

Toward single particle reconstruction without particle picking: Breaking the detection limit

Tamir Bendory^a, Nicolas Boumal^b, William Leeb^c, Eitan Levin^d, and Amit Singer^e

^aThe School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel

^bInstitute of Mathematics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

^cSchool of Mathematics, University of Minnesota, Minneapolis, Minnesota, USA

^dDepartment of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California, USA

^eThe Program in Applied and Computational Mathematics and Department of Mathematics, Princeton University, Princeton, NJ, USA

October 27, 2022

Abstract

Single-particle cryo-electron microscopy (cryo-EM) has recently joined X-ray crystallography and NMR spectroscopy as a high-resolution structural method to resolve biological macromolecules. In a cryo-EM experiment, the microscope produces images called micrographs. Projections of the molecule of interest are embedded in the micrographs at unknown locations, and under unknown viewing directions. Standard imaging techniques first locate these projections (detection) and then reconstruct the 3-D structure from them. Unfortunately, high noise levels hinder detection. When reliable detection is rendered impossible, the standard techniques fail. This is a problem, especially for small molecules. In this paper, we pursue a radically different approach: we contend that the structure could, in principle, be reconstructed directly from the micrographs, without intermediate detection. The aim is to bring small molecules within reach for cryo-EM. To this end, we design an autocorrelation analysis technique that allows to go directly from the micrographs to the sought structures. This involves only one pass over the micrographs, allowing online, streaming processing for large experiments. We show numerical results and discuss challenges that lay ahead to turn this proof-of-concept into a complementary approach to state-of-the-art algorithms.

1 Introduction

Cryo-electron microscopy (cryo-EM) is an imaging technique in structural biology used for single particle reconstruction of macromolecules. In a cryo-EM experiment, biological samples are rapidly frozen in a thin layer of vitreous ice. The microscope produces a 2-D tomographic image of the samples embedded in the ice, called a *micrograph*. Each micrograph contains tomographic projections of the samples at unknown locations and under unknown viewing directions. Figure 1 presents three experimental micrographs. The goal is to construct 3-D models of the molecular structure from the micrographs [19, 9, 56].

The signal-to-noise ratio (SNR) of the tomographic projections in the micrographs is a function of two dominating factors. On the one hand, the SNR is a function of the electron dose. To keep radiation




Figure 1: Excerpts from three data sets from the EMPIAR repository [29]: EMPIAR 10028 (Plasmodium Falciparum 80S ribosome) [64], EMPIAR 10061 (β -Galactosidase) [8], and EMPIAR 10249 (Alcohol Dehydrogenase) [28]. Detecting particles in the EMPIAR 10028 data set (left panel) is rather easy as the particles in the micrograph, corresponding to a large ribosome with a molecular weight of 4 MDa, are clearly visible. Doing so in the EMPIAR 10061 data set (middle panel) is harder, as the β -Galactosidase is a smaller protein with a molecular weight of 465 kDa, but some particles can still be identified visually. The particles in the EMPIAR 10249 data set (right panel) are tomographic projections of a small molecule that weighs 82 kDa, yielding low SNR. We note that the differences in particle contrast are also attributed to the fact that the EMPIAR 10249 data set was obtained with a microscope operating at 200 keV, whereas the EMPIAR 10028 and EMPIAR 10061 data sets were taken at 300 keV. In this work, we propose a method that aims to recover molecular structures when the SNR is very low, even below the SNR of the EMPIAR 10249 data set.

damage within acceptable bounds, the dose must be kept low, which leads to high noise levels. On the other hand, the SNR is a function of the molecule size. The smaller the molecules, the fewer detected electrons carry information about them.

All methods currently in use split the reconstruction procedure into two main stages. The first stage consists in extracting the various particle projections from the micrographs. This stage is called *particle picking* [50, 24, 14, 17]. The second stage aims to construct a 3-D model of the molecular structure from these projections. The quality of the reconstruction eventually hinges on the quality of the particle picking stage. Figure 1a shows an example of a data set of a large ribosome with a molecular weight of 4MDa, where the particles in the micrograph are clearly visible, and thus particle picking is rather easy. The particles in Figure 1b correspond to a smaller protein with a molecular weight of 465 kDa, but some particles can still be identified visually. The particles in Figure 1c are tomographic projections of a small molecule that weighs 82 kDa, yielding low SNR.

