of the sample covariance matrix be good approximations to the top eigenvectors of the population covariance? Since we might not know the value of r a priori, it is important to first determine for what values of τ_1, \dots, τ_r we can detect the presence of “spiked” population eigenvalues. In [5], the spectrum of the sample covariance matrix in the spiked model was investigated. It was found that the bulk of the distribution still obeys the MP law, whereas for each k such that

$$\frac{\tau_k^2}{\sigma^2} \geq \sqrt{\gamma}, \quad (2.21)$$

the sample covariance matrix will have an eigenvalue tending to $(\tau_k^2 + \sigma^2) \left(1 + \frac{\sigma^2}{\tau_k^2} \gamma\right)$. The signal eigenvalues below this threshold tend to the right edge of the noise distribution. Thus, (2.21) defines a criterion for detection of signal. In Figure 2(b), we illustrate these results with a numerical example. We choose $p = 800$, $n = 4000$, and a spectrum corresponding to $r = 3$, with τ_1, τ_2 above, but τ_3 below, the threshold corresponding to $\gamma = p/n = 0.2$. Figure

2(b) is a normalized histogram of the eigenvalues of the sample covariance matrix. The predicted MP distribution for the bulk is superimposed. We see that indeed we have two eigenvalues separated from this bulk. Moreover, the eigenvalue of $\tilde{\Sigma}_n$ corresponding to τ_3 does not pop out of the noise distribution.

It is also important to compare the top eigenvectors of the sample and population covariance matrices. Considering the simpler case of a spiked model with $r = 1$, [4, 37] showed a “phase transition” effect: as long as τ_1 is above the threshold in (2.21), the correlation of the top eigenvector (V_{PCA}) with the true principal component (V) tends to a limit between 0 and 1:

$$|\langle V_{PCA}, V \rangle|^2 \rightarrow \frac{\frac{1}{\gamma} \frac{\tau_1^4}{\sigma^4} - 1}{\frac{1}{\gamma} \frac{\tau_1^4}{\sigma^4} + \frac{\tau_1^2}{\sigma^2}}. \quad (2.22)$$

Otherwise, the limiting correlation is zero. Thus, high-dimensional PCA is inconsistent.

However, if τ_1^2/σ^2 is sufficiently high compared to $\sqrt{\gamma}$, then the top eigenvector of the sample covariance matrix is still a useful approximation.

While all the statements made so far have concerned the limiting case $n, p \rightarrow \infty$, similar (but slightly more complicated) statements hold for finite n, p as well (see, e.g., [37]). Thus, (2.21) has a practical interpretation. Again considering the case $r = 1$, note that the quantity τ_1^2/σ^2 is the SNR. When faced with a problem of the form (2.20) with a given p and SNR, one can determine how many samples one needs in order to detect the signal. If V represents a spatial object as in the cryo-EM case, then p can reflect the resolution to which we reconstruct V . Hence, if we have a dataset with a certain number of images n and a certain estimated SNR, then (2.21) determines the resolution to which V can be reconstructed from the data.

This information is important to practitioners (e.g., in cryo-EM), but as of now, the above theoretical results only apply to the case when \mathbf{P} is trivial. Of course, moving to the case of more general \mathbf{P} brings additional theoretical challenges. For example, with nontrivial \mathbf{P} , the empirical covariance matrix of \mathbf{X} is harder to disentangle from that of \mathbf{I} , because the operator L_n becomes nontrivial (see (2.6) and (2.8)). How can our knowledge about the spiked model be generalized to the setting of Problem 1.1? We raise some open questions along these lines.

1. In what high-dimensional parameter regimes (in terms of n, p, q) is there hope to detect and recover any signal from Σ_n ? With the addition of the parameter q , the traditional regime $p \approx n$ might no longer be appropriate. For example, in the random coordinate-selection case with the (extreme) parameter setting $q = 2$, it is expected that $n = p^2 \log p$ samples are needed just for L_n to be invertible (by the coupon collector problem).
2. In the case when there is no signal ($\mathbf{X} = 0$), we have $\mathbf{I} = \mathbf{E}$. In this case, what is the limiting eigenvalue distribution of Σ_n (in an appropriate parameter regime)? Is it

still the MP law? How does the eigenvalue distribution depend on the distribution of P ? This is perhaps the first step towards studying the signal-plus-noise model.

3. In the no-signal case, what is the limiting distribution of the largest eigenvalue of Σ_n ? Is it still Tracy–Widom? How does this depend on n, p, q , and P ? Knowing this distribution can provide p-values for signal detection, as is the case for the usual spiked model (see [21, p. 303]).
4. In the full model (1.1), if X takes values in a low-dimensional subspace of \mathbb{R}^p , is the limiting eigenvalue distribution of Σ_n a bulk distribution with a few separated eigenvalues? If so, what is the generalization of the SNR condition (2.21) that would guarantee separation of the top eigenvalues? What would these top eigenvalues be, in terms of the population eigenvalues? Would there still be a phase-transition phenomenon in which the top eigenvectors of Σ_n are correlated with the principal components as long as the corresponding eigenvalues are above a threshold?

Answering these questions theoretically would require tools from random matrix theory such as the ones used by [21, 5, 37]. We do not attempt to address these issues in this paper, but remark that such results would be very useful theoretical guides for practical applications of our estimator Σ_n . Our numerical results show that the spectrum of the cryo-EM estimator Σ_n has qualitative behavior similar to that of the sample covariance matrix.

At this point, we have concluded the part of our paper focused on the general properties of the estimator Σ_n . Next, we move on to the cryo-EM heterogeneity problem.

3. Covariance estimation in cryo-EM heterogeneity problem

Now that we have examined the general covariance matrix estimation problem, let us specialize to the cryo-EM case. In this case, the matrices P have a specific form: they are finite-dimensional versions of P (defined in (1.2)). We begin by describing the Fourier-domain counterpart of P , which will be crucial in analyzing the cryo-EM covariance estimation problem. Our Fourier transform convention is

$$\hat{f}(\xi) = \int_{\mathbb{R}^d} f(x) e^{-ix \cdot \xi} dx, \quad f(x) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{f}(\xi) e^{ix \cdot \xi} d\xi. \quad (3.1)$$

The following classical theorem in tomography (see, e.g., [38] for a proof) shows that the operator P takes on a simpler form in the Fourier domain.

Theorem 3.1 (Fourier projection slice theorem)

Suppose $Y \in L^2(\mathbb{R}^3) \cap L^1(\mathbb{R}^3)$ and $J : \mathbb{R}^2 \rightarrow \mathbb{R}$. Then

$$PY = \hat{J} \iff \hat{P}\hat{Y} = \hat{J}, \quad (3.2)$$

where $P : C(\mathbb{R}^3) \rightarrow C(\mathbb{R}^2)$ is defined by

$$\left(\hat{P}\hat{Y}\right)\begin{pmatrix}\hat{x} \\ \hat{y}\end{pmatrix}=\hat{Y}\left(R^T(\hat{x},\hat{y},0)^T\right)=\hat{y}\left(\hat{x}R^1+\hat{y}R^2\right). \quad (3.3)$$

Here, R^i is the i th row of R .

Hence, p rotates a function by R and then restricts it to the horizontal plane $\hat{z}=0$. If we let $\xi=(x,\hat{y},\hat{z})$, then another way of viewing p is that it restricts a function to the plane $\xi \cdot R^3=0$.

3.1. Infinite-dimensional heterogeneity problem

To build intuition for the Fourier-domain geometry of the heterogeneity problem, consider the following idealized scenario, taking place in Fourier space. Suppose detector technology improves to the point that images can be measured continuously and noiselessly and that we have access to the full joint distribution of \mathbf{R} and \hat{I} . We would like to estimate the mean $m_{\hat{0}}: \mathbb{R}^3 \rightarrow \mathbb{C}$ and covariance function $\hat{C}_0: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}$ of the random field \mathbf{X} , defined by

$$\hat{m}_0(\xi)=\mathbb{E}\left[\hat{X}(\xi)\right], \quad \hat{C}_0(\xi_1,\xi_2)=\mathbb{E}\left[\left(\hat{X}(\xi_1)-\hat{m}_0(\xi_1)\right)\overline{\left(\hat{X}(\xi_2)-\hat{m}_0(\xi_2)\right)}\right]. \quad (3.4)$$

Heuristically, we can proceed as follows. By the Fourier projection slice theorem, every image I provides an observation of $\mathbf{X}(\xi)$ for $\xi \in \mathbb{R}^3$ belonging to a central plane perpendicular to the viewing direction corresponding to P . By abuse of notation, let $\xi \in p$ if p carries the value of $P(\xi)$, and let $P(\xi)$ denote this value. Informally, we expect that we can recover $m_{\hat{0}}$ and \hat{C}_0

$$\hat{m}_0(\xi)=\mathbb{E}\left[\hat{I}(\xi)|\xi \in \hat{I}\right], \quad \hat{C}_0(\xi_1,\xi_2)=\mathbb{E}\left[\left(\hat{I}(\xi_1)-\hat{m}_0(\xi_1)\right)\overline{\left(\hat{I}(\xi_2)-\hat{m}_0(\xi_2)\right)}|\xi_1,\xi_2 \in \hat{I}\right]. \quad (3.5)$$

Now, let us formalize this problem setup and intuitive formulas for $m_{\hat{0}}$ and \hat{C}_0 .

Problem 3.2—Let $\hat{X}:\Omega \times \mathbb{R}^3 \rightarrow \mathbb{C}$ be a random field, where $(\Omega, \mathcal{F}, \nu)$ is a probability space. Here $\mathbf{X}(\omega, \cdot)$ is a Fourier volume for each $\omega \in \Omega$. Let $\mathbf{R}:\Omega \rightarrow SO(3)$ be a random rotation, independent of P , having the uniform distribution over $SO(3)$. Let $P=P(\mathbf{R})$ be the (random) projection operator associated with \mathbf{R} via (3.3). Define the random field $\mathbf{I}:\Omega \times \mathbb{R}^2 \rightarrow \mathbb{C}$ by

$$\hat{I}=\hat{P}\hat{X}. \quad (3.6)$$

Given the joint distribution of \mathbf{I} and \mathbf{R} , find the mean $m_{\hat{0}}$ and covariance function \hat{C}_0 of P , defined in (3.4). Let X be regular enough that

$$\hat{m}_0 \in C_0^\infty(\mathbb{R}^3), \quad \hat{C}_0 \in C_0^\infty(\mathbb{R}^3 \times \mathbb{R}^3). \quad (3.7)$$

In this problem statement, we do not assume that X has a discrete distribution. The calculations that follow hold for any \hat{I} satisfying (3.7).

We claim that $m_{\hat{0}}$ and \hat{C} can be found by solving

$$\hat{A}(\hat{m}_0) := \mathbb{E} [\hat{P}^* \hat{P}] \hat{m}_0 = \mathbb{E} [\hat{P}^* \hat{I}] \quad (3.8)$$

and

$$\hat{L}(\hat{C}_0) := \mathbb{E} [\hat{P}^* \hat{P} \hat{C}_0 \hat{P}^* \hat{P}] = \mathbb{E} [\hat{P}^* (\hat{I} - \hat{P} \hat{m}_0) (\hat{I} - \hat{P} \hat{m}_0)^* \hat{P}], \quad (3.9)$$

equations whose interpretations we shall discuss in this section. Note that (3.8) and (3.9) can be seen as the limiting cases of (2.5) and (2.6) for $\sigma^2 = 0$, $p \rightarrow \infty$, and $n \rightarrow \infty$.

In the equations above, we define

$$\begin{aligned} \hat{P}^*: C_0^\infty(\mathbb{R}^2) &\rightarrow C_0^\infty(\mathbb{R}^3)' \quad \text{by} \quad \langle \hat{P}^* \hat{J}, \hat{Y} \rangle: \\ &= \langle \hat{J}, \hat{P} \hat{Y} \rangle_{L^2(\mathbb{R}^2)} \quad \text{where} \quad \hat{J} \in C_0^\infty(\mathbb{R}^2), \hat{Y} \in C_0^\infty(\mathbb{R}^3), \quad \text{and} \quad C_0^\infty(\mathbb{R}^3)' \end{aligned}$$

is the space of continuous linear functionals $C_0^\infty(\mathbb{R}^3)$. Thus, both sides of (3.8) are elements of $C_0^\infty(\mathbb{R}^3)'$. To verify this equation, we apply both sides to a test function \hat{Y} :

$$\begin{aligned} \langle \mathbb{E} [\hat{P}^* \hat{I}], \hat{Y} \rangle &= \mathbb{E} [\langle \hat{I}, \hat{P} \hat{Y} \rangle_{L^2(\mathbb{R}^2)}] = \mathbb{E} [\mathbb{E} [\langle \hat{I}, \hat{P} \hat{Y} \rangle_{L^2(\mathbb{R}^2)} | \hat{P}]] \\ &= \mathbb{E} [\langle \hat{P} \hat{m}_0, \hat{P} \hat{Y} \rangle_{L^2(\mathbb{R}^2)}] = \langle \mathbb{E} [\hat{P}^* \hat{P} \hat{m}_0], \hat{Y} \rangle. \end{aligned} \quad (3.10)$$

Note that

$$\begin{aligned} \langle \hat{P}^* \hat{P} \hat{m}, \hat{Y} \rangle &= \langle \hat{P} \hat{m}, \hat{P} \hat{Y} \rangle_{L^2(\mathbb{R}^2)} = \int_{\mathbb{R}^2} \hat{m}(\hat{x}R^1 + \hat{y}R^2) \overline{\hat{Y}(\hat{x}R^1 + \hat{y}R^2)} d\hat{x}d\hat{y} \\ &= \int_{\mathbb{R}^3} \hat{m}(\xi) \overline{\hat{Y}(\xi)} \delta(\xi \cdot R^3) d\xi, \end{aligned} \quad (3.11)$$

from which it follows that in the sense of distributions,

$$(\hat{P}^* \hat{P} \hat{m})(\xi) = \hat{m}(\xi) \delta(\xi \cdot R^3) \quad (3.12)$$

Intuitively, this means that $\hat{P}^* \hat{P}$ inputs the volume \hat{m} and outputs a “truncated” volume that coincides with \hat{m} on a plane perpendicular to the viewing angle and is zero elsewhere. This reflects the fact that the image $\hat{I} = PX$ only gives us information about X on a single central plane. When we aggregate this information over all possible R , we obtain the operator \hat{A} :

$$\begin{aligned} \hat{A}\hat{m}(\xi) &= \mathbb{E} [\hat{m}(\xi) \delta(\xi \cdot R^3)] = \hat{m}(\xi) \frac{1}{4\pi} \int_{S^2} \delta(\xi \cdot \theta) d\theta \\ &= \frac{\hat{m}(\xi)}{|\xi|} \frac{1}{4\pi} \int_{S^2} \delta\left(\frac{\xi}{|\xi|} \cdot \theta\right) d\theta = \frac{\hat{m}(\xi)}{2|\xi|}. \end{aligned} \tag{3.13}$$

We used the fact that R^3 is uniformly distributed over S^2 if R is uniformly distributed over $SO(3)$. Here, $d\theta$ is the surface measure on S^2 (hence the normalization by 4π). The last step holds because the integral over S^2 is equal to the circumference of a great circle on S^2 , so it is 2π .

By comparing (3.8) and (2.7), it is clear that P is the analogue of A^{Pn} for infinite n and p . Also, (3.8) echoes the heuristic formula (3.5). The backprojection operator \hat{C} simply “inserts” a 2D image into 3D space by situating it in the plane perpendicular to the viewing direction of the image, and so the RHS of (3.8) at a point ξ is the accumulation of values $\hat{C}(\xi)$. Moreover, the operator P is diagonal, and for each ξ , P reflects the measure of the set $\xi \in \hat{C}$; i.e., the density of central planes passing through ξ under the uniform distribution of rotations. Thus, (3.8) encodes the intuition from the first equation in (3.5). Inverting P involves multiplying by the radial factor $2|\xi|$. In tomography, this factor is called the ramp filter [38]. Traditional tomographic algorithms proceed by applying the ramp filter to the

projection data and then backprojecting. Note that solving $\frac{1}{2|\xi|} \hat{m}_0(\xi) = E[\hat{P} * \hat{I}]$ implies performing these operations in the reverse order; however, backprojection and application of the ramp filter commute.

Now we move on to (3.9). Both sides of this equation are continuous linear functionals on $C_0^\infty(R^3) \times C_0^\infty(R^3)$. Indeed, for $\hat{Y}_1, \hat{Y}_2 \in C_0^\infty(R^3)$, the LHS of (3.9) operates on (\hat{Y}_1, \hat{Y}_2) through the definition

$$\left(\hat{P}^* \hat{P} \hat{C} \hat{P}^* \hat{P}\right) (\hat{Y}_1, \hat{Y}_2) = \left\langle \hat{C}, \left(\hat{P}^* \hat{P} \hat{Y}_1, \hat{P}^* \hat{P} \hat{Y}_2\right) \right\rangle, \tag{3.14}$$

where we view $\hat{C} \in C_0^\infty(R^3 \times R^3)$ as operating on pairs (η_1, η_2) of elements in $C_0^\infty(R^3)$ via

$$\left\langle \hat{C}, (\eta_1, \eta_2) \right\rangle := \int_{\mathbb{R}^3 \times \mathbb{R}^3} \overline{\eta_1(\xi_1)} \eta_2(\xi_2) \hat{C}(\xi_1, \xi_2) d\xi_1 d\xi_2. \tag{3.15}$$

Using these definitions, we verify (3.9):

$$\begin{aligned}
 & \mathbb{E} \left[\hat{P}^* \left(\hat{I} - \hat{P} \hat{m}_0 \right) \left(\hat{I} - \hat{P} \hat{m}_0 \right)^* \hat{P} \right] \left(\hat{Y}_1, \hat{Y}_2 \right) \\
 & := \mathbb{E} \left[\left\langle \hat{P}^* \left(\hat{I} - \hat{P} \hat{m}_0 \right), \hat{Y}_1 \right\rangle \overline{\left\langle \hat{P}^* \left(\hat{I} - \hat{P} \hat{m}_0 \right), \hat{Y}_2 \right\rangle} \right] \\
 & = \mathbb{E} \left[\left\langle \hat{P}^* \hat{P} \hat{Y}_1, \hat{X} - \hat{m}_0 \right\rangle \overline{\left\langle \hat{P}^* \hat{P} \hat{Y}_2, \hat{X} - \hat{m}_0 \right\rangle} \right] \\
 & = \mathbb{E} \left[\int_{\mathbb{R}^3 \times \mathbb{R}^3} \overline{\hat{P}^* \hat{P} \hat{Y}_1 \left(\xi_1 \right)} \left(\hat{X} \left(\xi_1 \right) - \hat{m}_0 \left(\xi_1 \right) \right) \hat{P}^* \hat{P} \hat{Y}_2 \left(\xi_2 \right) \overline{\left(\hat{X} \left(\xi_2 \right) - \hat{m}_0 \left(\xi_2 \right) \right)} d\xi_1 d\xi_2 \right] \quad (3.16) \\
 & = \mathbb{E} \left[\int_{\mathbb{R}^3 \times \mathbb{R}^3} \overline{\hat{P}^* \hat{P} \hat{Y}_1 \left(\xi_1 \right)} \hat{C}_0 \left(\xi_1, \xi_2 \right) \hat{P}^* \hat{P} \hat{Y}_2 \left(\xi_2 \right) d\xi_1 d\xi_2 \right] \\
 & = \mathbb{E} \left[\left\langle \hat{C}_0 \left(\hat{P}^* \hat{P} \hat{Y}_1, \hat{P}^* \hat{P} \hat{Y}_2 \right) \right\rangle \right] = \mathbb{E} \left[\hat{P}^* \hat{P} \hat{C}_0 \hat{P}^* \hat{P} \left(\hat{Y}_1, \hat{Y}_2 \right) \right].
 \end{aligned}$$

Substituting (3.12) into the last two lines of the preceding calculation, we find

$$\left(\hat{P}^* \hat{P} \hat{C}_0 \hat{P}^* \hat{P} \right) \left(\xi_1, \xi_2 \right) = \hat{C} \left(\xi_1, \xi_2 \right) \delta \left(\xi_1 \cdot R^3 \right) \delta \left(\xi_2 \cdot R^3 \right) \quad (3.17)$$

This reflects the fact that an image \hat{I} gives us information about $P \left(\xi, \xi \right)$ for $\xi, \xi \in \hat{C}$.

Taking the expectation over R , we find that

$$\begin{aligned}
 \left(\hat{L} \hat{C} \right) \left(\xi_2, \xi_2 \right) & = \mathbb{E} \left[\hat{C} \left(\xi_1, \xi_2 \right) \delta \left(\xi_1 \cdot R^3 \right) \delta \left(\xi_2 \cdot R^3 \right) \right] \\
 & = \hat{C} \left(\xi_1, \xi_2 \right) \frac{1}{4\pi} \int_{S^2} \delta \left(\xi_1 \cdot \theta \right) \theta \left(\xi_2 \cdot \theta \right) =: \hat{C} \left(\xi_1, \xi_2 \right) K \left(\xi_1, \xi_2 \right). \quad (3.18)
 \end{aligned}$$

Like \hat{A} , the operator P is diagonal. P is a fundamental operator in tomographic inverse problems involving variability; we term it the *projection covariance transform*. In the same way that (3.8) reflected the first equation of (3.5), we see that (3.9) resembles the second equation of (3.5). In particular, the kernel value $K \left(\xi_1, \xi_2 \right)$ reflects the density of central planes passing through ξ_1, ξ_2 .

To understand this kernel, let us compute it explicitly. We have

$$K \left(\xi_1, \xi_2 \right) = \frac{1}{4\pi} \int_{S^2} \delta \left(\xi_1 \cdot \theta \right) \delta \left(\xi_2 \cdot \theta \right) d\theta. \quad (3.19)$$

For fixed ξ_1 , note that $\delta \left(\xi_1 \cdot \theta \right)$ is supported on the great circle of S^2 perpendicular to ξ_1 . Similarly, $\delta \left(\xi_2 \cdot \theta \right)$ corresponds to a great circle perpendicular to ξ_2 . Choose $\xi_1, \xi_2 \in R^3$ so that $|\xi_1 \times \xi_2| \neq 0$. Then, note that these two great circles intersect in two antipodal points $\theta = \pm \left(\xi_1 \times \xi_2 \right) / |\xi_1 \times \xi_2|$, and the RHS of (3.19) corresponds to the total measure of $\delta \left(\xi_1 \cdot \theta \right) \delta \left(\xi_2 \cdot \theta \right)$ at those two points.

To calculate this measure explicitly, let us define the approximation to the identity

$\delta_\epsilon \left(t \right) = \frac{1}{2\epsilon} X_{[-\epsilon, \epsilon]} \left(t \right)$. Fix $E_1, E_2 > 0$. Note that $\delta_1 \left(\xi_1 \cdot \theta \right)$ is supported on a strip of width $2E_1 / |\xi_1|$ centered at the great circle perpendicular to ξ_1 . $\delta_2 \left(\xi_2 \cdot \theta \right)$ is supported on a strip of width $2E_2 / |\xi_2|$ intersecting the first strip transversely. For small E_1, E_2 , the intersection of

the two strips consists of two approximately parallelogram-shaped regions, S_1 and S_2 (see Figure 3).

The sine of the angle between the diagonals of each of these regions is $|\xi_1 \times \xi_2|/|\xi_1||\xi_2|$, and a simple calculation shows that the area of one of these regions is $2E_1 2E_2/|\xi_1 \times \xi_2|$. It follows that

$$\begin{aligned} K(\xi_1, \xi_2) &= \frac{1}{4\pi} \int_{S^2} \delta(\xi_1 \cdot \theta) \delta(\xi_2 \cdot \theta) d\theta = \lim_{\epsilon_1, \epsilon_2 \rightarrow 0} \frac{1}{4\pi} \int_{S^2} \delta(\xi_1 \cdot \theta) \delta(\xi_2 \cdot \theta) d\theta \\ &= \lim_{\epsilon_1, \epsilon_2 \rightarrow 0} \frac{1}{4\pi} \int_{S_1 \cup S_2} \frac{1}{2\epsilon_1} \frac{1}{2\epsilon_2} d\theta = \lim_{\epsilon_1, \epsilon_2 \rightarrow 0} \frac{1}{4\pi} 2 \frac{4\epsilon_1 \epsilon_2}{|\xi_1 \times \xi_2|} \frac{1}{2\epsilon_1} \frac{1}{2\epsilon_2} \\ &= \frac{1}{4\pi} \frac{2}{|\xi_1 \times \xi_2|}. \end{aligned} \tag{3.20}$$

This analytic form of K sheds light on the geometry of \hat{C} . Recall that $K(\xi_1, \xi_2)$ is a measure of the density of central planes passing through ξ_1 and ξ_2 . Note that this density is nonzero everywhere, which reflects the fact that there is a central plane passing through each pair of points in \mathbb{R}^3 . The denominator in K is proportional to the magnitudes $|\xi_1|$ and $|\xi_2|$, which indicates that there is a greater density of planes passing through pairs of points nearer the origin. Finally, note that K varies inversely with the sine of the angle between ξ_1 and ξ_2 ; indeed, a greater density of central planes pass through a pair of points nearly collinear with the origin. In fact, there is a singularity in K when ξ_1, ξ_2 are linearly dependent, reflecting the fact that infinitely many central planes pass through collinear points. As a way to sum up the geometry encoded in K , note that except for the factor of $1/4\pi$, $1/K$ is the area of the triangle spanned by the vectors ξ_1 and ξ_2 . For this reason, we call $1/K$ the *triangular area filter*.

Note that the triangular area filter is analogous to the ramp filter: it grows linearly with the frequencies $|\xi_1|$ and $|\xi_2|$ to compensate for the loss of high frequency information incurred by the geometry of the problem. So, this filter is a generalization of the ramp filter appearing in the estimation of the mean to the covariance estimation problem. The latter has a somewhat more intricate geometry, which is reflected in K .

The properties of K translate into the robustness of inverting P (supposing we added noise to our model). In particular, the robustness of recovering $P(\xi, \xi)$ grows with $K(\xi, \xi)$. For example, recovering higher frequencies in \hat{C} is more difficult. However, the fact that K is everywhere positive means that P is at least invertible. This statement is important in proving theoretical results about our estimators, as we saw in section 2.2. Note that an analogous problem of estimating the covariance matrix of 2D objects from their one-dimensional line projections would not satisfy this condition, because for most pairs of points in \mathbb{R}^2 , there is not a line passing through both points as well as the origin.

3.2. The discrete covariance estimation problem

The calculation in the preceding section shows that if we could sample images continuously and if we had access to projection images from all viewing angles, then P would become a diagonal operator. In this section, we explore the modifications necessary for the realistic case where we must work with finite-dimensional representations of volumes and images.

Our idea is to follow what we did in the fully continuous case treated above and estimate the covariance matrix in the Fourier domain. One possibility is to choose a Cartesian basis in the Fourier domain. With this basis, a tempting way to define PP^s would be to restrict the Fourier 3D grid to the pixels of a 2D central slice by nearest-neighbor interpolation. This would make PP^s a coordinate-selection operator, making L^{Pn} diagonal. However, this computational simplicity comes at a great cost in accuracy; numerical experiments show that the errors induced by such a coarse interpolation scheme are unacceptably large. Such an interpolation error should not come as a surprise, considering similar interpolation errors in computerized tomography [38]. Hence, we must choose other bases for the Fourier volumes and images.

The finite sampling rate of the images limits the 3D frequencies we can hope to reconstruct. Indeed, since the images are sampled on an $N \times N$ grid confining a disc of radius 1, the corresponding Nyquist bandlimit is $\omega_{\text{Nyq}} = N\pi/2$. Hence, the images carry no information past this 2D bandlimit. By the Fourier slice theorem, this means that we also have no information about X past the 3D bandlimit ω_{Nyq} . In practice, the exponentially decaying envelope of the CTF function renders even fewer frequencies possible to reconstruct. Moreover, we saw in section 3.1 and will see in section 6.2 that reconstruction of Σ_0 becomes more ill-conditioned as the frequency increases. Hence, it often makes sense to take a cutoff $\omega_{\text{max}} < \omega_{\text{Nyq}}$. We can choose ω_{max} to correspond to an effective grid size of N_{res} pixels, where $N_{\text{res}} \ll N$. In this case, we would choose $\omega_{\text{max}} = N_{\text{res}}\pi/2$. Thus, it is natural to search for X in a space of functions bandlimited in $B_{\omega_{\text{max}}}$ (the ball of radius ω_{max}) and with most of their energy contained in the unit ball. The optimal space B with respect to these constraints is spanned by a finite set of 3D Slepian functions [56]. For a given bandlimit ω_{max} , we have

$$p = \dim(B) = \frac{2}{9\pi} \omega_{\text{max}}^3. \quad (3.21)$$

This dimension is called the *Shannon number*, and is the trace of the kernel in [56, eq. 6].

For the purposes of this section, let us work abstractly with the finite-dimensional spaces $VP \subset C_0(B_{\omega_{\text{max}}})$ and $IP \subset C_0(D_{\omega_{\text{max}}})$, which represent Fourier volumes and Fourier images, respectively ($D_{\omega_{\text{max}}} \subset \mathbb{R}$ is the disc of radius ω_{max}). For example, VP could be spanned by the Fourier transforms of the 3D Slepian functions. Let

$$\hat{V} = \text{span} \{ \hat{h}_j \}, \quad \hat{I} = \text{span} \{ \hat{g}_i \} \quad (3.22)$$

with $\dim(\hat{V}) = pP$ and $\dim(\hat{I}) = qP$. Assume that for all R , $\hat{C}(\hat{V}) \subset IP$ (i.e., we do not need to worry about interpolation). Denote by PP the matrix expression of \hat{C} . Thus, $PP \in \mathbb{C}^{qP \times pP}$. Let X^{P1}, \dots, X^{Pn} be the representations of P, \dots, \hat{C} in the basis for \hat{V} .

Since we are given the images I_s in the pixel basis \mathbb{R}^q , let us consider how to map these images into \hat{I} . Let $Q_1 : \mathbb{R}^q \rightarrow IP$ be the mapping which fits (in the least-squares sense) an element of IP to the pixel values defined by a vector in \mathbb{R}^q . It is easiest to express Q_1 in

terms of the reverse mapping $Q_2 : IP \rightarrow \mathbb{R}^q$. The i th column of Q_2 consists of the evaluations of g_i at the real-domain grid points inside the unit disc. It is easy to see that the least-squares method of defining $Q_1 = Q_2^+ = (Q_2^H Q_2)^{-1} Q_2^H$

Now, note that

$$I = SPX + E \Rightarrow Q_1 I = Q_1 SPX + Q_1 E \approx \hat{P} \hat{X} + Q_1 E. \quad (3.23)$$

The last approximate equality is due to the Fourier slice theorem. The inaccuracy comes from the discretization operator S . Note that $Var [Q_1 E] = \sigma^2 Q_1 Q_1^H = \sigma^2 (Q_2^H Q_2)^{-1}$. We would like the latter matrix to be a multiple of the identity matrix so that the noise in the images remains white. Let us calculate the entries of $Q_2^H Q_2$ in terms of the basis functions g_i . Given the fact that we are working with volumes h_i which have most of their energy concentrated in the unit ball, it follows that g_i have most of their energy concentrated in the unit disc. If x_1, \dots, x_q are the real-domain image grid points, it follows that

$$\begin{aligned} (Q_2^H Q_2)_{ij} &= \sum_{r=1}^q \overline{g_1(x_r)} g_j(x_r) \approx \frac{q}{\pi} \int_{|x| \leq 1} \overline{g_i(x)} g_j(x) dx \\ &\approx \frac{q}{\pi} \langle g_i, g_j \rangle_{L^2(\mathbb{R}^2)} = \frac{q}{\pi} \frac{1}{(2\pi)^2} \langle \hat{g}_i, \hat{g}_j \rangle_{L^2(\mathbb{R}^2)}. \end{aligned} \quad (3.24)$$

It follows that in order for $Q_2^H Q_2$ to be (approximately) a multiple of the identity matrix, we should require $\{gP_i\}$ to be an orthonormal set in $L^2(\mathbb{R}^2)$. If we let $c_q = 4\pi^3/q$, then we find that

$$Q_1 Q_1^H \approx c_q I_{\hat{q}}. \quad (3.25)$$

It follows that, if we make the approximations in (3.23) and (3.25), we can formulate the heterogeneity problem entirely in the Fourier domain as follows:

$$\hat{I} = \hat{P} \hat{X} = \hat{E}, \quad (3.26)$$

where $Var[\hat{E}] = \sigma^2 c_q I_{qP}$. Thus, we have an instance of Problem (1.1) with σ^2 replaced by $\sigma^2 c_q$, q replaced by qP , and p replaced by pP . We seek $\mu P_0 = E[XP]$ and $\Sigma^{P0} = Var[XP]$. Equations (2.5) and (2.6) become

$$\hat{A}_n \hat{\mu}_n := \left(\frac{1}{n} \sum_{s=1}^n \hat{P}_s^H \hat{P}_s \right) \hat{\mu}_n = \frac{1}{n} \sum_{s=1}^n \hat{P}_s^H \hat{I}_s \quad (3.27)$$

and

$$\begin{aligned} \hat{L}_n \hat{\Sigma}_n &= \frac{1}{n} \sum_{s=1}^n \hat{P}_s^H \hat{P}_s \hat{\Sigma}_n \hat{P}_s^H \hat{P}_s \\ &= \frac{1}{n} \sum_{s=1}^n \hat{P}_s^H \left(\hat{I}_s - \hat{P}_s \hat{\mu}_n \right) \left(\hat{I}_s - \hat{P}_s \hat{\mu}_n \right)^H \hat{P}_s - \sigma^2 c_q \hat{A}_n =: \hat{B}_n. \end{aligned} \quad (3.28)$$

3.3. Exploring AP and $L\hat{C}$

In this section, we seek to find expressions for AP and $L\hat{C}$ like those in (3.13) and (3.18). The reason for finding these limiting operators is twofold. First of all, recall that the theoretical results in section 2.2 depend on the invertibility of these limiting operators. Hence, knowing AP and LP in the cryo-EM case will allow us to verify the assumptions of Propositions 2.1 and 2.2. Second, the law of large numbers guarantees that for large n , we have $A^{Pn} \approx AP$ and $L^{Pn} \approx L\hat{C}$. We shall see in section 5 that approximating A^{Pn} and L^{Pn} by their limiting counterparts makes possible the tractable implementation of our algorithm.

In section 3.1, we worked with functions $m : \mathbb{R}^3 \rightarrow \mathbb{C}$ and $P : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}$. Now, we are in a finite-dimensional setup, and we have formulated (3.27) and (3.28) in terms of vectors and matrices. Nevertheless, in the finite-dimensional case we can still work with functions as we did in section 3.1 via the identifications

$$\hat{\mu} \in C^{\hat{p}} \leftrightarrow \hat{m} = \sum_{i=1}^{\hat{p}} \hat{\mu}_i \hat{h}_i \in \hat{V}, \quad \hat{\Sigma} \in C^{\hat{p} \times \hat{p}} \leftrightarrow \hat{C} = \sum_{i,j=1}^{\hat{p}} \hat{\Sigma}_{i,j} \hat{h}_i \otimes \hat{h}_j \in \hat{V} \otimes \hat{V}, \quad (3.29)$$

where we define

$$\left(\hat{h}_i \otimes \hat{h}_j \right) (\xi_1, \xi_2) = \hat{h}_i (\xi_1) \overline{\hat{h}_j (\xi_2)}, \quad (3.30)$$

and $VP \otimes VP = \text{span}\{h^{Pi} \otimes h^{Pj}\}$. Thus, we identify C^{pP} and $C^{pP \times pP}$ with spaces of bandlimited functions. For these identifications to be isometries, we must endow VP with an inner product for which the h^{Pi} are orthonormal. We consider a family of inner products, weighted by radial functions $w(|\xi|)$:

$$\langle \hat{h}_i, \hat{h}_j \rangle_{L_w^2(\mathbb{R}^3)} = \int_{\mathbb{R}^3} \hat{h}_i(\xi) \overline{\hat{h}_j(\xi)} w(|\xi|) d\xi = \delta_{ij}. \quad (3.31)$$

The inner product on $VP \otimes VP$ is inherited from that of $V\hat{C}$.

Note that A^{Pn} and L^{Pn} both involve the projection-backprojection operator $P^{PH} P^{Ps}$. Let us see how to express $P^{PH} P^{Ps}$ as an operator on $V\hat{C}$. The i th column of P^{Ps} is the representation of δ in the orthonormal basis for I . Hence, using the isomorphism $C^{qP} \leftrightarrow I$ and reasoning along the lines of (3.11), we find that

$$\left(\hat{P}_s^H \hat{P}_s\right)_{i,j} = \overline{\langle \hat{P}_s \hat{h}_i, \hat{P}_s \hat{h}_j \rangle} L^2(R^2) = \int_{R^3} \overline{\hat{h}_i(\xi)} \hat{h}_j(\xi) \delta(\xi \cdot R_s^3) d\xi. \quad (3.32)$$

Note that here and throughout this section, we perform manipulations (like those in section 3.1) that involve treating elements of VP as test functions for distributions. We will ultimately construct VP so that its elements are continuous, but not in $C^\infty(R^3)$, as assumed in section 3.1. Nevertheless, since we are only dealing with distributions of order zero, continuity of the elements of VP is sufficient.

From (3.32), it follows that if $\mu P \in C^{pP} \leftrightarrow \hat{m}$

$$\begin{aligned} \left(\hat{P}_s^H \hat{P}_s\right) \hat{\mu} &\leftrightarrow \sum_{i=1}^{\hat{p}} \hat{h}_i \sum_{j=1}^{\hat{p}} \left(\hat{P}_s^H \hat{P}_s\right)_{ij} \mu_j &= \sum_{i=1}^{\hat{p}} \hat{h}_i \sum_{j=1}^{\hat{p}} \int_{R^3} \overline{\hat{h}_i(\xi)} \hat{h}_j(\xi) \delta(\xi \cdot R_s^3) d\xi \\ & &= \sum_{i=1}^{\hat{p}} \hat{h}_i \int_{R^3} (\hat{m}(\xi) \delta(\xi \cdot R_s^3)) \overline{\hat{h}_i(\xi)} d\xi \\ & &=: \pi_{\hat{V}}(\hat{m}(\xi) \delta(\xi \cdot R_s^3)), \end{aligned} \quad (3.33)$$

where $\pi_{\hat{\gamma}}: C_0^\infty(R^3)' \rightarrow \hat{V}$

$$\pi_{\hat{V}}(\eta) = \sum_i \hat{h}_i \langle \eta, \hat{h}_i \rangle, \quad \eta \in C_0^\infty(R^3)' \quad (3.34)$$

is a projection onto the finite-dimensional subspace $\hat{V}\hat{C}$.

In analogy with (3.8), we have

$$\hat{A}_{\hat{\mu}} \leftrightarrow E \left[\pi_{\hat{V}}(\hat{m}(\xi) \delta(\xi \cdot R^3)) \right] = \pi_{\hat{V}} \left(\frac{\hat{m}(\xi)}{2|\xi|} \right). \quad (3.35)$$

Note AP resembles the operator P obtained in (3.8), with the addition of the “low-pass filter” π_{VP} . As a particular choice of weight, one might consider $w(|\xi|) = 1/|\xi|$ in order to cancel the ramp filter. For this weight, note that

$$\hat{A}_{\hat{\mu}} \leftrightarrow \pi_{\hat{V}} \left(\frac{\hat{m}(\xi)}{2|\xi|} \right) = \pi_{\hat{V}}^w \left(\frac{1}{2} \hat{m}(\xi) \right) = \frac{1}{2} \hat{m}(\xi), \quad w(|\xi|) = 1/|\xi|, \quad (3.36)$$

where $\pi_{\hat{V}}^w$ is the *orthogonal* projection onto VP with respect to the weight w . Thus, for this

weight we find that $\hat{A} = \frac{1}{2} I_{\hat{p}}$

A calculation analagous to (3.33) shows that for $\Sigma P \in C^{pP \times pP} \leftrightarrow \hat{C}$

$$\hat{P}_s^H \hat{P}_s \hat{\Sigma} \hat{P}_s^H \hat{P}_s \leftrightarrow \pi_{\hat{V} \otimes \hat{V}} \left(\hat{C}(\xi_1, \xi_2) \delta(\xi_1 \cdot R_s^3) \delta(\xi_2 \cdot R_s^3) \right). \quad (3.37)$$

Then, taking the expectation over \mathbf{R}^3 , we find that

$$\hat{L} \hat{\Sigma} \leftrightarrow \pi_{\hat{V} \otimes \hat{V}} \left(\hat{C}(\xi_1, \xi_2) K(\xi_1, \xi_2) \right). \quad (3.38)$$

This shows that between $L\hat{C}$ is linked to P via the low-pass-filter π_P analogously to (3.34).