Crucially, it can be shown that reliable detection of individual particles is impossible below a certain critical SNR. This fact has been recognized early on by the cryo-EM community. In particular, in an influential paper from 1995, Nobel laureate Richard Henderson [26] investigates the following questions:

For the purposes of this review, I would like to ask the question: what is the smallest size of free-standing molecule whose structure can in principle be determined by phase-contrast electron microscopy? Given what has already been demonstrated in published work, this reduces to the question: what is the smallest size of molecule for which it is possible to determine from images of unstained molecules the five parameters needed to define accurately its orientation (three parameters) and position (two parameters) so that averaging can be performed?

In that paper and in others that followed (e.g., [22]), it was established that particle picking is impossible for molecules below a certain weight (below ~ 40 kDa). Joachim Frank voices a similar observation in

his 2017 Nobel Prize lecture: “*Using the ribosome as an example, it became clear from the formula we obtained that the single-particle approach to structure research was indeed feasible for molecules of sufficient size: Particle Size > 3/[Contrast² × Resolution (as length) × Critical Electron Dose]*” [20]. As these two leaders of the cryo-EM field point out, it is impossible to reconstruct small molecules by any of the existing computational pipelines for single particle analysis in cryo-EM, because the particles themselves cannot be picked from the micrographs. Given the fact that the vast majority of the proteins that make up the mammalian proteome weigh less than 100 kDa, whereas the molecular weights of almost all biological macromolecules whose structures have been determined using cryo-EM is greater than 100 kDa, new methodologies targeting macromolecules in this size are of pivotal importance [40]. In addition, recovering small structures is crucial for structure-guided drug design [49].

Recovering smaller molecular structures using cryo-EM is an active research effort [66, 65, 5], mostly focused on sample preparation techniques and hardware developments, such as Volta phase plates [31, 42], laser phase plates [51], and scaffolding cages [43], as well as steady improvements in the data processing pipeline. Despite this progress, detecting small molecules in the micrographs remains a challenge (see Figure 1c). We note that nuclear magnetic resonance (NMR) spectroscopy and X-ray crystallography are well suited to reconstruct small molecules. Yet, cryo-EM has a lot to offer even for molecules with already known structures obtained via NMR spectroscopy or X-ray crystallography, because these methods have limited ability to distinguish conformational variability [58].

In this paper, we argue that there is a gap between the two questions in Henderson’s quoted excerpt above, and that one may be able to exploit it to design better reconstruction algorithms. Specifically, the impossibility of particle picking does not necessarily imply impossibility of particle reconstruction. Indeed, the aim is only to reconstruct the molecule: estimating the locations of the particles in the micrograph is merely a helpful intermediate stage when it can be done. Our main message is that, while a molecule’s size may limit our ability to pick particles, that does not necessarily translate into a limit on our ability to reconstruct the molecule.

In order to recover the 3-D structure, we use autocorrelation analysis that relates the autocorrelations of the micrographs to the parameters describing the 3-D model. For any noise level, the autocorrelations of the micrographs can be estimated to any desired accuracy, provided we observe sufficiently many micrographs. The autocorrelations of the micrographs are straightforward to compute and require only one pass over the data. To estimate the 3-D structure itself from the estimated autocorrelations, we formulate and solve a nonlinear inverse problem. In particular, we show that the first order autocorrelation is just the average pixel value, the second-order autocorrelation is effectively a 1-D radial function, and the third-order is a 3-D function. We may therefore hope (but not guarantee) to recover the 3-D volume from the third-order autocorrelation, as it provides sufficiently many equations to solve for the parameters. Importantly, there is no need to detect individual projection images. We show a few numerical examples and outline the future developments required to make this method applicable to experimental data.

Another interesting feature of the described approach pertains to model bias, whose importance in cryo-EM was stressed by a number of authors [54, 61, 27, 60]. In the classical “Einstein from noise” experiment, multiple realizations of pure noise are aligned to a picture of Einstein using template matching and then averaged. In [54], it was shown that the averaged noise rapidly becomes remarkably similar to the Einstein template. In the context of cryo-EM, this experiment exemplifies how prior assumptions about the particles may influence the reconstructed structure. This model bias is common to all particle picking methods based on template matching. In our approach, we do not attempt to match a template to noisy data. Instead, we use an optimization algorithm to try to solve the inverse problem relating autocorrelations to the sought volume. This nonconvex problem may have spurious local minima, hence the algorithm’s initialization (a form of template) may affect the outcome. However, we find empirically that this poses a lesser risk because we are able to obtain meaningful reconstructions from random initializations.