3.4. Properties of AP and $L\hat{C}$

In this section, we will prove several results about AP and $L\hat{C}$, defined in (3.35) and (3.38). We start by proving a useful lemma.

Lemma 3.3—For $\eta \in C_0^\infty(\mathbf{R}^3)'$ and $\hat{Y} \hat{C}$, we have

$$\langle \pi_{\hat{V}} \eta, \hat{Y} \rangle_{L_w^2(\mathbf{R}^3)} = \langle \eta, \hat{Y} \rangle. \quad (3.39)$$

Likewise, if $\eta \in C_0^\infty(\mathbf{R}^3 \times \mathbf{R}^3)'$ and $\hat{C} \in \hat{V} \otimes \hat{V}$, we have

$$\langle \pi_{\hat{V} \otimes \hat{V}} \eta, \hat{C} \rangle_{L_w^2(\mathbf{R}^3 \times \mathbf{R}^3)} = \langle \eta, \hat{C} \rangle. \quad (3.40)$$

Proof: Indeed, we have

$$\begin{aligned} \langle \pi_{\hat{V}} \eta, \hat{Y} \rangle_{L_w^2(\mathbf{R}^3)} &= \sum_{i=1}^{\hat{p}} \langle \eta, \hat{h}_i \rangle \langle \hat{h}_i, \hat{Y} \rangle_{L_w^2(\mathbf{R}^3)} \\ &= \left\langle \eta, \sum_{i=1}^{\hat{p}} \langle \hat{Y}, \hat{h}_i \rangle_{L_w^2(\mathbf{R}^3)} \hat{h}_i \right\rangle = \langle \eta, \hat{Y} \rangle. \end{aligned} \quad (3.41)$$

The proof of the second claim is similar.

Note that AP and LP are self-adjoint and PSD because each A^{Pn} and L^{Pn} satisfies this property. In the next proposition, we bound the minimum eigenvalues of these two operators from below.

Proposition 3.4—Let $M_w(\omega_{\max}) = \max_{|\xi| \leq \omega_{\max}} |\xi| w(|\xi|)$. Then,

$$\lambda_{\min}(\hat{A}) \geq \frac{1}{2M_w(\omega_{\max})}, \quad \lambda_{\min}(\hat{L}) \geq \frac{1}{2M_w^2(\omega_{\max})}. \quad (3.42)$$

Proof: Let $\mu P \in C^{pP} \leftrightarrow m$ find

$$\begin{aligned} \langle \hat{A}\hat{\mu}, \hat{\mu} \rangle_{C^{\hat{P}}} &= \langle \pi_V \left(\hat{m} \frac{1}{2|\xi|} \right), \hat{m} \rangle_{L_w^2(R^3)} = \langle \hat{m} \frac{1}{2|\xi|}, \hat{m} \rangle \\ &= \int_{B_{\omega_{max}}} |\hat{m}(\xi)|^2 \frac{1}{2|\xi|w(|\xi|)} w(|\xi|) d\xi \geq \frac{1}{2M_w(\omega_{max})} \|\hat{m}\|_{L_w^2(R^3)}^2 = \frac{1}{2M_w(\omega_{max})} \|\hat{\mu}\|^2. \end{aligned} \quad (3.43)$$

The bound on the minimum eigenvalue of LP follows from a similar argument, using (3.38) and the following bound:

$$\min_{\xi_1, \xi_2 \in B_{\omega_{max}}} \frac{K(\xi_1, \xi_2)}{w(|\xi_1|)w(|\xi_2|)} = \min_{\xi_1, \xi_2 \in B_{\omega_{max}}} \frac{1}{2\pi|\xi_1 \times \xi_2|w(|\xi_1|)w(|\xi_2|)} \geq \frac{1}{2\pi M_w^2(\omega_{max})}. \quad (3.44)$$

By inspecting $M_w(\omega_{max})$, we see that choosing $w = 1/|\xi|$ leads to better conditioning of both AP and $L\hat{C}$, as compared to $w = 1$. This is because the former weight compensates for the loss of information at higher frequencies. We see from (3.36) that for $w = 1/|\xi|$, AP is perfectly conditioned. This weight also cancels the linear growth of the triangular area filter with radial frequency. However, it does not cancel K altogether, since the dependency on $\sin \gamma$ in the denominators in (3.44) remains, where γ is the angle between ξ_1 and ξ_2 .

The maximum eigenvalue of LP cannot be bounded as easily, since the quotient in (3.44) is not bounded from above. A bound on $\lambda_{\max}(L\hat{C})$ might be obtained by using the fact that a bandlimited P can only be concentrated to a limited extent around the singular set $\{\xi_1, \xi_2 : |\xi_1 \times \xi_2| = 0\}$.

Finally, we prove another property of AP and $L\hat{C}$: they commute with rotations. Let us define the group action of $SO(3)$ on functions $R^3 \rightarrow C$ as follows: for $R \in SO(3)$ and $P : R^3 \rightarrow C$, let $R \cdot \hat{C}(\xi) = \hat{C}(R^T \xi)$. Likewise, define the group action of $SO(3)$ on functions $P : R^3 \times R^3 \rightarrow C$ via $R \cdot \hat{C}(\xi_1, \xi_2) = \hat{C}(R^T \xi_1, R^T \xi_2)$.

Proposition 3.5—*Suppose that the subspace VP is closed under rotations. Then, for any $Y \in V, C \in V \otimes V$, and $R \in SO(3)$, we have*

$$R \cdot (\hat{A}\hat{Y}) = \hat{A}(R \cdot \hat{Y}), \quad R \cdot (\hat{L}\hat{C}) = \hat{L}(R \cdot \hat{C}), \quad (3.45)$$

where APX and LPX are understood via the identifications (3.29).

Proof: We begin by proving the first half of (3.45). First of all, extend the group action of

$SO(3)$ to the space $C_0^\infty(R^3)'$, via

$$\langle R \cdot \eta \hat{Y}, \hat{Y} \rangle := \langle \eta, R^{-1} \cdot \hat{Y} \rangle, \quad \hat{Y} \in C_0^\infty(R^3). \quad (3.46)$$

We claim that for any $\eta \in C_0^\infty(R^3)'$, we have $R \cdot (\pi_P \eta) = \pi_P(R \cdot \eta)$. Since VP is closed under rotations, both sides of this equation are elements of $V\hat{C}$. We can verify their equality by

taking an inner product with an arbitrary element $\hat{C} \in V\hat{C}$. Using Lemma 3.3 and the fact that VP is closed under rotations, we obtain

$$\begin{aligned} \langle R.(\pi_{\hat{V}}\eta), \hat{Y} \rangle_{L_w^2(R^3)} &= \langle \pi_{\hat{V}}\eta, R^{-1}.\hat{Y} \rangle_{L_w^2(R^3)} = \langle \eta, R^{-1}.\hat{Y} \rangle \\ &= \langle R.\eta, \hat{Y} \rangle = \langle \pi_{\hat{V}}(R.\eta), \hat{Y} \rangle_{L_w^2(R^3)}. \end{aligned} \quad (3.47)$$

Next, we claim that for any $\hat{C} \in V\hat{C}$, we have $R.(P\hat{C}) = \hat{C}(R.\hat{C})$. To check whether these two elements of $\eta \in C_0^\infty(R^3)$ are the same, we apply them to a test function $\hat{Z} \in C_0^\infty(R^3)$:

$$\begin{aligned} \langle R.(\hat{A}\hat{Y}), \hat{Z} \rangle &= \langle \hat{A}\hat{Y}, R^{-1}.\hat{Z} \rangle = \int_{R^3} \frac{\hat{Y}(\xi)}{2|\xi|} \hat{Z}(R\xi) d\xi \\ &= \int_{R^3} \frac{\hat{Y}(R^H\xi)}{2|\xi|} \hat{Z}(\xi) d\xi = \langle \hat{A}(R.\hat{Y}), \hat{Z} \rangle. \end{aligned} \quad (3.48)$$

Putting together what we have, we find that

$$R.(\hat{A}\hat{Y}) = R.(\pi_{\hat{V}}(\hat{A}\hat{Y})) = \pi_{\hat{V}}(R.(\hat{A}\hat{Y})) = \hat{A}(R.\hat{Y}), \quad (3.49)$$

which proves the first half of (3.45). The second half is proved analogously.

This property of AP and LP is to be expected, given the rotationally symmetric nature of these operators. This suggests that LP can be studied further using the representation theory of $SO(3)$.

Finally, let us check that the assumptions of Propositions 2.1 and 2.2 hold in the cryo-EM case. It follows from Proposition 3.4 that as long as $M_w(\omega_{\max}) < \infty$, the limiting operators AP and LP are invertible. Of course, it is always possible to choose such a weight w . In particular the weights already considered, $w = 1, 1/|\xi|$ satisfy this property. Moreover, by rotational symmetry, $|P\hat{C}(R)|$ is independent of R , and so of course this quantity is uniformly bounded. Thus, we have checked all the necessary assumptions to arrive at the following conclusion.

Proposition 3.6—If we neglect the errors incurred in moving to the Fourier domain and assume that the rotations are drawn uniformly from $SO(3)$, then the estimators μP_n and Σ^{Pn} obtained from (3.27) and (3.28) are consistent.

4. Using $\hat{\mu}_n, \hat{\Sigma}_n$ to determine the conformations

To solve Problem 1.2, we must do more than just estimate μP_0 and Σ^{P0} . We must also estimate C, X^{Pc} , and p_c , where X^{Pc} is the coefficient vector of Pc in the basis for $V\hat{C}$. Once we solve (3.27) and (3.28) for μP_n and Σ^{Pn} , we perform the following steps.

From the discussion on high-dimensional PCA in section 2.3, we expect to determine the number of structural states by inspecting the spectrum of Σ^{Pn} . We expect the spectrum of Σ^{Pn} to consist of a bulk distribution along with $C - 1$ separate eigenvalues (assuming the

SNR is sufficiently high), a fact confirmed by our numerical results. Hence, given Σ^{Pn} , we can estimate C .

Next, we discuss how to reconstruct X^{P1}, \dots, X^{PC} and p_1, \dots, p_C . Our approach is similar to Penczek, Kimmel, and Spahn's [43]. By the principle of PCA, the leading eigenvectors of $\hat{\Sigma}_n$ span the space of mean subtracted volumes

$\hat{X}^1 - \hat{\mu}_0, \dots, \hat{X}^C - \hat{\mu}_0$. If $\hat{V}_n^1, \dots, \hat{V}_n^{C-1}$ are the leading eigenvectors of $\hat{\Sigma}_n$, we can write

$$\hat{X}_s \approx \hat{\mu}_n + \sum_{c'=1}^{C-1} \alpha_{s,c'} \hat{V}_n^{c'}. \quad (4.1)$$

Note that there is only approximate equality because we have replaced the mean μP_0 by the estimated mean μP_n , and the eigenvectors of Σ^{P0} by those of Σ^{Pn} . We would like to recover the coefficients $\alpha_s = (\alpha_{s,1}, \dots, \alpha_{s,C-1})$, but the X^{Ps} are unknown. Nevertheless, if we project the above equation by P^{Ps} , then we get

$$\sum_{c'=1}^{C-1} \alpha_{s,c'} \hat{P}_s \hat{V}_n^{c'} \approx \hat{P}_s \hat{X}_s - \hat{P}_s \hat{\mu}_n = (\hat{I}_s - \hat{P}_s \hat{\mu}_n) - \hat{\epsilon}_s. \quad (4.2)$$

For each s , we can now solve this equation for the coefficient vector α_s in the least-squares sense. This gives us n vectors in C^{C-1} . These should be clustered around C points

$\alpha^c = (\alpha_1^c, \dots, \alpha_{C-1}^c)$ for $c = 1, \dots, C$, corresponding to the C underlying volumes. At this point, Penczek, Kimmel, and Spahn propose to perform K -means clustering on α_s in order to deduce which image corresponds to which class. However, if the images are too noisy, then it would be impossible to separate the classes via clustering. Note that in order to reconstruct the original volumes, all we need are the means of the C clusters of coordinates. If the mean volume and top eigenvectors are approximately correct, then the main source of noise in the coordinates is the Gaussian noise in the images. It follows that the distribution of the coordinates in C^{C-1} is a mixture of Gaussians. Hence, we can find the means α^c of each cluster using either an EM algorithm (of which the K -means algorithm used by Penczek is a limiting case [8]) or the method of moments, e.g., [23]. In the current implementation, we use an EM algorithm. Once we have the C mean vectors, we can reconstruct the original volumes using (4.1). Putting these steps together, we arrive at a high-level algorithm to solve the heterogeneity problem (see Algorithm 1).

ALGORITHM 1. HIGH-LEVEL ALGORITHM FOR HETEROGENEITY PROBLEM (PROBLEM 1.2).

- 1: **Input:** n images I_s and the corresponding rotations R_s .
 - 2: Estimate the noise level σ^2 from the corner regions of the images.
 - 3: Choose bases for \mathcal{S} and \mathcal{V} .
 - 4: Map the images I_s into $\hat{I}_s \in \mathbb{C}^{\mathcal{V}}$.
 - 5: Estimate $\hat{\mu}_n, \hat{\Sigma}_n$ by solving (3.27) and (3.28).
 - 6: Compute the eigendecomposition of $\hat{\Sigma}_n$ and estimate its rank r . Set $C = r + 1$.
 - 7: Estimate each $\alpha_s \in \mathbb{C}^{C-1}$ by solving (4.2) using least squares.
 - 8: Find α^c and p_c by applying either EM or a method of moments algorithm to α_s .
 - 9: Using α^c , find $\hat{X}^1, \dots, \hat{X}^C$ from (4.1). Map volumes back to real domain for visualization.
 - 10: **Output:** $C, \hat{X}^1, \dots, \hat{X}^C, p_1, \dots, p_C$.
-

5. Implementing Algorithm 1

In this section, we confront the practical challenges of implementing Algorithm 1. We consider different approaches to addressing these challenges and choose one approach to explore further.

5.1. Computational challenges and approaches

The main computational challenge in Algorithm 1 is solving for Σ^{P^n} in

$$\hat{L}_n(\hat{\Sigma}_n) = \hat{B}_n, \quad (5.1)$$

given the immense size of this problem. Two possibilities for inverting L^{P^n} immediately come to mind. The first is to treat (5.1) as a large system of linear equations, viewing Σ^{P^n} as a vector in C^{pP^2} and $L\hat{C}$ as a matrix in $C^{pP^2 \times pP^2}$. In this scheme, the matrix $L\hat{C}$ could be computed once and stored. However, this approach has an unreasonably large storage

requirement. Since $\hat{p} = O(N_{res}^3)$, it follows that L^{P^n} has size $N_{res}^6 \times N_{res}^6$. Even for a small N_{res} value such as 17, each dimension of L^{P^n} is 1.8×10^6 . Storing such a large L^{P^n} requires over 23 terabytes. Moreover, inverting this matrix naively is completely intractable.

The second possibility is to abandon the idea of forming L^{P^n} as a matrix, and instead to use an iterative algorithm, such as the conjugate gradient (CG) algorithm, based on repeatedly applying L^{P^n} to an input matrix. From (3.28), we see that applying L^{P^n} to a matrix is dominated by n multiplications of a $qP \times pP$ matrix by a $pP \times pP$ matrix, which costs $n\hat{q}\hat{p}^2 = O(nN_{res}^o)$. If κ_n is the condition number of L^{P^n} , then CG will converge in $O(\kappa_n)$ iterations (see, e.g., [58]). Hence, while the storage requirement of this alternative algorithm is only $O\hat{p}^2 = O(nN_{res}^o)$, the computational complexity is $O(nN^8 \kappa_n)$. Thus, the price to pay for reducing the storage requirement is that n matrix multiplications must be performed at each iteration. While this computational complexity might render the algorithm impractical for a regular computer, one can take advantage of the fact that the n matrix multiplications can be performed in parallel.

We propose a third numerical scheme, one which requires substantially less storage than the first scheme above and does not require $O(n)$ operations at each iteration. We assume that the R_s are drawn from the uniform distribution over $SO(3)$, and so for large n , the operator L^{P^n} does not differ much from its limiting counterpart LP (defined in (3.38)). Hence, if we replace L^{P^n} by LP in (5.1), we would not be making too large an error. Of course, LP is a matrix of the same size as L^{P^n} , so it is also impossible to store on a computer. However, we leverage the analytic form of $L\hat{C}$ in order to invert it more efficiently. At this point, we have not yet chosen the spaces VP and $I\hat{C}$, and by constructing these carefully we give LP a special structure. This approach also entails a tradeoff: in practice the approximation $L^{P^n} \approx L\hat{C}$ is accurate to the extent that R^3, \dots, R^3 are uniformly distributed on S^2 . Hence, we must extract a subset of the given rotations whose viewing angles are approximately uniformly distributed on the sphere. Thus, the sacrifice we make in this approach is a reduction in the sample size. Moreover, since the subselected viewing directions are no longer statistically

independent, the theoretical consistency result stated in Proposition 3.6 does not necessarily extend to this numerical scheme.

Nevertheless, the latter approach is promising because the complexity of inverting LP is independent of the number of images, and this computation might be tractable for reasonable values of N_{res} if LP has enough structure. It remains to construct VP and IP to induce a special structure in $L\hat{C}$, which we turn to next.

5.2. Choosing VP to make LP sparse and block diagonal—In this section, we write down an expression for an individual element of $L\hat{C}$, and discover that for judiciously chosen basis functions $\hat{C}h_i$, the matrix LP becomes sparse and block diagonal.

First, let us fix a functional form for the basis elements h^{Pi} : let

$$\hat{h}_i(r, \alpha) = f_i(r) a_i(\alpha), \quad r \in R^+, \alpha \in S^2, \quad (5.2)$$

where $f_i : R^+ \rightarrow R$ are radial functions and $a_i : S^2 \rightarrow C$ are spherical harmonics. Note, for example, that the 3D Slepian functions have this form [56, eq. 110]. If the h^{Pi} are orthogonal with respect to the weight w , then

$$\langle f_i, f_j \rangle_{L^2_{r^2w(r)}} \langle a_i, a_j \rangle_{L^2(S^2)} = \delta_{ij}, \quad (5.3)$$

where we use L^2_w as a shorthand for $L^2_w(R^+)$. The 3D Slepian functions satisfy the above condition with $w = 1$, because they are orthogonal in $L^2(R^3)$.

Next, we write down the formula for an element $L^{Pi1, i2, j1, j2}$ (here, j_1, j_2 are the indices of the input matrix, and i_1, i_2 are the indices of the output matrix). From (3.38) and Lemma 3.3,

we find

$$\begin{aligned} \hat{L}_{i_1, i_2, j_1, j_2} &= \left\langle \pi_{\hat{V} \otimes \hat{V}} \left(\left(\hat{h}_{j_1} \otimes \hat{h}_{j_2} \right) K \right), \hat{h}_{i_1} \otimes \hat{h}_{i_2} \right\rangle_{L^2_w(R^3 \times R^3)} \\ &= \int_{R^3 \times R^3} \left(\hat{h}_{j_1} \otimes \hat{h}_{j_2} \right) (\xi_1, \xi_2) K(\xi_1, \xi_2) \overline{\left(\hat{h}_{i_1} \otimes \hat{h}_{i_2} \right) (\xi_1, \xi_2)} d\xi_1 d\xi_2 \\ &= \int_{S^2 \times S^2} \int_{R^+ \times R^+} \left(\hat{h}_{j_1} \otimes \hat{h}_{j_2} \right) (\xi_1, \xi_2) \overline{\left(\hat{h}_{i_1} \otimes \hat{h}_{i_2} \right) (\xi_1, \xi_2)} \frac{1}{2\pi r_1 r_2 |\alpha \times \beta|} r_1^2 r_2^2 dr_1 dr_2 d\alpha d\beta \\ &= \langle f_{j_1}, f_{i_1} \rangle_{L^2_r} \langle f_{j_2}, f_{i_2} \rangle_{L^2_r} \int_{S^2 \times S^2} (a_{j_1} \otimes a_{j_2})(\alpha, \beta) \overline{(a_{i_1} \otimes a_{i_2})(\alpha, \beta)} \frac{1}{2\pi |\alpha \times \beta|} d\alpha d\beta. \end{aligned} \quad (5.4)$$

Thus, to make many of the radial inner products in LP correct weight is vanish, we see from (5.3) that the

$$w(r) = \frac{1}{r}. \quad (5.5)$$

Recall that this is the weight needed to cancel the ramp filter in AP (see (3.36)). We obtain a cancellation in LP as well because the kernel of this operator also grows linearly with radial

frequency. From this point on, w will represent the weight above, and we will work in the corresponding weighted L^2 space.