The rest of the paper is organized as follows. Section 2 shows that it is impossible to detect the particle images in the micrograph in extremely low SNR regimes. In Section 3, which is the main contribution of this paper, we develop an autocorrelation analysis framework to recover molecular structures directly from the micrograph, circumventing particle picking. Section 4 shows a few numerical experiments, and Section 5 concludes the paper and delineates challenges that lay ahead to turn this proof-of-concept into a complementary approach to state-of-the-art algorithms.

2 The detection limit

In the low SNR regime—even if the 3-D structure is known—one cannot reliably detect the projection images in the micrograph. To support this claim, we consider a strictly simpler problem: suppose an oracle identifies for us one patch in the micrograph that either contains a full signal occurrence (plus noise), or contains just noise. Our task is to determine which one it is. The oracle further provides the true signal x , the noise variance σ^2 , and the probability q that we observe signal-plus-noise as opposed to pure noise.

This decision problem can be abstracted as follows: We have two known vectors $\theta_0 = x$ and $\theta_1 = 0$ in \mathbb{R}^L . There is a random variable η taking values 0 or 1 with probabilities q and $1 - q$, respectively. We observe a random vector $X \in \mathbb{R}^L$ (akin to an extract of the micrograph) with the following distribution: if $\eta = 0$, then $X \sim \mathcal{N}(\theta_0, \sigma^2 I_L)$ and if $\eta = 1$, then $X \sim \mathcal{N}(\theta_1, \sigma^2 I_L)$.

We observe a realization of X , and our task is to estimate η . How reliably can this be done? If $q \geq 1/2$, the constant estimator $\hat{\eta} = 0$ is correct with probability q ; likewise, if $q \leq 1/2$, the constant estimator $\hat{\eta} = 1$ is correct with probability $1 - q$. The question is, can we do better than this? We prove that, as $\sigma \rightarrow \infty$, the answer is no. The result is proved in Appendix A.

Proposition 2.1. *For any deterministic estimator $\hat{\eta}$ of η ,*

$$\lim_{\sigma \rightarrow \infty} \text{Prob}[\hat{\eta} = \eta] \leq \max(q, 1 - q); \quad (2.1)$$

that is, as the SNR deteriorates, the probability of success is no better than random chance.

Proposition 2.1 implies that, at low SNR, particle picking is impossible: in order to estimate the 3-D structure, we must consider methods that aim to estimate the structure directly, without estimating the nuisance locations of the projection images as an intermediate step. In the next section, we develop an autocorrelation analysis technique for that purpose.

3 Autocorrelation analysis for cryo-EM

3.1 Autocorrelation analysis

For a random signal $z \in \mathbb{R}^{N \times N}$, the autocorrelation of order p is given by

$$a_z^p[\ell_1, \dots, \ell_{p-1}] = \mathbb{E}_z \left\{ \frac{1}{N^2} \sum_{\mathbf{i}} z[\mathbf{i}] z[\mathbf{i} + \ell_1] \cdots z[\mathbf{i} + \ell_{p-1}] \right\}, \quad (3.1)$$

where the expectation is taken with respect to the distribution of z , the summation is for \mathbf{i} ranging over the N^2 pixels of the signal, and $\ell_1, \dots, \ell_{p-1}$ are two-dimensional integer shifts. Indexing out of bounds is zero-padded, that is, $z[\mathbf{i}] = 0$ for \mathbf{i} out of the range $[0, N - 1] \times [0, N - 1]$. In our case, z will represent the micrographs, whose distribution is parameterized by the sought 3-D structure. In the next sections, we show that as $N \rightarrow \infty$, the first three autocorrelations of the micrograph (i.e., $p = 1, 2, 3$) converge

to explicit polynomial functions of the parameters describing the sought volume. Then, we wish to estimate the volume by finding a set of parameters which are consistent with the autocorrelations; this entails solving a non-convex optimization problem.

We mention that autocorrelation analysis for cryo-EM images was first proposed in 1980 by Zvi Kam [30]. However, Kam’s method, as well as some of its recent extensions [41, 53], assumes picked particles. To break this fundamental barrier, we suggest computing the autocorrelations of the micrographs directly, completely bypassing particle picking.

3.2 Model and autocorrelation functions

Let $\phi: \mathbb{R}^3 \rightarrow \mathbb{R}$ be the Coulomb potential representing the molecule we aim to recover. We assume that the molecule is smooth. Specifically, in spherical coordinates, its 3-D Fourier transform $\widehat{\phi}$ admits a finite expansion of the form

$$\widehat{\phi}(ck, \theta, \varphi) = \sum_{\ell=0}^{L_{\max}} \sum_{m=-\ell}^{\ell} \sum_{s=1}^{S(\ell)} x_{\ell,m,s} Y_{\ell}^m(\theta, \varphi) j_{\ell,s}(k), \quad k \leq 1, \quad (3.2)$$

where $0 < c \leq 0.5$ is the bandlimit in Fourier space (a standard assumption in cryo-EM) $\{S(\ell)\}$ are determined using the Nyquist criterion as described in [15, 33], $j_{\ell,s}$ is the normalized spherical Bessel functions given by

$$j_{\ell,s}(k) = \frac{4}{|j_{\ell+1}(u_{\ell,s})|} j_{\ell}(u_{\ell,s}k),$$

j_{ℓ} is the spherical Bessel function of order ℓ and $u_{\ell,s}$ is the s th positive zero of j_{ℓ} . We use the complex spherical harmonics Y_{ℓ}^m defined by

$$Y_{\ell}^m(\theta, \varphi) := \sqrt{\frac{2\ell+1}{4\pi} \cdot \frac{(\ell-m)!}{(\ell+m)!}} P_{\ell}^m(\cos \theta) e^{im\varphi},$$

where P_{ℓ}^m are the associated Legendre polynomials with the Condon-Shortley phase. Representing the molecule using a spherical harmonics expansion, akin to (3.2), is a ubiquitous practice in the cryo-EM literature, see for example [7]. Sampling at the Nyquist rate dictates $c = 1/2$ [41]. Because ϕ is real-valued, $\widehat{\phi}$ is conjugate-symmetric and thus the expansion coefficients satisfy $x_{\ell,-m,s} = (-1)^{\ell+m} \overline{x_{\ell,m,s}}$. Therefore, we only need to recover coefficients $x_{\ell,m,s}$ with $m \geq 0$ in order to recover the molecular structure.

Let I_{ω} denote the tomographic projection obtained from viewing direction $\omega \in SO(3)$. The Fourier slice theorem states that the 2-D Fourier transform of a tomographic projection is equal to a 2-D slice of the volume’s 3-D Fourier transform [45]. Specifically,

$$\widehat{I}_{\omega}(ck, \varphi) = \sum_{\ell,m,m',s} x_{\ell,m,s} D_{m',m}^{\ell}(\omega) Y_{\ell}^{m'}\left(\frac{\pi}{2}, \varphi\right) j_{\ell,s}(k), \quad (3.3)$$

where $D_{m',m}^{\ell}(\omega)$ is a Wigner-D matrix. This implies that the projections are also c -bandlimited. In practice, the projections are further affected by additional factors, such as the microscope’s point spread function, which are neglected in this paper.

Let $\mathcal{I} \in \mathbb{R}^{N \times N}$ denote a micrograph. We assume it consists of shifted copies of projections contaminated by additive white Gaussian noise:

$$\mathcal{I} = \sum_{t=1}^M I_{\omega_t} * \delta_{\mathbf{s}_t} + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I), \quad (3.4)$$

where $*$ denotes convolution, and the viewing directions ω_t are assumed to be drawn from the uniform distribution over $SO(3)$. We assume the projections are discretized on a Cartesian grid of size $P \times P$ and $\mathbf{s}_t \in [P, N - P]^2$ denotes the location of the upper-left corner of the t -th projection in the micrograph.

We impose a *separation condition* so that any two projections are separated by at least $2P - 1$ pixels between their upper-left corners, in each direction. Thus, their end points (in each direction) are necessarily separated by at least $P - 1$ signal-free entries in the micrograph.

We stress that the assumptions of the generative model (3.4) are unrealistic for experimental cryo-EM data. For example, in practice, the particles are not likely to satisfy the separation condition (but they do not overlap and thus are separated by at least $P - 1$ pixels in each direction), and the distribution over $SO(3)$ is typically non-uniform; these assumptions were made to ease the analysis. Section 5 discusses how our autocorrelation analysis technique can be extended to include these important aspects that are essential to reach high-resolution reconstructions.