What are sets of functions of the form (5.2) that are orthonormal in $L^2(\mathbb{R}^3)$? If we chose 3D Slepian functions, we would get the functional form

$$\hat{h}_{k,l,m}(r, \alpha) = f_{k,l}(r) Y_l^m(\alpha). \quad (5.6)$$

However, these functions are orthonormal with weight $w = 1$ instead of $w = 1/r$. Consider modifying this construction by replacing the $f_{k,R}(r)$ by the radial functions arising in the 2D Slepian functions. These satisfy the property

$$\langle f_{k_1,l_1}, f_{k_2,l_2} \rangle_{L_r^2} = 0 \quad \text{if } l_1 \neq l_2, k_1 \neq k_2. \quad (5.7)$$

With this property (5.6) becomes orthonormal in $L^2(\mathbb{R}^3)$. This gives LP a certain degree of sparsity. However, note that the construction (5.6) has different families of L^2 -orthogonal radial functions corresponding to each angular function. Thus, we only have orthogonality of the radial functions $f_k^{1,R1}$ and $f_k^{2,R2}$ when $l_1 = .e_2$. Thus, many of the terms f_j, f_i in (5.4) are still nonzero.

A drastic improvement on (5.6) would be to devise an orthogonal basis in L^2 that used one set of r -weighted orthogonal functions f_k for all the angular functions, rather than a separate set for each angular function. Namely, suppose we chose

$$\hat{h}_{k,l,m}(\tau, \alpha)(r, \alpha) = f_k(r) Y_l^m(\alpha), \quad (k, l, m) \in J, \quad (5.8)$$

where J is some indexing set. Note that f_k and J need to be carefully constructed so that $\text{span}\{h_{k,R,m}\} \approx B$ (see section 5.3 for this construction). We have

$$f_k(r) Y_{l,m}(\alpha) = \hat{h}_{k,l,m}(r, \alpha) = \hat{h}_{k,l,m}(-r, -\alpha) = f_k(-r) Y_{l,m}(-\alpha) = (-1)^l f_k(-r) Y_{l,m}(\alpha). \quad (5.9)$$

Here, we assume that each f_k is either even or odd at the origin, and we extend $f_k(r)$ to $r \in \mathbb{R}$ according to this parity. The above calculation implies that f_k should have the same parity as $.e$. Let us suppose that f_k has the same parity as k . Then, it follows that $(k, .e, m) \in J$ only if $k = .e \pmod 2$. Thus, $h_{k,R,m}$ will be orthonormal in L^2 if

$$\{f_k : k \equiv 0 \pmod 2\} \quad \text{and} \quad \{f_k : k \equiv 1 \pmod 2\} \quad \text{are orthonormal in } L_r^2. \quad (5.10)$$

If we let k_i be the radial index corresponding to i , then we claim that the above construction implies

$$\hat{L}_{i_1, i_2, j_1, j_2}^{k_1, k_2} = \delta_{k_{i_1} k_{j_1}} \delta_{k_{i_2} k_{j_2}} \int_{S^2 \times S^2} (a_{j_1} \otimes a_{j_2})(\alpha, \beta) \overline{(a_{i_1} \otimes a_{i_2})(\alpha, \beta)} \frac{1}{2\pi|\alpha \times \beta|} d\alpha d\beta. \quad (5.11)$$

This statement does not follow immediately from (5.10), because we still need to check the case when $k_{i_1} \neq k_{j_1} \pmod 2$. Note that in this case, the dependence on a in the integral over $S^2 \times S^2$ is odd, and so indeed $L^{P_i, i, j, j} = 0$ in that case as well. If $\hat{V}\mathcal{C}$ is the space spanned by $f_k(r)Y^m(a)$ for all $.e, m$, then the above implies that LP operates separately on each $V^{Pk} \otimes V^{Pk_2}$. In the language of matrices, this means that if we divide Σ^{Pn} into blocks Σ^{Pk_1, k_2} based on radial indices, LP operates on these blocks separately. We denote each of the corresponding ‘‘blocks’’ of LP by L^{Pk_1, k_2} . Let us reindex the angular functions so that a^k denotes the i th angular basis function paired with f_k . From (5.11), we have

$$\hat{L}_{i_1, i_2, j_1, j_2}^{k_1, k_2} = \int_{S^2 \times S^2} (a_{j_1}^{k_1} \otimes a_{j_2}^{k_2})(\alpha, \beta) \overline{a_{i_1}^{k_1} \otimes a_{i_2}^{k_2}(\alpha, \beta)} \frac{1}{2\pi|\alpha \times \beta|} d\alpha d\beta. \quad (5.12)$$

This block diagonal structure of LP makes it much easier to invert. Nevertheless, each block L^{Pk_1, k_2} is a square matrix with dimension $O(k_1^2 k_2^2)$. Hence, inverting the larger blocks of $L\hat{\mathcal{C}}$ can be difficult. Remarkably, it turns out that each block of LP is sparse. In Appendix C, we simplify the above integral over $S^2 \times S^2$. Then, (5.12) becomes

$$\hat{L}_{i_1, i_2, j_1, j_2}^{k_1, k_2} = \sum_{l, m} c(l) C_{l, m} \left(\overline{a_{i_1}^{k_1} a_{j_1}^{k_1}} \right) \overline{C_{l, m} \left(\overline{a_{i_2}^{k_2} a_{j_2}^{k_2}} \right)}, \quad (5.13)$$

where the constants $c(.e)$ are defined in (C.8) and $C_{R, m}(\psi\hat{\mathcal{C}})$ is the $.e, m$ coefficient in the spherical harmonic expansion of $\psi P : S^2 \rightarrow \mathbb{C}$. It turns out that the above expression is zero for most sets of indices. To see why, recall that the functions a^k are spherical harmonics. It is known that the product $Y^m Y^{m^*}$ can be expressed as a linear combination of harmonics Y^M , where $M = m + m^1$ and $|.e - .e^1| \leq |.e + .e^1|$. Thus, $C^m(a; a_j)$ are sparse vectors, which shows that each block L^{Pk_1, k_2} is sparse. For example, $L^{P15, 15}$ has each dimension approximately 2×10^4 . However, only about 10^7 elements of this block are nonzero, which is only about 3% of its total number of entries. This is about the same number of elements as a 3000×3000 full matrix.

Thus, we have found a way to tractably solve the covariance matrix estimation problem: reconstruct Σ^{Pn} (approximately) by solving the sparse linear systems

$$\hat{L}^{k_1, k_2} \hat{\Sigma}_n^{k_1, k_2} = \hat{B}_n^{k_1, k_2}, \quad (5.14)$$

where we recall that B^{Pn} is the RHS of (3.28). Also, using the fact that $\hat{A}_n \approx \hat{A} = \frac{1}{2} I_{\hat{q}}$, we can estimate μP_n from

$$\hat{\mu}_n = \frac{2}{n} \sum_{s=1}^n \hat{P}_s^H \hat{I}_s. \quad (5.15)$$

In the next two sections, we discuss how to choose the radial components $f_k(r)$ and define IP and VP more precisely.

5.3. Constructing $f_k(r)$ and the space VP —We have discussed so far that

$$\hat{V} = \text{span} (\{f_k(r) Y_l^m(\theta, \varphi) : (k, l, m) \in J\}) \quad (5.16)$$

with $(k, .e, m) \in J$ only if $k = .e \bmod 2$. Moreover, we have required the orthonormality condition (5.10). However, recall that we initially assumed that the real-domain functions X_s belonged to the space of 3D Slepian functions B . Thus, we must choose VP to approximate the image of B under the Fourier transform. Hence, the basis functions $f_k(r) Y^m(\theta, \varphi)$ should be supported in the ball of radius ω_{\max} and have their inverse Fourier transforms concentrated in the unit ball. Moreover, we must have $\dim(\hat{V}) \approx \dim(B)$. Finally, the basis functions h^{Pi} should be analytic at the origin (they are the truncated Fourier transforms of compactly supported molecules). We begin by examining this condition.

Expanding h^{Pi} in a Taylor series near the origin up to a certain degree, we can approximate it locally as a finite sum of homogeneous polynomials. By [57, Theorem 2.1], a homogeneous polynomial of degree d can be expressed as

$$H_d(\xi) = r^d (c_d Y_d(\alpha) + c_{d-2} Y_{d-2}(\alpha) + \dots), \quad (5.17)$$

where each Y_R represents a linear combination of spherical harmonics of degree $.e$. Hence, if $(k, .e, m) \in J$, then we require that $f_k(r) = a_R r^R + a_{R+2} r^{R+2} + \dots$, where some coefficients can be zero. We satisfy this requirement by constructing f_0, f_1, \dots so that

$$f_k(r) = \alpha_{k,k} r^k + \alpha_{k,k+2} r^{k+2} + \dots \quad (5.18)$$

for small r with $\alpha_{k,k} \neq 0$, and combine f_k with Y^m if $k = .e \bmod 2$ and $.e = k$. This leads to the following set of 3D basis functions:

$$\{\hat{h}_i\} = \left\{ f_0 Y_0^0, f_1 Y_1^{-1}, f_1 Y_1^0, f_1 Y_1^1, f_2 Y_0^0, f_2 Y_2^{-2}, \dots, f_2 Y_2^2, \dots \right\}. \quad (5.19)$$

Written another way, we define

$$\hat{V} = \text{span} (\{f_k(r) Y_l^m(\theta, \varphi) : 0 \leq k \leq K, \quad l = k \pmod{2}, \quad 0 \leq l \leq k, \quad |m| \leq l\}). \quad (5.20)$$

Following the reasoning preceding (5.17), it can be seen that near the origin, this basis spans the set of polynomial functions up to degree K .

Now, consider the real- and Fourier-domain content of h^{Pi} . The bandlimitedness requirement on X_s is satisfied if and only if the functions f_k are supported in the interval $[0, \omega_{max}]$. To deal with the real domain requirement, we need the inverse Fourier transform of $f_k(r)Y^m(\theta, \varphi)$. With the Fourier convention (3.1), it follows from [2] that

$$F^{-1}(f_k(r)Y_l^m(\theta, \varphi))(r_x, \theta_x, \varphi_x) = \frac{1}{2\pi^2} i^l \left(\int_0^\infty f_k(r) j_l(rr_x) r^2 dr \right) Y_l^m(\theta_x, \varphi_x) = \frac{1}{2\pi^2} i^l (S_l f_k)(r_x) Y_l^m(\theta_x, \varphi_x). \tag{5.21}$$

Here, j_R is the spherical Bessel function of order $.e$, and S_R is the spherical Hankel transform. Also note that (r, θ, φ) are Fourier-domain spherical coordinates, while $(r_x, \theta_x, \varphi_x)$ are their real-domain counterparts. Thus, satisfying the real-domain concentration requirement amounts to maximizing the percentage of the energy of $S_R f_k$ that is contained in $[0, 1]$ for $0 \leq k \leq K, 0 \leq .e \leq k, .e = k \pmod 2$.

Finally, we have arrived at the criteria we would like $f_k(r)$ to satisfy:

1. $\text{supp } f_k \subset [0, \omega_{max}]$;
2. $\{f_k : k \text{ even}\}$ and $\{f_k : k \text{ odd}\}$ orthonormal in $L^2(\mathbb{R}^+, r)$;
3. $f_k(r) = a_{k,k}r^k + a_{k,k+2}r^{k+2} + \dots$ near $r = 0$;
4. under the above conditions, maximize the percentage of the energy of $S_R f_k$ in $[0, 1]$, for $0 \leq k \leq K, 0 \leq .e \leq k, .e = k \pmod 2$.

While it might be possible to find an optimal set of such functions $\{f_k\}$ by solving an optimization problem, we can directly construct a set of functions that satisfactorily satisfies the above criteria.

Note that since $.e$ ranges in $[0, k]$, it follows that for larger k , we need to have higher-order spherical Hankel transforms $S_R f_k$ remain concentrated in $[0, 1]$. Since higher-order spherical Hankel transforms tend to be less concentrated for oscillatory functions, it makes sense to choose f_k to be less and less oscillatory as k increases. Note that the functions f_k cannot all have only few oscillations because the even and odd functions must form orthonormal sets. Using this intuition, we construct f_k as follows. Since the even and odd f_k can be constructed independently, we will illustrate the idea by constructing the even f_k . For simplicity, let us assume that K is odd, with $K = 2K_0 + 1$. define the cutoff $\chi = \chi([0, \omega_{max}])$. First, consider the sequence

$$J_0(z_{0,k_0+1}r/\omega_{max})\chi, J_2(z_{2,K_0}r/\omega_{max})\chi, \dots, J_{2K_0}(z_{2K_0,1}r/\omega_{max})\chi \tag{5.22}$$

where $z_{k,m}$ is the m th positive zero of J_k (the k th-order Bessel function). Note that the functions in this list satisfy criteria 1 (by construction) and 3 (due to the asymptotics of the Bessel function at the origin). Also note that we have chosen the scaling of the arguments of the Bessel functions so that the number of zero crossings decreases as the list goes on. Thus, the functions become less and less oscillatory, which is the pattern that might lead to satisfying criterion 4. However, since these functions might not be orthogonal with respect

to the weight r , we need to orthonormalize them with respect to this weight (via Gram–Schmidt). We need to be careful to orthonormalize them in such a way as to preserve the properties that they already satisfy. This can be achieved by running the (r -weighted) Gram–Schmidt algorithm from higher k towards lower k . This preserves the supports of the functions, their asymptotics at the origin, and the oscillation pattern. Moreover, the orthogonality property now holds as well. See Figure 4 for the first several even radial basis functions. Constructing the odd radial functions requires following an analogous procedure. Also, changing the parity of K requires the obvious modifications.

It remains to choose K . We do this based on how well criterion 4 is satisfied. For example, we can calculate how much energy of $S_R f_k$ is contained in the unit interval for all $0 \leq k \leq K$, $0 \leq e \leq k$, $e = k \bmod 2$. Numerical experiments show that $K = N_{\text{res}} - 2$ is a reasonable value. For each value of N_{res} that we tested, this choice led to $S_R f_k$ having at least 80% of its energy concentrated in the unit interval for each relevant (e, k) , and at least 95% on average over all such pairs (e, k) . Thus our experiments show that for our choice of f_k , choosing roughly $K \approx N_{\text{res}}$ leads to acceptable satisfaction of criterion 4. A short calculation yields

$$\hat{p} = \dim(\hat{V}) = \sum_{k=0}^K \frac{(k+1)(k+2)}{2} = \frac{(K+1)(K+2)(K+3)}{6} \approx \frac{N_{\text{res}}^3}{6} = \frac{4\omega_{\text{max}}^3}{3\pi^3}. \quad (5.23)$$

$$\hat{P}_s(f_k(r) Y_l^m(\theta, \varphi)) = f_k(r) \hat{P}_s(Y_l^m(\theta, \varphi)). \quad (5.24)$$

Hence, we have $pP/p = 6/\pi^2 \approx 0.6$. Hence, the dimension of the space VP we have constructed is within a constant factor of the dimension of B . This factor is the price we pay for the computational simplicity VP provides.

Note that a different construction of f_k might have even better results. Choosing better radial functions can be the topic of further research. In any case, the specific choice of f_k does not affect the structure of our algorithm at all because LP is independent of these functions, as can be seen from (5.12). Thus, the selection of the radial basis functions can be viewed as an independent module in our algorithm. The radial functions we choose here work well in numerical experiments; see section 7.

5.4. Constructing \mathcal{P} —Finally, the remaining piece in our construction is the finite dimensional space of Fourier images, \hat{IC} . To motivate our construction, consider applying P_s to a basis element of V . The first observation to make is that the radial components $f_k(r)$ factor through \hat{C} completely:

Recall from (3.21) that

$$\hat{P}_s(Y_l^m(\theta, \varphi)) = \sum_{\substack{|m'| \leq l \\ m' \equiv l \pmod{2}}} c_{l,m,m'}(R_s) \frac{1}{\sqrt{2\pi}} e^{im'\varphi}, \quad (5.25)$$

Note that the \hat{C} on the LHS should be interpreted as $C(\mathbb{R}^3) \rightarrow C(\mathbb{R}^2)$, whereas the one on the RHS is the restricted map $C(S^2) \rightarrow C(S^1)$, which we also call P . The correct interpretation should be clear in each case. Viewed in this new way, $\hat{C} : C(S^2) \rightarrow C(S^1)$ rotates a function on the sphere by $R_s \in SO(3)$, and then restricts the result to the equator.

By the rotational properties of spherical harmonics, a short calculation shows that

$$f_k(r) Y_l^m(\theta, \varphi) \in \hat{V} \Rightarrow \frac{1}{\sqrt{2\pi}} f_k(r) e^{im\varphi} \in \hat{J}, \quad m = -l, -l+2, \dots, l-2, l. \quad (5.26)$$

where the constants c_{R,m,m^*} depend on the Wigner D matrices D^R [36]. Hence, $P(\hat{V}\hat{C}) \subset IP$ if

$$(\hat{g}_i) = \left\{ \frac{1}{\sqrt{2\pi}} f_0(r), \frac{1}{\sqrt{2\pi}} f_1(r) e^{-i\varphi}, \frac{1}{\sqrt{2\pi}} f_1(r) e^{i\varphi}, \frac{1}{\sqrt{2\pi}} f_2(r) e^{-2i\varphi}, \frac{1}{\sqrt{2\pi}} f_2(r), \frac{1}{\sqrt{2\pi}} f_2(r) e^{2i\varphi}, \dots \right\}. \quad (5.27)$$

Thus, we construct IP by pairing f_k with $\frac{1}{\sqrt{2\pi}} e^{im\varphi}$ if $k = m \pmod 2$ and $m \leq k$. This leads to the 2D basis functions

$$\hat{J} = \text{span} \left(\left\{ \frac{1}{\sqrt{2\pi}} f_k(r) e^{im\varphi} : 0 \leq k \leq K, m = k \pmod 2, |m| \leq k \right\} \right). \quad (5.28)$$

Written another way, we construct

$$\hat{q} = \dim(\hat{J}) = \sum_{k=0}^K (k+1) = \frac{(K+1)(K+2)}{2} \approx \frac{N_{res}^2}{2} = \frac{2\omega_{max}^2}{\pi^2}. \quad (5.29)$$

If P^k is the subspace of IP spanned by the basis functions with radial component f_k , (5.24) shows that $P(VP) \subset \hat{I}\hat{C}$ for each k . Thus, $P\hat{C}$ has a block diagonal structure, as depicted in Figure 5.

Let us now compare the dimension of IP to that of the corresponding space of 2D Slepian functions, as we did the previous section. We have

$$nnz(\hat{L}^{k_1, k_2}) \leq \frac{1}{k_1 + k_2 + 1} \left(\frac{(k_1 + 1)(k_1 + 2)(k_2 + 1)(k_2 + 2)}{4} \right)^2, \quad (6.1)$$

The Shannon number in 2D corresponding to the bandlimit ω_{max} is $\omega^2/4$. Thus, we are short of this dimension by a constant factor of $8/\pi^2 \approx 0.8$. Another comparison to make is that the number of grid points in the disc inscribed in the $N_{res} \times N_{res}$ grid is $\pi N^2 = \omega^2/\pi$. Thus, $\dim(\hat{I}\hat{C})$ is short of this number by a factor of 2 . Note that this is the same factor that was obtained in a similar situation in [69], so IP is comparable in terms of approximation to the Fourier–Bessel space constructed there.

Thus, by this point we have fully specified our algorithm for the heterogeneity problem. After finding Σ^{P^n} numerically via (5.14), we can proceed as in steps 6–9 of Algorithm 1 to solve Problem 1.2.

6. Algorithm complexity

In this section, we explore the consequences of the constructions of VP and IP for the complexity of the proposed algorithm. We also compare this complexity with that of the straightforward CG approach discussed in section 5.1.

To calculate the computational complexity of inverting the sparse matrix L^{Pk_1,k_2} via the CG algorithm, we must bound the number of nonzero elements of this matrix and its condition number.

6.1. Sparsity of LP and storage complexity—Preliminary numerical experiments confirm the following conjecture.

Conjecture 6.1

$$\text{maximum eigenvalue of } \hat{L}^{k,k} = 0.2358 + 0.1357k. \quad (6.2)$$

where $\text{nnz}(A)$ is the number of nonzero elements in a matrix A , and the term involving the square is the total number of elements in L^{Pk_1,k_2} .

Hence, the percentage of nonzero elements in each block of $L\hat{C}$ decays linearly with the frequencies associated with that block. This conjecture remains to be verified theoretically.

We pause here to note the storage complexity of the proposed algorithm, which is dominated by the cost of storing $L\hat{C}$. In fact, since we process all the blocks separately, only storing one L^{Pk_1,k_2} at a time will suffice. Hence, the storage complexity is the memory required to store the largest block of $L\hat{C}$, which is $\text{nnz}(L^{PK,K}) = O(K^7) = O(N^7)$ required storage for a full matrix of the size of $L\hat{C}$, which is $\text{nnz}(\hat{L}^{K,K}) = O(K^7) = O(N_{res}^7)$. Compare this to the required storage for a full matrix of the size of L , which is $O(N_{res}^{12})$.

6.2. Condition number of $L\hat{C}$ —Here we find the condition number of each L^{Pk_1,k_2} . We already proved in Proposition 3.4 that $\lambda_{\min}(L\hat{C}) = 1/2\pi$. For any k_1, k_2 , this implies that $\lambda_{\min}(L^{Pk_1,k_2}) = 1/2\pi$. This is confirmed by a numerical experiment: in Figure 6(a) are plotted the minimum eigenvalues of $L^{Pk,k}$ for $0 \leq k \leq 15$. Note that the eigenvalues actually approach the value $1/2\pi$ (marked with a horizontal line) as k increases. We remarked in section 3.4 that an upper bound on the maximum eigenvalue is harder to find. Nevertheless, numerical experiments have led us to the following conjecture.

Conjecture 6.2: *The maximal eigenvalue of L^{Pk_1,k_2} grows linearly with $\min(k_1, k_2)$.*

Moreover, a plot of the maximal eigenvalue of $L^{Pk,k}$ shows a clear linear dependence on k . See Figure 6(b). The line of best fit is approximately

$$k \left(\hat{L}^{k_1, k_2} \right) \leq 1.4818 + 0.8524 \min(k_1, k_2). \quad (6.3)$$

Taken together, Proposition 3.4 and Conjecture 6.2 imply the following conjecture about the condition number of $L^{P^{k_1, k_2}}$, which we denote by $\kappa(L^{P^{k_1, k_2}})$.

Conjecture 6.3

$$k \left(\hat{L} \right) \leq 1.4818 + .852K. \quad (6.4)$$

In particular, this implies that

$$\begin{aligned} & \text{complexity of inverting } \hat{L} \\ & \lesssim \sum_{k_1, k_2=0}^K \sqrt{k \left(\hat{L}^{k_1, k_2} \right) \text{nnz} \left(\hat{L}^{k_1, k_2} \right)} \\ & \lesssim \sum_{k_1, k_2=0}^K \sqrt{\min(k_1, k_2=0)} \frac{1}{k_1 + k_2 + 1} \left(\frac{(k_1+1)(k_1+2)(k_2+1)(k_2+2)}{4} \right)^2 \quad (6.5) \\ & \lesssim \sum_{k_1, k_2=0}^K (k_1 k_2)^{1/4} \frac{1}{\sqrt{k_1 k_2}} k_1^4 k_2^4 \\ & \lesssim \sum_{k_1=0}^K k_1^{3.75} \sum_{k_2=0}^K k_2^{3.75} \lesssim K^{4.75} K^{4.75} = K^{9.5}. \end{aligned}$$

6.3. Algorithm complexity—Using the above results, we estimate the computational complexity of Algorithm 1. We proceed step by step through the algorithm and estimate the complexity at each stage. Before we do so, note that due to the block diagonal structure of P^{Ps} (depicted in Figure 5), it can be easily shown that an application of P^{Ps} or P^{PH} costs $O(K^4)$.

Sending the images from the pixel domain into IP requires n applications of the matrix $Q_1 \in C^{qP \times q}$, which costs $O(nqqP) = O(nN^2N^2)$. Note that this complexity can be improved using an algorithm of the type [39], but in this paper we do not delve into the details of this alternative.

Finding μP_n from (5.15) requires n applications of the matrix \hat{P}_s^H , and so has complexity $O(nK^4) = O(nN_{res}^4)$.

Next, we must compute the matrix B^{Pn} . Note that the second term in B^{Pn} can be replaced by a multiple of the identity matrix by (3.36), so only the first term of B^{Pn} must be computed.

Note that B^{Pn} is a sum of n matrices, and each matrix can be found as the outer product of $P_s (P^{Ps} - P^{Ps} \mu P_n) \in C^{pP}$ with itself. Calculating this vector has complexity $O(K^4)$, from which it follows that calculating B^{Pn} costs $O(nK^4) = O(nN^4)$.

Next, we must invert $L\hat{C}$. As mentioned in section 5.1, the inversion of a matrix A via CG takes $\kappa(A)$ iterations. If A is sparse, then applying it to a vector has complexity $\text{nnz}(A)$.

Hence, the total complexity for inverting a sparse matrix is $\kappa(A)\text{nnz}(A)$. Conjectures 6.1 and 6.3 imply that

$$O\left(nN^2N_{res}^2+N_{res}^{9.5}\right). \quad (6.6)$$

Since LP has size of the order $K^6 \times K^6$, note that the complexity of inverting a full matrix of this size would be K^{18} . Thus, our efforts to make LP sparse have saved us a $K^{8.5}$ complexity factor. Moreover, the fact that LP is block diagonal makes its inversion parallelizable.

Assuming that $C = O(1)$, solving each of the n least-squares problems (4.2) is dominated by a constant number of applications of P^{Ps} to a vector. Thus, finding α_s for $s = 1, \dots, n$ costs

$$O\left(N_{res}^{12}\right)$$

Next, we must fit a mixture of Gaussians to α_s to find a^c . An EM approach to this problem requires $O(n)$ operations per iteration. Assuming that the number of iterations is constant, finding a^c has complexity $O(n)$.

Finally, reconstructing X^{P^c} via (4.1) has complexity $O(N^3)$.

Hence, neglecting lower-order terms, we find that the total complexity of our algorithm is

$$O\left(nN_{res}^{7.5}\right). \quad (6.7)$$

6.4. Comparison to straightforward CG approach—We mentioned in section 5.1 that a CG approach is possible in which at each iteration, we apply L^{P^n} to $\Sigma\hat{C}$ using the definition (3.28). This approach has the advantage of not requiring uniformly spaced viewing directions. While the condition number of L^{P^n} depends on the rotations R_1, \dots, R_n , let us assume here that $\kappa(L^{P^n}) \approx \kappa(L\hat{C})$. We estimated the computational complexity of this approach in section 5.1, but at that point we assumed that each P^{Ps} was a full matrix. If we use the bases VP and $I\hat{C}$, we reap the benefit of the block diagonal structure of P^{Ps} . Hence, for each s , evaluating $P^{PH} P^{Ps} \Sigma P^{PH} P^{Ps}$ is dominated by the multiplication $P^{Ps} \Sigma P$, which has complexity N^7 . Hence, applying L^{P^n} to $\Sigma\hat{C}$ has complexity nN^7 . By (6.4), we assume that $\kappa(L^{P^n}) = O(N_{res})$. Hence, the full complexity of inverting LP using the conjugate gradient approach is (6.7) $O(nN^{7.5})$.

$$SNR = \frac{P(\text{signal})}{P(\text{noise})}, \quad (7.1)$$

Compare this to a complexity of $O(N^{9.5})$ for inverting $L\hat{C}$. Given that n is usually on the order of 10^5 or 10^6 , for moderate values of N_{res} we have $N^{9.5} \gg nN^{7.5}$. Nevertheless, both algorithms have possibilities for parallelization, which might change their relative complexities. As for memory requirements, note that the straightforward CG algorithm only

requires $O(N^6)$ storage, whereas we saw in section 6.1 that the proposed algorithm requires $O(N^7)$ storage.

In summary, these two algorithms each have their strengths and weaknesses, and it would be interesting to write parallel implementations for both and compare their performances. In the present paper, we have implemented and tested only the algorithm based on inverting $L\hat{C}$.

7. Numerical results

Here, we provide numerical results illustrating Algorithm 1, with the bases IP and VP chosen so as to make $L\hat{C}$ sparse, as discussed in section 5. The results presented below are intended for proof-of-concept purposes, and they demonstrate the qualitative behavior of the algorithm. They are not, however, biologically significant results. We have considered an idealized setup in which there is no CTF effect, and have assumed that the rotations R_s (and translations) have been estimated perfectly. In this way, we do not perform a “full-cycle” experiment, starting from only the noisy images. Therefore, we cannot gauge the overall effect of noise on our algorithm because we do not account for its contribution to the misspecification of rotations; we investigate the effect of noise on the algorithm only after the rotation estimation step. Moreover, we use simulated data instead of experimental data. The application of our algorithm to experimental datasets is left for a separate publication.

7.1. An appropriate definition of SNR—Generally, the definition of SNR is

$$P(\text{noise}) = \frac{N_{\text{res}}^2}{N^2} \sigma^2. \quad (7.2)$$

where P denotes power. In our setup, we will find appropriate definitions for both P (signal) and P (noise). Let us consider first the noise power. The standard definition is $P(\text{noise}) = \sigma^2$. However, note that in our case, the noise has a power of σ^2 in each pixel of an $N \times N$ grid, but we reconstruct the volumes to a bandlimit ω_{max} , corresponding to N_{res} . Hence, if we downsampled the $N \times N$ images to size $N_{\text{res}} \times N_{\text{res}}$, then we would still obey the Nyquist criterion (assuming the volumes actually are bandlimited by ω_{max}). This would have the effect of reducing the noise power by a factor of N^2/N_{res}^2 . Hence, in the context of our problem, we define

$$P(\text{signal}) = \frac{1}{n} \sum_{s=1}^n \frac{1}{q} \|P_s X_s\|^2. \quad (7.3)$$

Now, consider P (signal). In standard SPR, a working definition of signal power is

$$P(\text{signal}_{\text{het}}) = \frac{1}{n} \sum_{s=1}^n \frac{1}{q} \|P_s (X_s - \mu_0)\|^2. \quad (7.4)$$

However, in the case of the heterogeneity problem, the object we are trying to reconstruct is not the volume itself, but rather the deviation from the average volume, due to

heterogeneity. Thus, the relevant signal to us is not the images themselves, but the parts of the images that correspond to projections of the deviations of X_s from μ_0 . Hence, a natural definition of signal power in our case is

$$SNR_{het} = \frac{P(signal_{het})}{P(noise)} = \frac{\frac{1}{qn} \sum_{s=1}^n \|P_s(X_s - \mu_0)\|^2}{\sigma^2 N_{res}^2 / N^2}. \quad (7.5)$$

Using the above definitions, let us define SNR_{het} in our problem by

$$SNR = \frac{P(signal)}{P(noise)} = \frac{\frac{1}{n} \sum_{s=1}^n \frac{1}{q} \|P_s X_s\|^2}{\sigma^2 N_{res}^2 / N^2}. \quad (7.6)$$

Even with the correction factor N_{res}^2 / N^2 values are lower than the SNR values usually encountered in structural biology. Hence, we also define

$$X^c(r) = \sum_{i=1}^{M_c} a_{i,c} \exp\left(-\frac{\|r - r_{i,c}\|^2}{2\sigma_{i,c}^2}\right), \quad \tau_{i,c} \in R^3, a_{i,c}, \sigma_{i,c} \in R_+, \quad c=1, \dots, C. \quad (7.7)$$

We will present our numerical results primarily using SNR_{het} , but we will also provide the corresponding SNR values in parentheses.

To get a sense of the difference between this definition of SNR and the conventional one, compare the signal strength in a projection image to that in a mean-subtracted projection image in Figure 7.

7.2. Experimental procedure—We performed three numerical experiments: one with two heterogeneity classes, one with three heterogeneity classes, and one with continuous variation along the perimeter of a triangle defined by three volumes. The first two demonstrate our algorithm in the setup of Problem 1.2, and the third shows that we can estimate the covariance matrix and discover a low-dimensional structure in more general setups than the discrete heterogeneity case.

As a first step in each of the experiments, we created a number of phantoms analytically. We chose the phantoms to be linear combinations of Gaussian densities:

$$\begin{aligned} D_{z^*} \left(z^H a \right) &= a, \\ D_{z^*} \left(z^H A z \right) &= A z, \\ D_z \left(\text{tr} (A Z) \right) &= A, \\ D_z \left(\text{tr} \left(Z A Z^H A \right) \right) &= A Z^H A. \end{aligned} \quad (A.1)$$

For the discrete heterogeneity cases, we chose probabilities p_1, \dots, p_C and generated X_1, \dots, X_n by sampling from X^1, \dots, X^C accordingly. For the continuous heterogeneity case, we

generated each X_s by choosing a point uniformly at random from the perimeter of the triangle defined by X^1, X^2, X^3 .

For all of our experiments, we chose $n = 10000$, $N = 65$, $N_{\text{res}} = 17$, $K = 15$, and selected the set of rotations R_s to be approximately uniformly distributed on $SO(3)$. For each R_s , we calculated the clean continuous projection image $P_s X_s$ analytically, and then sampled the result on an $N \times N$ grid. Then, for each SNR level, we used (7.5) to find the noise power σ^2 to add to the images.

After simulating the data, we ran Algorithm 1 on the images I_s and rotations R_s on an Intel i7-3615QM CPU with 8 cores, and 8 GB of RAM. The runtime for the entire algorithm with the above parameter values (excluding precomputations) is 257 seconds. For the continuous heterogeneity case, we stopped the algorithm after computing the coordinates α_s (we did not attempt to reconstruct individual volumes in this case). To quantify the resolution of our reconstructions, we use the Fourier shell correlation (FSC), defined as the correlation of the reconstruction with the ground truth on each spherical shell in Fourier space [48]. For the discrete cases, we calculated FSC curves for the mean, the top eigenvectors, and the mean-subtracted reconstructed volumes. We also plotted the correlations of the mean, top eigenvectors, and mean-subtracted volumes with the corresponding ground truths for a range of SNR values. Finally, we plotted the coordinates α_s . For the continuous heterogeneity case, we tested the algorithm on only a few different SNR values. By plotting α_s in this case, we recover the triangle used in constructing X_s .

7.3. Experiment: Two classes—In this experiment, we constructed two phantoms X^1 and X^2 of the form (7.7), with $M_1 = 1$, $M_2 = 2$. Cross sections of X^1 and X^2 are depicted in the top row panels (c) and (d) in Figure 8. We chose the two heterogeneity classes to be equiprobable: $p_1 = p_2 = 1/2$. Note that the theoretical covariance matrix in the two-class heterogeneity problem has rank 1, with dominant eigenvector proportional to the difference between the two volumes.

Figure 8 shows the reconstructions of the mean, top eigenvector, and two volumes for $\text{SNR}_{\text{het}} = 0.013, 0.003, 0.0013$ (0.25, 0.056, 0.025). In Figure 9, we display eigenvalue histograms of the reconstructed covariance matrix for the above SNR values. Figure 10 shows the FSC curves for these reconstructions. Figure 11 shows the correlations of the computed means, top eigenvectors, and (mean-subtracted) volumes with their true values for a broader range of SNR values. In Figure 12, we plot a histogram of the coordinates α_s from step 7 of Algorithm 1.

Our algorithm was able to meaningfully reconstruct the two volumes for SNR_{het} as low as about 0.003 (0.06). Note that the means were always reconstructed with at least a 94% correlation to their true values. On the other hand, the eigenvector reconstruction shows a phase-transition behavior, with the transition occurring between SNR_{het} values of 0.001

Regarding the coefficients α_s depicted in Figure 12, note that in the noiseless case, there should be a distribution composed of two spikes. By adding noise to the images, the two spikes start blurring together. For SNR values up to a certain point, the distribution is still

visibly bimodal. However, after a threshold the two spikes coalesce into one. The proportions p_c are reliably estimated until this threshold.

7.4. Experiment: Three classes—In this experiment, we constructed three phantoms X^1, X^2, X^3 of the form (7.7), with $M_1 = 2, M_2 = 2, M_3 = 1$. The cross sections of X^1, X^2, X^3 are depicted in Figure 13 (top row, panels (d)–(f)). We chose the three classes to be equiprobable: $p_1 = p_2 = p_3 = 1/3$. Note that the theoretical covariance matrix in the three-class heterogeneity problem has rank 2.

Figures 13, 14, 15, 16, 17 are the three-class analogues of Figures 8, 9, 10, 11, 12 in the two-class case.

Qualitatively, we observe behavior similar to that in the two-class case. The mean is reconstructed with at least 90% accuracy for all SNR values considered, while both top eigen-vectors experience a phase-transition phenomenon (Figure 16(a)). As with the two-class case, we see that the disappearance of the eigengap coincides with the phase-transition behavior in the reconstruction of the top eigenvectors. However, in the three-class case we have two eigenvectors, and we see that the accuracy of the second eigenvector decays more quickly than that of the first eigenvector. This reflects the fact that the top eigenvalue of the true covariance Σ^{P^0} is 2.1×10^5 , while the second eigenvalue is 1.5×10^5 . These two eigenvalues differ because X^3 has greater norm than $X^2 - X^1$, which means that the two directions of variation have different associated variances. Hence, recovering the second eigenvector is less robust to noise. In particular, there are SNR values for which the top eigenvector can be recovered, but the second eigenvector cannot. $\text{SNR}_{\text{het}} = 0.0044$ (0.03) is such an example. We see in Figure 14 that for this SNR value, only the top eigenvector pops out of the bulk distribution. In this case, we would incorrectly estimate the rank of the true covariance as 1, and conclude that $C = 2$.

The coefficients a_s follow a similar trend to those in the two-class case. For high SNRs, there is a clearly defined clustering of the coordinates around three points, as in Figure 17(a). As the noise is increased, the three clusters become increasingly less defined. In Figure 17(b), we see that in this threshold case, the three clusters begin merging into one. As in the two-class case, this is the same threshold up to which the p_c are accurately estimated. By the time $\text{SNR} = 0.0044$ (0.03), there is no visible cluster separation, just as we observed in the two-class case. Although the SNR threshold for finding p_c from the a_s coefficients comes earlier than the one for the eigengap, the quality of volume reconstruction roughly tracks the quality of the eigenvector reconstruction. This suggests that the estimation of cluster means is more robust than that of the probabilities p_c .

7.5. Experiment: Continuous variation—In this experiment, we sampled X_s uniformly from the perimeter of the triangle determined by volumes X^1, X^2, X^3 (from the three-class discrete heterogeneity experiment). This setup is more suitable to model the case when the molecule can vary continuously between each pair X^i and X^j . Despite the fact this experiment does not fall under Problem 1.2, Figure 18 shows that we still recover the rank two structure. Indeed, it is clear that all the clean volumes still belong to a subspace of dimension 2. Moreover, we can see the triangular pattern of heterogeneity in the scatter plots

of α_s (Figure 19). However, note that once the images get moderately noisy, the triangular structure starts getting drowned out. Thus, in practice, without any prior assumptions, just looking at the scatter plots of α_s will not necessarily reveal the heterogeneity structure in the dataset. To detect continuous variation, a new algorithmic step must be designed to follow covariance matrix estimation. Nevertheless, this experiment shows that by solving the general Problem 1.1, we can estimate covariance matrices beyond those considered in the discrete case of the heterogeneity problem.

8. Discussion

In this paper, we proposed a covariance matrix estimator from noisy linearly projected data and proved its consistency. The covariance matrix approach to the cryo-EM heterogeneity problem is essentially a special case of the general statistical problem under consideration, but has its own practical challenges. We overcame these challenges and proposed a methodology to tractably estimate the covariance matrix and reconstruct the molecular volumes. We proved the consistency of our estimator in the cryo-EM case and also began the mathematical investigation of the projection covariance transform. We discovered that inverting the projection covariance transform involves applying the triangular area filter, a generalization of the ramp filter arising in tomography. Finally, we validated our methodology on simulated data, producing accurate reconstructions at low SNR levels. Our implementation of this algorithm is now part of the ASPIRE package at spr.math.princeton.edu. In what follows, we discuss several directions for future research.

As discussed in section 2.3, our statistical framework and estimators have opened many new questions in high-dimensional statistics. While a suite of results are already available for the traditional high-dimensional PCA problem, generalizing these results to the projected data case would require new random matrix analysis. Our numerical experiments in the cryo-EM case have shown many qualitative similarities between the estimated covariance matrix in the cryo-EM case and the sample covariance matrix in the spiked model. There is again a bulk distribution with eigenvalues separated from it. Moreover, there is a phase-transition phenomenon in the cryo-EM case, in which the top eigenvectors of the estimated covariance lose correlation with those of the population covariance once the corresponding eigenvalues are absorbed by the bulk distribution. Answering the questions posed in section 2.3 would be very useful in quantifying the theoretical limitations of our approach.

As an additional line of further inquiry, note that the optimization problem (2.4) for the covariance matrix is amenable to regularization. If $n \gg f(p, q)$ is the high-dimensional statistical regime in which the unregularized estimator still carries a signal, then of course we need regularization when $n \sim f(p, q)$. Here, f is a function depending on the distribution of the operators P_s . Moreover, regularization increases robustness to noise, so in applications like cryo-EM, this could prove useful. Tikhonov regularization does not increase the complexity of our algorithm, but has the potential to make L^{Pn} invertible. Under what conditions can we still achieve accurate recovery in a regularized setting? Other regularization schemes can take advantage of a priori knowledge of Σ_0 , such as using nuclear norm regularization in the case when Σ_0 is known to be low rank. See [25] for an application of nuclear norm minimization in the context of dealing with heterogeneity in

cryo-electron tomography. Another special structure Σ_0 might have is that it is sparse in a certain basis. For example, the localized variability assumption in the case of the heterogeneity problem is such an example; in this case, the covariance matrix is sparse in the real Cartesian basis or a wavelet basis. This sparsity can be encouraged using a matrix 1-norm regularization term. Other methods, such as sparse PCA [22] or covariance thresholding [7] might be applicable in certain cases when we have sparsity in a given basis.

We developed our algorithm in an idealized environment, assuming that the rotations R_s (and in-plane translations) are known exactly and correspond to approximately uniformly distributed viewing directions, and that the molecules belong to B . Moreover, we did not account for the CTF effect of the electron microscope. In practice, of course rotations and translations are estimated with some error. Also, certain molecules might exhibit a preference for a certain orientation, invalidating the uniform rotations assumption. Note that as long as L^{Pn} is invertible, our framework produces a valid estimator, but without the uniform rotations assumption, the computationally tractable approach to inverting this matrix proposed in section 5 no longer holds. Moreover, molecules might have higher frequencies than those we reconstruct, which could potentially lead to artifacts. Thus, an important direction of future research is to investigate the stability of our algorithm to perturbations from the idealized assumptions we have made. An alternative research direction is to devise numerical schemes to invert L^{Pn} without replacing it by $L\hat{C}$, which could allow incorporation of CTF and obviate the need to assume uniform rotations. We proposed one such scheme in section 5.1.

As we discussed in the introduction, our statistical problem (1.1) is actually a special case of the matrix sensing problem. In future work, it would be interesting to test matrix sensing algorithms on our problem. In the cryo-EM case, it would be useful to compare our approach with matrix sensing algorithms. It would also be interesting to explore the applications of our methodology to other tomographic problems involving variability. For example, the field of four-dimensional (4D) electron tomography focuses on reconstructing a 3D structure that is a function of time [26]. This 4D reconstruction is essentially a movie of the molecule in action.

The methods developed in this paper can in principle be used to estimate the covariance matrix of a molecule varying with time. This is another kind of “heterogeneity” that is amenable to the same analysis we used to investigate structural variability in cryo-EM.

Acknowledgments

E. Katsevich thanks Jane Zhao, Lanhui Wang, and Xiuyuan Cheng (PACM, Princeton University) for their valuable advice on several theoretical and practical issues. Parts of this work have appeared in E. Katsevich’s undergraduate Independent Work at Princeton University.

The authors are also indebted to Philippe Rigollet (ORFE, Princeton), as this work benefited from discussions with him regarding the statistical framework. Also, the authors thank Joachim Frank (Columbia University) and Joakim Anden (PACM, Princeton University) for providing helpful comments about their manuscript. They also thank Dr. Frank and Hstau Liao (Columbia University) for allowing them to reproduce Figure 2 from [29] as our Figure 1. Finally, they thank the editor and the referees for their many helpful comments.

The research of this author was partially supported by Award DMS-1115615 from NSF.

The research of this author was partially supported by Award R01GM090200 from the NIGMS, by Awards FA9550-12-1-0317 and FA9550-13-1-0076 from AFOSR, and by Award LTR DTD 06-05-2012 from the Simons Foundation.

Appendix A. Matrix derivative calculations

The goal of this appendix is to differentiate the objective functions of (2.3) and (2.4) to verify formulas (2.5) and (2.6). In order to differentiate with respect to vectors and matrices, we appeal to a few results from [17]. The results are as follows:

$$\|I_s - P_s \mu\|^2 = \left(I_s^H - \mu^H P_s^H \right) (I_s - P_s \mu) = \mu^H P_s^H P_s \mu - \mu^H P_s^H I_s - I_s^H P_s \mu + \text{const} \quad (\text{A.2})$$

Here, the lowercase letters represent vectors and the uppercase letters represent matrices. Also note that z^* denotes the complex conjugate of z . The general term of (2.3) is

$$D_{\mu^*} \|I_s - P_s \mu\|^2 = P_s^H P_s \mu - P_s^H I_s. \quad (\text{A.3})$$

We can differentiate this with respect to μ^* by using the first two formulas of (A.1). We get

$$\begin{aligned} & \| (I_s - P_s \mu_n) (I_s - P_s \mu_n)^H - (P_s \Sigma P_s^H + \sigma^2 I) \|_F^2 \\ &= \| A_s - P_s \Sigma P_s^H \|_F^2 \\ &= \text{tr} \left(A_s^H - P_s \Sigma^H P_s^H \right) \left(A_s - P_s \Sigma P_s^H \right) \\ &= \text{tr} \left(P_s \Sigma^H P_s^H P_s \Sigma P_s^H \right) - \text{tr} \left(P_s \Sigma^H P_s^H A_s \right) - \text{tr} \left(A_s^H P_s \Sigma P_s^H \right) + \text{const}, \\ &= \text{tr} \left(\Sigma P_s^H P_s \Sigma^H P_s^H P_s \right) - \text{tr} \left(P_s^H A_s P_s \Sigma^H \right) - \text{tr} \left(P_s^H A_s^H P_s \Sigma \right) + \text{const}, \end{aligned}$$

Summing in s gives us (2.5).

If we let $A_s = (I_s - P_s \mu_n)(I_s - P_s \mu_n)^H - \sigma^2 I$, then the general term of (2.4) is

$$P_s^H P_s \Sigma^H P_s^H P_s - P_s^H A_s^H P_s.$$

Using the last two formulas of (A.1), we find that the derivative of this expression with respect to Σ is

$$E[Y_s] = 0 \quad \text{and} \quad \|Y_s\| \leq R \quad \text{almost surely.} \quad (\text{B.1})$$

Taking a Hermitian and summing in s gives us (2.6).

Appendix B. Consistency of μ_n and Σ_n

In this appendix, we will prove the consistency results about μ_n and Σ_n stated in section 2.2. Recall μ_n and Σ_n are defined nontrivially if $I_A^{-1} \geq 2 I_A^{-1}$ and $I_L^{-1} \geq 2 I_L^{-1}$. As a necessary step towards our consistency results, we must first prove that the probability of

these events tends to 1 as $n \rightarrow \infty$. Such a statement follow from a matrix concentration argument based on Bernstein's inequality [59, Theorem 1.4], which we reproduce here for the reader's convenience as a lemma.

Lemma B.1 (matrix Bernstein's inequality)

Consider a finite sequence \mathbf{Y}_s of independent, random, self-adjoint matrices with dimension p . Assume that each random matrix satisfies

$$P \left\{ \left\| \sum_s \mathbf{Y}_s \right\| \geq t \right\} \leq p \cdot \exp \left(\frac{-t^2/2}{\sigma^2 + Rt/3} \right), \quad \text{where } \sigma^2 := \left\| \sum_s E \left(\mathbf{Y}_s^2 \right) \right\|. \quad (\text{B.2})$$

Then, for all $t > 0$,

$$P \left\{ \left\| \frac{1}{n} \sum_{s=1}^n \mathbf{Z}_s - E[Z] \right\| \geq t \right\} \leq d \exp \left(\frac{-3nt^2}{6B^2 + 4Bt} \right). \quad (\text{B.3})$$

Next, we prove another lemma, which is essentially the Bernstein inequality in a more convenient form.

Lemma B.2

Let \mathbf{Z} be a symmetric $d \times d$ random matrix, with $\|\mathbf{Z}\| \leq B$ almost surely. If $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ are i.i.d. samples from \mathbf{Z} , then

$$E \left\| \frac{1}{n} \sum_{s=1}^n \mathbf{Z}_s - E[Z] \right\| \leq CB \max \left(\sqrt{\frac{\log d}{n}}, \frac{2 \log d}{n} \right), \quad (\text{B.4})$$

Moreover,

$$\|\mathbf{Y}_s\| \leq \frac{1}{n} (\|\mathbf{Z}_s\| + E[\|\mathbf{Z}\|]) \leq \frac{2B}{n} =: R \quad \text{almost surely.} \quad (\text{B.5})$$

where C is an absolute constant.

Proof

The proof is an application of the matrix Bernstein inequality. Let $\mathbf{Y}_s = \frac{1}{n} (\mathbf{Z}_s - E\mathbf{Z})$. Then, note that $E[\mathbf{Y}_s] = 0$ and

$$E \left[\mathbf{Y}_s^2 \right] = \frac{1}{n^2} E \left[\mathbf{Z}_s^2 - \mathbf{Z}_s E[Z] - E[Z] \mathbf{Z}_s + E[Z]^2 \right] = \frac{1}{n^2} \left(E \left[\mathbf{Z}_s^2 \right] - E[Z]^2 \right) \leq \frac{1}{n^2} E \left[\mathbf{Z}_s^2 \right]. \quad (\text{B.6})$$

Next, we have

$$\sigma^2 := \left\| \sum_{s=1}^n E \left[Y_s^2 \right] \right\| \leq \sum_{s=1}^n \|E \left[Y_s^2 \right]\| \leq \sum_{s=1}^n \frac{1}{n^2} \|E \left[Z_s^2 \right]\| \leq \sum_{s=1}^n \frac{1}{n^2} E \left[\|Z_s\|^2 \right] \leq \frac{B^2}{n}. \quad (\text{B.7})$$

It follows that

$$P \left\{ \left\| \frac{1}{n} \sum_{s=1}^n Z_s - E[Z] \right\| \geq t \right\} = P \left\{ \left\| \sum_{s=1}^n Y_s \right\| \geq nt \right\} \leq \text{dexp} \left(\frac{-t^2/2}{\sigma^2 + Rt/3} \right) \leq \text{dexp} \left(\frac{-3nt^2}{6B^2 + 4Bt} \right). \quad (\text{B.8})$$

Now, by the matrix Bernstein inequality, we find that

$$P \{ \|A_n - A\| \leq t \} \leq p \exp \left(\frac{-3nt^2}{6B_p^4 + 4B_p^2 t} \right). \quad (\text{B.9})$$

$$E \|A_n - A\| \leq C B_p^2 \max \left(\sqrt{\frac{\log p}{n}}, \frac{2 \log p}{n} \right) = C B_p^2 \sqrt{\frac{\log p}{n}}, \quad (\text{B.10})$$

This proves (B.3). The bound (B.4) follows from [59, Remark 6.5].

$\mathbf{P}^{H\mathbf{P}_s}$, where $\mathbf{P}_1, \dots, \mathbf{P}_n$ are i.i.d. samples from \mathbf{P} . Then,

Corollary B.3

Let \mathbf{P} be a random $q \times p$ matrix such that $\|\mathbf{P}\| \leq B_P$ almost surely. Let $A = E[\mathbf{P}^{H\mathbf{P}}]$ and let

$$P \{ \|L_n - L\| \geq t \} \leq p^2 \exp \left(\frac{-3nt^2}{6q^4 B_p^8 + 4q^2 B_p^4 t} \right). \quad (\text{B.11})$$

Moreover,

$$E \|L_n - L\| \leq C q^2 B_p^4 \max \left(\sqrt{\frac{2 \log p}{n}}, \frac{4 \log p}{n} \right) = C_q^2 B_p^4 \sqrt{\frac{2 \log p}{n}}, \quad (\text{B.12})$$

where the last equality holds if $n \geq 4 \log p$.

Proof

These bounds follow by letting $\mathbf{Z} = \mathbf{P}^{H\mathbf{P}}$ in Lemma B.2 and noting that $\|\mathbf{Z}\| \leq B^2$ almost surely.

Corollary B.4

Let \mathbf{P} be a random $q \times p$ matrix such that $\mathbb{P} \{ \|\mathbf{P}\|_F \leq B_P \}$ almost surely. Let

$L\Sigma = E \left[P^H P \Sigma P^H P \right]$ and let $L_n \Sigma = \frac{1}{n} \sum_{s=1}^n P_s^H P_s \Sigma P_s^H P_s$, where $\mathbf{P}_1, \dots, \mathbf{P}_n$ are i.i.d. samples from \mathbf{P} . Then,

$$\begin{aligned} \|Z\| &= \max_{\|\text{vec}(\Sigma)\|=1} \|Z_{\text{vec}}(\Sigma)\| = \max_{\|\Sigma\|_F=1} \|Z\Sigma\|_F \\ &= \max_{\|\Sigma\|_F=1} \|P^H P \Sigma P^H P\|_F \leq \|P\|_F^4 \leq q^2 \|P\|^4 \leq q^2 B_P^4. \end{aligned} \quad (\text{B.13})$$

$$P \left[\varepsilon_n^A \geq 1 - \alpha_n^A, \quad P \left[\varepsilon_n^L \geq 1 - \alpha_n^L \right], \quad (\text{B.14})$$

Moreover,

where the last equality holds if $n \geq 8 \log p$.

Proof

We wish to apply Lemma B.2 again, this time for $Z\Sigma = P^H P \Sigma P^H P$. In this case we must be careful because Z is an operator on the space of $p \times p$ matrices. We can view it as a $p^2 \times p^2$ matrix if we represent its argument (a $p \times p$ matrix Σ) as a vector of length p^2 (denoted by $\text{vec}(\Sigma)$). Then, almost surely,

$$\alpha_n^A = p \exp \left(\frac{-3n \lambda_{\min}(A)^2 / 4}{6B_P^4 + 2B_P^2 \lambda_{\min}(A)} \right) \quad \text{and} \quad \alpha_n^L = p^2 \exp \left(\frac{-3n \lambda_{\min}(L)^2 / 4}{6q^4 B_P^8 + 2q^2 B_P^4 \lambda_{\min}(L)} \right). \quad (\text{B.15})$$

In the penultimate inequality above we used the fact that $\|A\|_F \leq \sqrt{\text{rank}(A)} \|A\|$ for an arbitrary matrix A . Now, (B.11) follows from (B.3) by setting $B = q^2 B_P^4$ and $d = p^2$.

Proposition B.5

Let E^A be the event that $\|A_n^{-1}\| \leq 2 \|A^{-1}\|$, and let E^L be the event that

$$P \left[\|A_n^{-1}\| > 2 \|A^{-1}\| \right] = P \left[\lambda_{\min}(A_n) < \frac{1}{2} \lambda_{\min}(A) \right] \leq P \left[\|A_n - A\| > \frac{1}{2} \lambda_{\min}(A) \right]. \quad (\text{B.16})$$

where

$$P \left[\|A_n - A\| > \frac{1}{2} \lambda_{\min}(A) \right] \leq p \exp \left(\frac{-3n \lambda_{\min}(A)^2 / 4}{6B_P^4 + 2B_P^2 \lambda_{\min}(A)} \right) = \alpha_n^A. \quad (\text{B.17})$$

Proof

Note that $\lambda_{\min}(A_n) \leq \lambda_{\min}(A) - \|A_n - A\|$. It follows that

$$P \left[\|L_n - L\| > \frac{1}{2} \lambda_{\min}(L) \right] \leq p^2 \exp \left(\frac{-3n \lambda_{\min}(L)^2 / 4}{6q^4 B_p^8 + 2q^2 B_p^4 \lambda_{\min}(L)} \right) = \alpha_n^A. \quad (\text{B.18})$$

By Corollary B.3, it follows that

$$B_I^2 := E \left[\|I - P\mu_0\|^2 \right]. \quad (\text{B.19})$$

Analogously, Corollary B.4 implies that

$$B_I^2 \leq B_p^2 E \left[\|X - \mu_0\|^2 \right] + E[\|E\|^2]. \quad (\text{B.20})$$

Now, we prove the consistency results, which we restate for convenience. In the following propositions, define

$$\|V\|_m = E[\|V - E[V]\|^m]^{\frac{1}{m}}, \quad (\text{B.21})$$

Note that

$$E\|\mu_n - \mu_0\| = O\left(\frac{1}{\sqrt{n}}\right). \quad (\text{B.22})$$

Also, recall the following notation introduced in section 2.2:

$$\begin{aligned} E[\|\mu_n - \mu_0\|] &= P \left[\epsilon_n^A \right] E \left[\|\mu_n - \mu_0\| \mid \epsilon_n^A \right] + \left(1 - P \left[\epsilon_n^A \right] \right) E \left[\|\mu_n - \mu_0\| \mid \overline{\epsilon_n^A} \right] \\ &\leq P \left[\epsilon_n^A \right] E \left[\|A_n^{-1} b_n - \mu_0\| \mid \epsilon_n^A \right] + \alpha_n^A \|\mu_0\| \\ &\leq P \left[\epsilon_n^A \right] E \left[\|A_n^{-1} (b_n - A_n \mu_0)\| \mid \epsilon_n^A \right] + \alpha_n^A \|\mu_0\| \\ &\leq P \left[\epsilon_n^A \right] a \|A^{-1}\| E \left[\|b_n - A_n \mu_0\| \mid \epsilon_n^A \right] + \alpha_n^A \|\mu_0\| \\ &\leq 2 \|A^{-1}\| E \left[\|b_n - A_n \mu_0\| \right] + \alpha_n^A \|\mu_0\|. \end{aligned} \quad (\text{B.23})$$

where \mathbf{V} is a random vector. For example, (B.20) can be written as $B_I^2 \leq B_p^2 \|X\|_2^2 + \|E\|_2^2$

Proposition B.6

Suppose A (defined in (2.10)) is invertible, that $\mathbf{IP} 1$ B_p almost surely, and that $\|X\|_2, \|E\|_2 < \infty$. Then, for fixed p, q we have

$$E \left[\|b_n - A_n \mu_0\|^2 \right] \leq E \left[\|b_n - A_n \mu_0\|^2 \right] = \frac{1}{n} E \left[\|P^H (I - P\mu_0)\|^2 \right] \leq \frac{1}{n} B_p^2 B_I^2. \quad (\text{B.24})$$

Hence, under these assumptions, μ_n is consistent.

Proof

Since $P[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\| \leq t] \leq t^{-1} E[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\|]$ by Markov's inequality, it is sufficient to prove that $E[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\|] \rightarrow 0$ as $n \rightarrow \infty$. Note that by the definition of $\boldsymbol{\mu}_n$ and Proposition B.5,

$$E[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\|] \leq \frac{2\|A^{-1}\|B_P B_I + \alpha_n^A \|\boldsymbol{\mu}_0\|}{\sqrt{n}}. \quad (\text{B.25})$$

$$P[\epsilon_n^A] E[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\|^2 | \epsilon_n^A] \leq \frac{4\|A^{-1}\|^2}{n} B_P^2 B_I^2. \quad (\text{B.26})$$

where these summands are i.i.d., we find

Since

$$E\left\|\frac{1}{n} \sum_{s=1}^n V_s V_s^H - \Sigma_V\right\| \leq C \|\Sigma_V\| \|\Sigma_V^{-1/2}\| \frac{\sqrt{\log p}}{\sqrt{n}} \left(E\|V\|^{\log n}\right)^{1/\log n}, \quad (\text{B.27})$$

Putting together what we have, we arrive at

$$\|X\|_j, \|E\|_j \leq Q(j), \quad j \in N. \quad (\text{B.28})$$

Inspecting this bound reveals that $E[\|\boldsymbol{\mu}_n - \boldsymbol{\mu}_0\|] \rightarrow 0$ as $n \rightarrow \infty$, as needed.

Remark B.7

Note that with a simple modification to the above argument, we obtain

$$E\|\Sigma_n - \Sigma_0\| = O\left(\frac{Q(\log n)}{\sqrt{n}}\right). \quad (\text{B.29})$$

This bound will be useful later.

Before proving the consistency of Σ_n , we state a lemma.

Lemma B.8

Let \mathbf{V} be a random vector on C^p with $E[\mathbf{V}\mathbf{V}^H] = \Sigma_V$, and let $\mathbf{V}_1, \dots, \mathbf{V}_n$ be i.i.d. samples from \mathbf{V} . Then, for some absolute constant C ,

$$\begin{aligned}
E[\|\Sigma_n - \Sigma_0\|] &= P[\epsilon_n^A \cap \epsilon_n^L] E[\|\Sigma_n - \Sigma_0\| | \epsilon_n^A \cap \epsilon_n^L] + (1 - P[\epsilon_n^A \cap \epsilon_n^L]) E[\|\Sigma_n - \Sigma_0\| | \overline{\epsilon_n^A \cap \epsilon_n^L}] \\
&\leq P[\epsilon_n^A \cap \epsilon_n^L] E[\|L_n^{-1} B_n - \Sigma_0\| | \epsilon_n^A \cap \epsilon_n^L] + (\alpha_n^A + \alpha_n^L) \|\Sigma_0\| \\
&\leq P[\epsilon_n^A \cap \epsilon_n^L] E[\|L_n^{-1} B_n - L_n \Sigma_0\| | \epsilon_n^A \cap \epsilon_n^L] + (\alpha_n^A + \alpha_n^L) \|\Sigma_0\| \\
&\leq 2\|L^{-1}\| P[\epsilon_n^A \cap \epsilon_n^L] E[\|(B_n - L_n \Sigma_0)\| | \epsilon_n^A \cap \epsilon_n^L] + (\alpha_n^A + \alpha_n^L) \|\Sigma_0\| \\
&\leq 2\|L^{-1}\| P[\epsilon_n^A] E[\|B_n - L_n \Sigma_0\| | \epsilon_n^A] + (\alpha_n^A + \alpha_n^L) \|\Sigma_0\|.
\end{aligned} \tag{B. 30}$$

provided the RHS does not exceed $\|\Sigma_V\|$.

Proof

This result is a simple modification of [47, Theorem 1].

Proposition B.9

Suppose A and L (defined in 2.10) are invertible, that $\mathbf{IP} 1$ B_p almost surely, and that there is a polynomial Q for which

$$\begin{aligned}
B_n - L_n \Sigma_0 &= \left(\frac{1}{n} \sum_{s=1}^n P_s^H (I_s - P_s \mu_n) (I_s - P_s \mu_n)^H P_s - \frac{1}{n} \sum_{s=1}^n P_s^H (I_s - P_s \mu_0) (I_s - P_s \mu_0)^H P_s \right) \\
&+ \left(\frac{1}{n} \sum_{s=1}^n P_s^H (I_s - P_s \mu_0) (I_s - P_s \mu_0)^H P_s - (\sigma^2 A + L \Sigma_0) \right) + \sigma^2 (A - A_n) + (L - L_n) \Sigma_0 \\
&=: D_1 + D_2 + D_3 + D_4.
\end{aligned} \tag{B. 31}$$

Then, for fixed p, q , we have

$$E[\|D_1\| | \epsilon_n^A] \leq B_p^2 \frac{1}{n} \sum_{s=1}^n E[\|(I_s - P_s \mu_n) (I_s - P_s \mu_n)^H - (I_s - P_s \mu_0) (I_s - P_s \mu_0)^H\| | \epsilon_n^A]. \tag{B. 32}$$

Hence, under these assumptions, Σ_n is consistent.

Proof

In parallel to the proof of Proposition B.6, we will prove that $E[\|\Sigma_n - \Sigma_0\|] \rightarrow 0$ as $n \rightarrow \infty$.

We compute

$$\begin{aligned}
&(I_s - P_s \mu_n) (I_s - P_s \mu_n)^H - (I_s - P_s \mu_0) (I_s - P_s \mu_0)^H \\
&= \{(I_s - P_s \mu_0) + P_s (\mu_0 - \mu_n)\} \{(I_s - P_s \mu_0) + P_s (\mu_0 - \mu_n)\}^H - (I_s - P_s \mu_0) (I_s - P_s \mu_0)^H \\
&= (I_s - P_s \mu_0) (\mu_0 - \mu_n)^H P_s^H + P_s (\mu_0 - \mu_n) (I_s - P_s \mu_0)^H + P_s (\mu_0 - \mu_n) (\mu_0 - \mu_n)^H P_s^H.
\end{aligned} \tag{B. 33}$$

Now, we will bound $E E[\|B_n - L_n \Sigma_n\| | \epsilon_n^A]$. To do this, we write

$$\begin{aligned}
& E \left[\left\| (I_s - P_s \mu_0) (\mu_0 - \mu_n)^H P_s^H \right\| \epsilon_n^A \right]^2 \\
& \leq B_P^2 E \left[\left\| I_s - P_s \mu_0 \right\| \left\| \mu_0 - \mu_n \right\| \epsilon_n^A \right]^2 \\
& \leq B_P^2 E \left[\left\| I_s - P_s \mu_0 \right\|^2 \epsilon_n^A \right] E \left[\left\| \mu_0 - \mu_n \right\|^2 \epsilon_n^A \right] \quad (\text{B.34}) \\
& \leq \frac{4 \|A^{-1}\|^2}{nP[\epsilon_n^A]^2} B_P^4 B_I^4.
\end{aligned}$$

Let us consider each of these four difference terms in order. Note that

$$E \left[\left\| P_s (\mu_0 - \mu_n) (\mu_0 - \mu_n)^H P_s^H \right\| \epsilon_n^A \right] \leq B_P^2 E \left[\left\| \mu_0 - \mu_n \right\|^2 \epsilon_n^A \right] \leq \frac{4 \|A^{-1}\|^2}{nP[\epsilon_n^A]} B_P^4 B_I^2. \quad (\text{B.35})$$

Moreover,

$$\begin{aligned}
P \left[\epsilon_n^A \right] E \left[\left\| D_1 \right\| \epsilon_n^A \right] & \leq P \left[\epsilon_n^A \right] B_P^2 \left(2 \frac{\|A^{-1}\|}{\sqrt{nP[\epsilon_n^A]}} B_P^2 B_I^2 + 4 \frac{\|A^{-1}\|^2}{nP[\epsilon_n^A]} B_P^4 B_I^2 \right) \\
& = \frac{4 B_P^4 B_I^2 \|A^{-1}\|}{n} \left(\sqrt{n} + \|A^{-1}\| B_P^2 \right). \quad (\text{B.36})
\end{aligned}$$

Using the Cauchy–Schwarz inequality and (B.26), we find

$$\Sigma_V = E \left[V V^H \right] = E \left[P^H P (X - \mu_0) (X - \mu_0)^H P^H P \right] + E \left[P^H E E^H P \right] = L \Sigma_0 + \sigma^2 A. \quad (\text{B.37})$$

Here, we used (B.26). This bound also holds for the second term in the last line of (B.33). As for the third term,

$$P \left[\epsilon_n^A \right] E \left[\left\| D_2 \right\| \epsilon_n^A \right] \leq E \left[\frac{1}{n} \sum_{s=1}^n \left\| V_s V_s^H - \Sigma_V \right\| \right] \leq C \|\Sigma_V\| \|\Sigma_V^{-1}\|^{1/2} \frac{\sqrt{\log p}}{\sqrt{n}} \left(E \|V\| \log n \right)^{\frac{1}{\log n}}. \quad (\text{B.38})$$

Putting these bounds together, we arrive at

$$\|L \Sigma_0\| \leq \|L \Sigma_0\|_F \leq q^4 B_P^4 \|\Sigma_0\|_F \leq q^4 B_P^4 \sqrt{\text{rank}(\Sigma_0)} \|\Sigma_0\| q^4 B_P^4 \sqrt{\text{rank}(\Sigma_0)} \|X\|_2^2. \quad (\text{B.39})$$

$$\begin{aligned}
\left(E \|V\| \log n \right)^{\frac{1}{\log n}} & \leq B_p \left(B_p \left(E \left[\|X - \mu_0\| \log n \right]^{\frac{1}{\log n}} + E \left[\|E\| \log n \right]^{\frac{1}{\log n}} \right) \right) \\
& = B_p \left(B_p \left(\|X\| \log n + \|E\| \log n \right) \right). \quad (\text{B.40})
\end{aligned}$$

Next, we move on to analyzing D_2 . If $V = P^H (I - P \mu_0)$, note that

$$\begin{aligned}
P \left[\epsilon_n^A \right] E \left[\left\| D_2 \right\| \epsilon_n^A \right] & \leq C B_P^3 \left(q^4 B_P^2 \sqrt{\text{rank}(\Sigma_0)} \|X\|_2^2 + \sigma^2 \right) \\
& \quad \times \left\| (L \Sigma_0 + \sigma^2 A)^{-1/2} \right\| \frac{\sqrt{\log p}}{\sqrt{n}} \left(B_p \left(\|X\| \log n + \|E\| \log n \right) \right). \quad (\text{B.41})
\end{aligned}$$

By Lemma (B.8), we find (B.38)

Since $\Sigma_0 = E[(X - \mu_0)(X - \mu_0)^H]$, it follows that $\|\Sigma_0\| = E[\|X - \mu_0\|^2] = \|X\|_2$. Further, the calculation (B.13) implies that

$$P \left[\epsilon_n^A \right] E \left[\|D_3\| \mid \epsilon_n^A \right] \leq E [\|D_3\|] = \sigma^2 E [\|A - A_n\|] \leq \sigma^2 C' B_P^2 \sqrt{\frac{\log n}{p}}. \quad (\text{B.42})$$

Also, it is clear that $\|A\| \leq B_P^2$. Furthermore, Minkowski inequality implies that

$$\begin{aligned} P \left[\epsilon_n^A \right] E \left[\|D_4\| \mid \epsilon_n^A \right] &\leq E [\|D_4\|] \leq E [\|L - L_n\|] \|\Sigma_0\|_F \\ &\leq \sigma^2 C' q^2 B_P^4 \sqrt{\frac{2 \log p}{n}} \sqrt{\text{rank}(\Sigma_0)} \|X\|_2^2. \end{aligned} \quad (\text{B.43})$$

Hence, (B.38) becomes

$$\begin{aligned} E [\|\Sigma_n - \Sigma_0\|] &\leq 2 \|L^{-1}\| \frac{4 B_P^4 (B_P^2 \|X\|_2^2 + \|E\|_2^2) \|A^{-1}\|}{n} \left(\sqrt{n} + \|A^{-1}\| B_P^2 \right) \\ &\quad + C B_P^3 \left(q^4 B_P^2 \sqrt{\text{rank}(\Sigma_0)} \|X\|_2^2 + \sigma^2 \right) \|(L \Sigma_0 + \sigma^2 A)^{-1/2}\| \\ &\quad \times \frac{\sqrt{\log p}}{\sqrt{n}} (B_P (\|X\|_{\log n} + \|E\|_{\log n})) \\ &\quad + \sigma^2 C' B_P^2 \sqrt{\frac{\log p}{n}} + \sigma^2 C' q^2 B_P^4 \sqrt{\frac{2 \log p}{n}} \sqrt{\text{rank}(\Sigma_0)} \|X\|_2^2 \\ &\quad + (\alpha_n^A + \alpha_n^L) \|X\|_2^2. \end{aligned} \quad (\text{B.44})$$

Next, a bound for D_3 follows immediately from (B.10):

$$\hat{L}_{i_1 i_2, j_1 j_2}^{k_1, k_2} = \int_{S^2 \times S^2} \left(a_{j_1}^{k_1} \otimes a_{j_2}^{k_2} \right) (\alpha, \beta) \overline{\left(a_{j_1}^{k_1} \otimes a_{j_2}^{k_2} \right) (\alpha, \beta)} K(\alpha, \beta) d\alpha d\beta. \quad (\text{C.1})$$

Similarly, (B.12) gives

$$\hat{L}_{i_1 i_2, j_1 j_2}^{k_1, k_2} = \int_{S^2 \times S^2} A_{i_1 j_1}^{k_1}(\alpha) \overline{A_{i_2 j_2}^{k_2}(\beta)} K(\alpha, \beta) d\alpha d\beta. \quad (\text{C.2})$$

Combining the four bounds (B.36), (B.39), (B.42), (B.43) with (B.30) and (B.31), we arrive at

$$A_{i_1 j_1}^{k_1}(\alpha) = \sum_{l=0}^{2k_1} \sum_{|m| \leq l} C_{l,m} \left(A_{i_1 j_1}^{k_1} \right) Y_l^m(\alpha), \quad A_{i_2 j_2}^{k_2}(\beta) = \sum_{l=0}^{2k_2} \sum_{|m'| \leq l'} C_{l,m'} \left(A_{i_2 j_2}^{k_2} \right) Y_{l'}^{m'}(\beta). \quad (\text{C.3})$$

Fixing all the variables except n , we see that the largest term is the one in the second line, and it decays as $Q(\log n)/n$ due to the moment growth condition (B.28).

Appendix C. Simplifying (5.12)

Here, we simplify the expression for an element of L^{Pk_1, k_2} :

$$\hat{L}_{i_1 i_2, j_1 j_2}^{k_1, k_2} = \sum_{l, m, l', m'} C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)} \int_{S^2} \int_{S^2} Y_l^m(\alpha) K(\alpha, \beta) \overline{Y_{l'}^{m'}(\beta)} d\alpha d\beta. \quad (C.4)$$

Let $A_{i, j}^k = \overline{a_i^k} a_j^k$. Then, (C.1) becomes

$$\int_{S^2} Y_l^m(\alpha) K(\alpha, \beta) d\alpha = c(l) Y_l^m(\beta), \quad (C.5)$$

Recall from section 5.3 that a_i^k is a spherical harmonic of order up to k . It follows that $A_{i_1 j_1}^{k_1}$ has a spherical harmonic expansion up to order $2k_1$ (using the formula for the product of two spherical harmonics, which involves the Clebsch–Gordan coefficients). The same holds for $A_{i_2 j_2}^{k_2}$ where the order goes up to $2k_2$. Let us write $C_l^m \left(A_{ij}^k \right)$ for the l, m coefficient of the spherical harmonic expansion of A_{ij}^k . Thus, we have

$$c(l) = \frac{2\pi}{P_l(1)} \int_{-1}^1 K(t) P_l(t) dt. \quad (C.6)$$

It follows that

$$c(l) = 2 \int_0^1 \frac{1}{\sqrt{1-t^2}} P_l(t) dt. \quad (C.7)$$

Since $K(\alpha, \beta)$ depends only on $\alpha \cdot \beta$, by an abuse of notation we can write $K(\alpha, \beta) = K(\alpha \cdot \beta)$. Thus, the Funk–Hecke theorem applies [38], so we may write

$$c(l) = 2 \int_0^1 \frac{1}{\sqrt{1-t^2}} P_l(t) dt = \pi \left(\frac{l!}{2^l \left(\frac{l!}{2}\right)^2} \right)^2. \quad (C.8)$$

where

$$\begin{aligned} \hat{L}_{i_1 i_2, j_1 j_2}^{k_1 k_2} &= \sum_{l, m, l', m'} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)} \int_{S^2} Y_l^m(\beta) \overline{Y_{l'}^{m'}(\beta)} d\beta \\ &= \sum_{l, m} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)}. \end{aligned} \quad (C.9)$$

Note that the P_R are the Legendre polynomials. Since K is an even function of t and P_R has the same parity as $.e$, it follows that $c(.e) = 0$ for odd $.e$. For even $.e$, we have

$$\begin{aligned}\hat{L}_{i_1 i_2, j_1 j_2}^{k_1 k_2} &= \sum_{l, m} \sum_{l', m'} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)} \int_{S^2} Y_l^m(\beta) \overline{Y_{l'}^{m'}(\beta)} d\beta \\ &= \sum_{l, m} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)}.\end{aligned}\tag{C.9}$$

It follows from formula 3 on p. 423 of [45] that

$$\begin{aligned}\hat{L}_{i_1 i_2, j_1 j_2}^{k_1 k_2} &= \sum_{l, m} \sum_{l', m'} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)} \int_{S^2} Y_l^m(\beta) \overline{Y_{l'}^{m'}(\beta)} d\beta \\ &= \sum_{l, m} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)}.\end{aligned}\tag{C.9}$$

Using Stirling's formula, we can find that $c(e) \sim .e^{-1}$ for large $.e$.

Finally, plugging the result of Funk–Hecke into (C.4), we obtain

$$\begin{aligned}\hat{L}_{i_1 i_2, j_1 j_2}^{k_1 k_2} &= \sum_{l, m} \sum_{l', m'} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)} \int_{S^2} Y_l^m(\beta) \overline{Y_{l'}^{m'}(\beta)} d\beta \\ &= \sum_{l, m} c(l) C_{l, m} \left(A_{i_1 j_1}^{k_1} \right) \overline{C_{l', m'} \left(A_{i_2 j_2}^{k_2} \right)}.\end{aligned}\tag{C.9}$$

Thus, we have verified (5.13).

REFERENCES

- [1]. Amunts A, Brown A, Bai X, Llaácer J, Hussain T, Emsley P, Long F, Murshudov G, Scheres S, Ramakrishnan V. Structure of the yeast mitochondrial large ribosomal subunit. *Science*. 2014; 343:1485–1489. [PubMed: 24675956]
- [2]. Baddour N. Operational and convolution properties of three dimensional Fourier transforms in spherical polar coordinates. *J. Opt. Soc. Amer. A*. 2010; 27:2144–2155.
- [3]. Bai X, Fernandez I, McMullan G, Scheres S. Ribosome structures to near-atomic resolution from thirty thousand cryo-em particles. *eLife*. 2013; 2:e00461. [PubMed: 23427024]
- [4]. Baik J, Ben Arous G, Páecháe S. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* 2005; 33:1643–1697.
- [5]. Baik J, Silverstein JW. Eigenvalues of large sample covariance matrices of spiked population models. *J. Multivariate Anal.* 2006; 97:1382–1408.
- [6]. Bennett, J.; Lanning, S. The Netflix prize. 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; San Jose, CA, ACM, New York. 2007.
- [7]. Bickel PJ, Levina E. Covariance regularization by thresholding. *Ann. Statist.* 2008; 36:2577–2604.
- [8]. Bishop, C. *Inf. Sci. Statist.* Springer-Verlag; New York: 2006. Pattern Recognition and Machine Learning.
- [9]. Candes E, Plan Y. Matrix completion with noise. *Proc. IEEE*. 2010; 98:925–936.
- [10]. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 1977; 39:1–38.
- [11]. Donoho, D. *Math Challenges of the 21st Century*. Los Angeles: 2000. High-dimensional data analysis: The curses and blessings of dimensionality.
- [12]. Frank, J. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. Oxford University Press; Oxford: 2006.

- [13]. Frank, J. Exploring the Dynamics of Supramolecular Machines with Cryo-Electron Microscopy. Proceedings of the 23rd International Solvay Conference on Chemistry; Brussels: International Solvay Institutes; 2013.
- [14]. Frank J. Story in a sample – the potential (and limitations) of cryo-electron microscopy applied to molecular machines. *Biopolymers*. 2013; 99:832–836. [PubMed: 23640776]
- [15]. Henderson R. Realizing the potential of electron cryo-microscopy. *Quart. Rev. Biophys.* 2004; 37:3–13.
- [16]. Herman G, Kalinowski M. Classification of heterogeneous electron microscopic projections into homogeneous subsets. *Ultramicroscopy*. 2008; 108:327–338. [PubMed: 17574340]
- [17]. Hjørungnes A, Gesbert D. Complex-valued matrix differentiation: Techniques and key results. *IEEE Trans. Signal Process.* 2007; 55:2740–2746.
- [18]. Ilin A, Raiko T. Practical approaches to principal component analysis in the presence of missing values. *J. Mach. Learn. Res.* 2010; 11:1957–2000.
- [19]. Jain, P.; Netrapalli, P.; Sanghavi, S. Low-rank matrix completion using alternating minimization. Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13, ACM; New York. 2013. p. 665-674.
- [20]. Jin Q, Sorzano COS, de la Rosa-Trevlin JM, Bilbao-Castro JR, Núñez-Ramírez R, Llorca O, Tama F, Joni S. Iterative elastic 3D-to-2D alignment method using normal modes for studying structural dynamics of large macromolecular complexes. *Structure*. 2014; 22:496–506. [PubMed: 24508340]
- [21]. Johnstone I. On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* 2001; 29:295–327.
- [22]. Johnstone I, Lu A. On consistency and sparsity for principal components analysis in high dimensions. *J. Amer. Statist. Assoc.* 2009; 104:682–693.
- [23]. Kalai AT, Moitra A, Valiant G. Disentangling Gaussians. *Commun. ACM*. 2012; 55:113–120.
- [24]. Kühlbrandt W. The resolution revolution. *Science*. 2014; 343:1443–1444. [PubMed: 24675944]
- [25]. Kuybeda O, Frank GA, Bartesaghi A, Borgnia M, Subramaniam S, Sapiro G. A collaborative framework for 3D alignment and classification of heterogeneous subvolumes in cryoelectron tomography. *J. Struct. Biol.* 2013; 181:116–127. [PubMed: 23110852]
- [26]. Kwon O, Zewail AH. 4D electron tomography. *Science*. 2010; 328:1668–1673. [PubMed: 20576886]
- [27]. Leger, F.; Yu, G.; Sapiro, G. Efficient matrix completion with Gaussian models. IEEE 2011 International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE; Piscataway, NJ. 2011. p. 1113-1116.
- [28]. Li X, Mooney P, Zheng S, Booth C, Braunfeld M, Gubbens S, Agard D, Cheng Y. Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-em. *Nature Methods*. 2013; 10:584–590. [PubMed: 23644547]
- [29]. Liao, H.; Frank, J. Classification by bootstrapping in single particle methods. Proceedings of the 2010 IEEE International Conference on Biomedical Imaging: From Nano to Macro, IEEE; Piscataway, NJ. 2010. p. 169-172.
- [30]. Liao M, Cao E, Julius D, Cheng Y. Structure of the TRPV 1 ion channel determined by electron cryo-microscopy. *Nature*. 2013; 504:107–124. [PubMed: 24305160]
- [31]. Little, R.; Rubin, D. Wiley Ser. Probab. Stat. 2nd. John Wiley; Hoboken, NJ: 2002. Statistical Analysis with Missing Data.
- [32]. Loh P, Wainwright M. High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *Ann. Statist.* 2012; 40:1637–1664.
- [33]. Lounici K. High-dimensional covariance matrix estimation with missing observations. *Bernoulli*. 2014; 20:1029–1058.
- [34]. Ludtke S, Baker M, Chen D, Song J, Chuang D, Chiu W. De novo backbone trace of GroEL from single particle electron cryomicroscopy. *Structure*. 2008; 16:441–448. [PubMed: 18334219]
- [35]. Mar enko VA, Pastur LA. Distribution of eigenvalues of some sets of random matrices. *Math. USSR Sb.* 1967; 1:507–536.

- [36]. Morrison MA, Parker GA. A guide to rotations in quantum mechanics. *Aust. J. Phys.* 1987; 40:465–497.
- [37]. Nadler B. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *Ann. Statist.* 2008; 36:2791–2817.
- [38]. Natterer, F. *Classics Appl. Math.* SIAM; Philadelphia: 2001. *The Mathematics of Computerized Tomography.*
- [39]. O’Neil M, Woolfe F, Rokhlin V. An algorithm for the rapid evaluation of special function transforms. *Appl. Comput. Harmon. Anal.* 2010; 28:203–226.
- [40]. Pearson K. On lines and planes of closest fit to systems of points in space. *Philos. Mag.* 1901; 2:559–572.
- [41]. Penczek, P.; Liang, ZP. Variance in three-dimensional reconstructions from projections. In: Unser, M., editor. *Proceedings of the 2002 IEEE International Symposium on Biomedical Imaging*; Piscataway, NJ. 2002. p. 749-752. IEEE
- [42]. Penczek P, Chao Y, Frank J, Spahn CMT. Estimation of variance in single-particle reconstruction using the bootstrap technique. *J. Struct. Biol.* 2006; 154:168–183. [PubMed: 16510296]
- [43]. Penczek P, Kimmel M, Spahn C. Identifying conformational states of macromolecules by eigenanalysis of resampled cryo-EM images. *Structure.* 2011; 19:1582–1590. [PubMed: 22078558]
- [44]. Penczek P, Renka R, Schomberg H. Gridding-based direct Fourier inversion of the three-dimensional ray transform. *J. Opt. Soc. Amer. A.* 2004; 21:499–509.
- [45]. Prudnikov, AP.; Brychkov, YA.; Marychev, OI. *Integrals and Series: Special Functions.* Gordon and Breach; Amsterdam: 1983.
- [46]. Recht B, Fazel M, Parrilo PA. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.* 2010; 52:471–501.
- [47]. Rudelson M. Random vectors in the isotropic position. *J. Funct. Anal.* 1999; 164:60–72.
- [48]. Saxton WO, Baumeister W. The correlation averaging of a regularly arranged bacterial cell envelope protein. *J. Microscopy.* 1982; 127:127–138.
- [49]. Scheres S. Relion: Implementation of a Bayesian approach to cryo-EM structure determination. *J. Struct. Biol.* 2012; 180:519–530. [PubMed: 23000701]
- [50]. Scheres S. Maximum-likelihood methods in cryo-EM. Part II: Application to experimental data. *J. Struct. Biol.* 2013; 181:195–206. [PubMed: 23246781]
- [51]. Schneider T. Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values. *J. Climate.* 2001; 14:853–871.
- [52]. Shatsky M, Hall R, Nogales E, Malik J, Brenner S. Automated multi-model reconstruction from single-particle electron microscopy data. *J. Struct. Biol.* 2010; 170:98–108. [PubMed: 20085819]
- [53]. Sigworth F, Doerschuk P, Carazo J, Scheres S. Maximum-likelihood methods in cryo EM. Part I: Theoretical basis and overview of existing approaches. *Methods Enzymology.* 2010; 482:263–294.
- [54]. Silverstein JW, Bai ZD. On the empirical distribution of eigenvalues of a class of large dimensional random matrices. *J. Multivariate Anal.* 1995; 54:175–192.
- [55]. Singer A, Shkolnisky Y. Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming. *SIAM J. Imag. Sci.* 2011; 4:543–572.
- [56]. Slepian D. Prolate spheroidal wave functions. Fourier analysis and uncertainty – IV: Extensions to many dimensions; generalized prolate spheroidal functions, *Bell System Tech. J.* 1964; 43:3009–3057.
- [57]. Stein, EM.; Weiss, GL. *Introduction to Fourier Analysis on Euclidean Spaces.* Princeton University Press; Princeton, NJ: 1971.
- [58]. Trefethen, L.; Bau, D, III. *Numerical Linear Algebra.* SIAM; Philadelphia: 1997.
- [59]. Tropp J. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.* 2012; 12:389–434.
- [60]. van Heel M. Principles of Phase Contrast (Electron) Microscopy. 2009 http://www.singleparticles.org/methodology/MvH_Phase_Contrast.pdf.

- [61]. van Heel M, Gowen B, Matadeen R, Orlova EV, Finn R, Pape T, Cohen D, Stark H, Schmidt R, Patwardhan A. Single particle electron cryo-microscopy: Towards atomic resolution. *Quart. Rev. Biophys.* 2000; 33:307–369.
- [62]. Vershynin, R. Introduction to the non-asymptotic analysis of random matrices, in *Compressed Sensing, Theory and Applications*. In: Idar, Y.; Kutyniok, G., editors. Cambridge University Press; Cambridge: 2012. p. 210-268.
- [63]. Wang L, Sigworth FJ. Cryo-EM and single particles. *Physiology (Bethesda)*. 2006; 21:13–18. [PubMed: 16443818]
- [64]. Wang L, Singer A, Wen Z. Orientation determination of cryo-EM images using least squares deviations. *SIAM J. Imag. Sci.* 2013; 6:2450–2483.
- [65]. Wang Q, Matsui T, Domitrovic T, Zheng Y, Doerschuk P, Johnson J. Dynamics in cryo EM reconstructions visualized with maximum-likelihood derived variance maps. *J. Struct. Biol.* 2013; 181:195–206. [PubMed: 23246781]
- [66]. Wilks SS. Moments and distributions of estimates of population parameters from fragmentary samples. *Ann. Math. Statist.* 1932; 3:163–195.
- [67]. Zhang W, Kimmel M, Spahn CM, Penczek P. Heterogeneity of large macromolecular complexes revealed by 3d cryo-em variance analysis. *Structure*. 2008; 16:1770–1776. [PubMed: 19081053]
- [68]. Zhang X, Settembre E, Xu C, Dormitzer P, Bellamy R, Harrison S, Grigorieff N. Near-atomic resolution using electron cryomicroscopy and single-particle reconstruction. *Proc. Natl. Acad. Sci. USA*. 2008; 105:1867–1872. [PubMed: 18238898]
- [69]. Zhao Z, Singer A. Fourier-Bessel rotational invariant eigenimages. *J. Opt. Soc. Amer. A*. 2013; 30:871–877.
- [70]. Zhao Z, Singer A. Rotationally invariant image representation for viewing direction classification in cryo-EM. *J. Struct. Biol.* 2014; 186:153–166. [PubMed: 24631969]






Figure 1. Classical (left) and hybrid (right) states of 70S E. Coli ribosome (image source: [29]).



(a) MP distributions (2.19) for $\sigma^2 = 1$



(b) Empirical spectrum for spiked model

Figure 2.
Illustrations of high-dimensional PCA

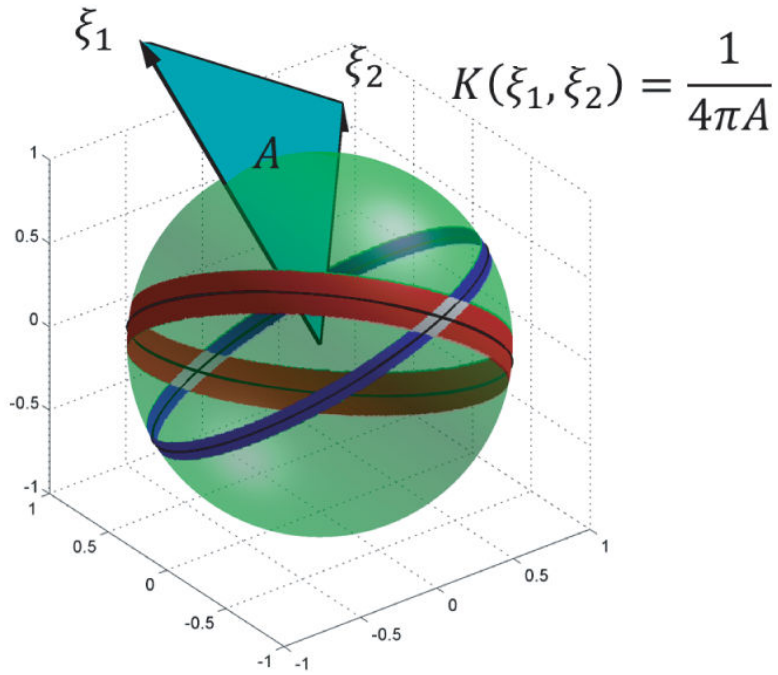


Figure 3.

The triangular area filter. ξ_1 induces a strip on S^2 of width proportional to $1/|\xi_1|$ (blue); ξ_2 induces a strip of width proportional to $1/|\xi_2|$ (red). The strips intersect in two parallelogram-shaped regions (white), each with area proportional to $1/|\xi_1 \times \xi_2|$. Hence, $K(\xi_1, \xi_2)$ is inversely proportional to the area of the triangle spanned by ξ_1, ξ_2 (cyan).

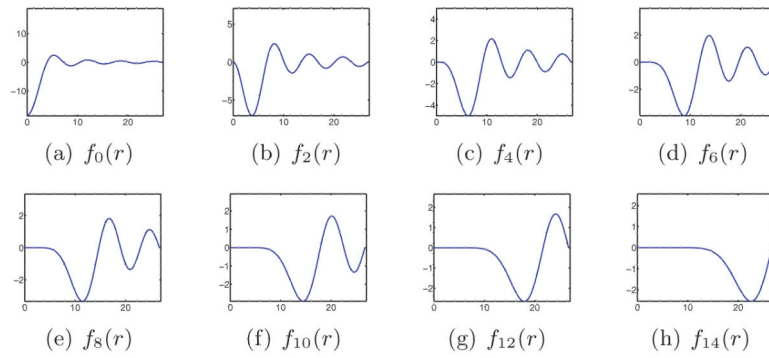


Figure 4.

The even basis functions up to $f_{14}(r)$. Note that they become less oscillatory as k increases, and that $f^k(r) \sim r^k$ at the origin. The odd basis functions have a similar structure and so are not pictured.




Figure 5. Block diagonal structure of PPs. The shaded rectangles represent the nonzero entries. For an explanation of the specific pairing of angular and radial functions, see (5.27) and (5.19) and the preceding discussion. A short calculation shows that the k th block of \hat{P}_S has size

$$(k+1) \times \frac{(k+1)(k+2)}{2}.$$

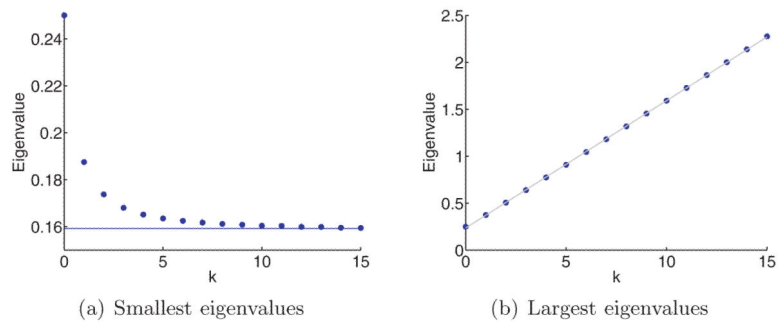


Figure 6.

The smallest and largest eigenvalues of (the continuous version of) LPk,k , for $0 \leq k \leq 15$.

The smallest eigenvalues approach their theoretical lower bound of $1/2\pi$ as k increases. The largest eigenvalues show a clear linear dependence on k .




Figure 7.

This figure depicts the effect of mean subtraction on projection images in the context of a two-class heterogeneity. The bottom row projections are obtained from the top row by mean subtraction. Columns (a) and (b) are clean projection images of the two classes from a fixed viewing angle. Columns (c) and (d) are both noisy versions of column (a). The image in the top row of column (c) has an SNR of 0.96, but the SNR of the corresponding mean-subtracted image is only 0.05. In column (d), the top image has an SNR of 0.19, but the mean-subtracted image has SNR 0.01. Note: the SNR values here are not normalized by N^2/N^2 in order to illustrate the signal present in a projection image.




Figure 8.

Cross-sections of reconstructions of the mean, top eigenvector, and two volumes for three different SNR values. The top row is clean, the second row corresponds to $SNR_{het} = 0.013$ (0.25), the third row to $SNR_{het} = 0.003$ (0.056), and the last row to $SNR_{het} = 0.0013$ (0.025). (a) $SNR_{het} = 0.013$ (0.25) (b) $SNR_{het} = 0.003$ (0.056) (c) $SNR_{het} = 0.0013$ (0.025)




Figure 9. *Eigenvalue histograms of $\Sigma P n x$ in the two-volume case for three SNR values. Note that as the SNR decreases, the distribution of eigenvalues associated with noise comes increasingly closer to the top eigenvalue that corresponds to the structural variability, and eventually the latter is no longer distinguishable.*


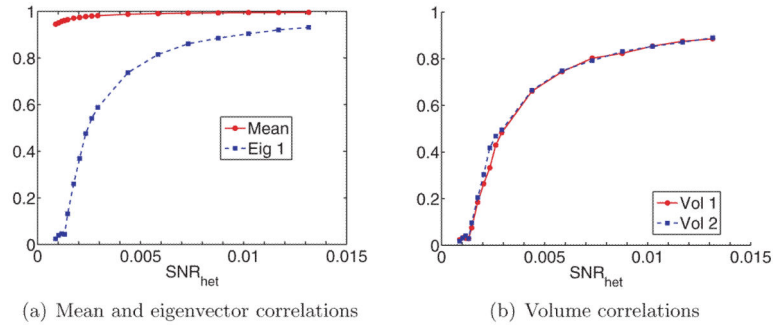


Figure 10.

FSC curves for the mean volume, top eigenvector, and one mean-subtracted volume at the same three SNR levels as in Figure 8. Note that the mean volume is reconstructed successfully for all three SNR levels. On the other hand, the top eigenvector and volume are recovered at the highest two SNR levels but not at the lowest SNR.

**Figure 11.**

Correlations of computed quantities with their true values for different SNRs (averaged over 10 experiments) for the two-volume case. Note that in the two-volume case, the mean-subtracted volume correlations are essentially the same as the eigenvector correlation (the only small discrepancy is that we subtract the true mean rather than the computed mean to obtain the former).




Figure 12.

Histograms of α_s for two-class case. Note that (a) has a bimodal distribution corresponding to two heterogeneity classes, but these two distributions merge as SNR decreases. (0.002) and 0.003 (0.006). Note that this behavior is tied to the spectral gap (separation of top eigenvalues from the bulk) of $\Sigma^P n$. Indeed, the disappearance of the spectral gap going from panel (b) to panel (c) of Figure 9 coincides with the estimated top eigenvector becoming uncorrelated with the truth, as reflected in Figures 10(b) and 11(a). This phase transition behavior is very similar to that observed in the usual high-dimensional PCA setup, described in section 2.3.

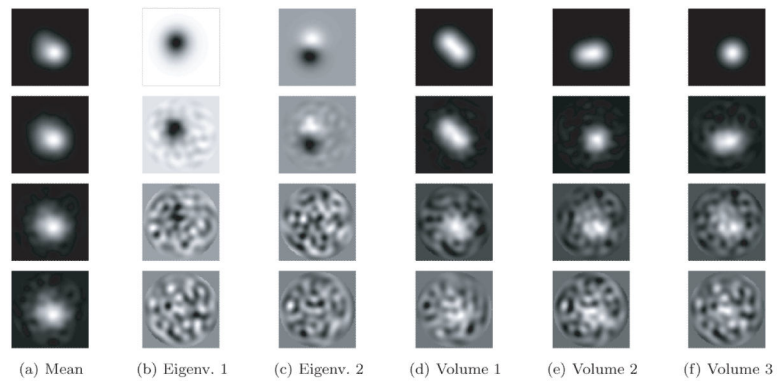


Figure 13. Cross sections of clean and reconstructed objects for the three-class experiment. The top row is clean, the second row corresponds to $SNR_{het} = 0.044$ (0.3), the third row to $SNR_{het} = 0.0044$ (0.03), and the last row to $SNR_{het} = 0.0015$ (0.01).





Figure 14.

Eigenvalue histograms of reconstructed covariance matrix in the three-class case for three SNR values. Note that the noise distribution initially engulfs the second eigenvalue, and eventually the top eigenvalue as well.

**Figure 15.**

FSC curves for the mean volume, top eigenvector, and one mean-subtracted volume at the same three SNRs as in Figure 13. Note that the mean volume is reconstructed successfully for all three SNR levels, and that the second eigenvector is recovered less accurately than the first.

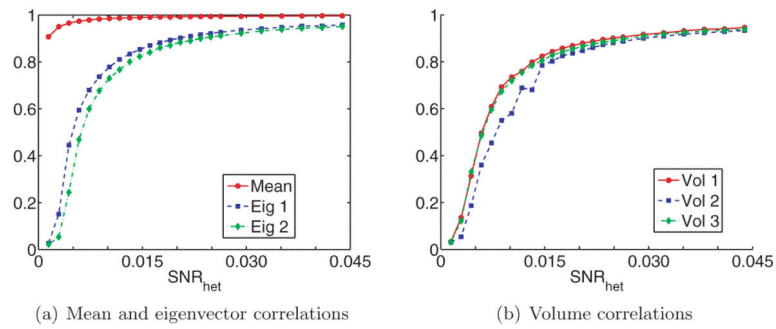


Figure 16.

Correlations of computed means, eigenvectors, and mean-subtracted volumes with their true values for different SNRs (averaged over 30 experiments). Note that the mean volume is consistently recovered well, whereas recovery of the eigenvectors and volumes exhibits a phase-transition behavior.




Figure 17.

The coordinates α_s for the three-class case, colored according to true class. The middle scatter plot is near the transition at which the three clusters coalesce.




Figure 18. Eigenvalue histograms of covariance matrix reconstructed in continuous variation case.




Figure 19. Scatter plots (with some outliers removed) of α_s for high SNR values.