Define the p th (empirical) autocorrelation of \mathcal{I} as

$$a_{\mathcal{I}}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] := \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}_{p-1}], \quad (3.5)$$

where the summation is for \mathbf{i} ranging over the N^2 pixels of the micrograph, using zero padding when indices exceed the micrograph's edge. Computing the autocorrelations of the micrograph is straightforward. Let $\mathcal{I}_1, \dots, \mathcal{I}_K$ denote a set of K micrographs. Under the specified conditions, we show in the next section that the first three autocorrelations of the micrographs are related to those of the projections by

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{i=1}^K a_{\mathcal{I}_i}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] = \gamma \langle a_{\mathcal{I}_\omega}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] \rangle_\omega + b_p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}], \quad (3.6)$$

$$p = 1, 2, 3, \quad \boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1} \in [-(P - 1), P - 1]^2,$$

where $\langle \cdot \rangle_\omega$ denotes averaging over all possible viewing directions ω with respect to the uniform measure, $\gamma = \frac{MP^2}{N^2} \in (0, 1)$ is a scalar that encodes the density of the particle projections in the data, and b_p is a bias term. Specifically, $b_1 = 0$ and therefore the mean is unbiased. The bias term of the second-order autocorrelation b_2 depends only on σ^2 , the variance of the noise. Hence, if the noise level can be accurately estimated from the micrographs, this bias can be removed. Finally, the bias term of the third-order autocorrelation b_3 depends on the mean of the micrograph and σ^2 . Therefore, given sufficiently many projections, we can accurately estimate the quantities $\gamma \langle a_{\mathcal{I}_\omega}^p \rangle_\omega$ directly from the micrographs. These quantities are functions of the unknown coefficients $x_{l,m,s}$ and we could proceed to invert their relation.

In practice, we want to leverage one more feature of the 3-D reconstruction problem. Since all in-plane rotations of the micrographs are equally likely observations, it is desirable in (3.6) to average over all in-plane rotations as well. This can be done efficiently using Prolate Spheroidal Wave Functions (PSWFs), which have been used before for cryo-EM data processing [38, 39, 53, 25]. In particular, as we show next, averaging over all in-plane rotations using the PSWFs reduces the dimensionality of the problem, without any loss of information. We use autocorrelations up to and including the third order. Indeed, second-order autocorrelations are not enough, as was observed already in [30] for a simpler problem where the input is not micrographs but rather picked, perfectly centered particles.

3.3 Autocorrelation derivation

In this section, we prove relation (3.6). We note that mathematically taking infinitely many micrographs is equivalent to taking one infinitely large micrograph with fixed density γ . Hence, we consider the

moments of one micrograph \mathcal{I} in the limit $N \rightarrow \infty$ and $\gamma = \lim_{N \rightarrow \infty} \frac{MP^2}{N^2} \in (0, 1)$. The separation condition guarantees that if $\mathbf{i} = (i, j)$ is in the support of some projection, then $\mathbf{i} + \boldsymbol{\ell}$ for $\boldsymbol{\ell} \in [-(P-1), P-1]^2$ is either in the support of the same projection or outside the support of any projection.

We begin by calculating the relation between the p th autocorrelation of the clean micrograph and the averaged autocorrelation of the projections. Let us denote the clean micrograph by $\tilde{\mathcal{I}} = \mathcal{I} - \varepsilon$, where \mathcal{I} and ε are given in (3.4). Denote by \mathcal{S}_t the support of the projection of the t -th particle in the micrograph. Then, we have

$$\begin{aligned}
a_{\tilde{\mathcal{I}}}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] &= \frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\
&= \frac{1}{N^2} \sum_{t=1}^M \sum_{\mathbf{i} \in \mathcal{S}_t} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\
&= \frac{MP^2}{N^2} \cdot \frac{1}{M} \sum_{t=1}^M \frac{1}{P^2} \sum_{i,j=0}^{P-1} I_{\omega_t}[\mathbf{i}] I_{\omega_t}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots I_{\omega_t}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\
&= \frac{MP^2}{N^2} \frac{1}{M} \sum_{t=1}^M a_{I_{\omega_t}}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] \\
&\rightarrow \gamma \langle a_{I_\omega}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] \rangle_\omega,
\end{aligned} \tag{3.7}$$

where the average is taken over ω with respect to the distribution of viewing directions. Here, we assume it to be uniform.

In the presence of noise, we get additional bias terms denoted by b_p in (3.6). The mean ($p = 1$) is unbiased, since the noise is assumed to have zero mean. For the second-order autocorrelation ($p = 2$), we have

$$\begin{aligned}
a_{\mathcal{I}}^2[\boldsymbol{\ell}] &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}] \\
&= \frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}] + \frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}] + \frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}] + \frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}].
\end{aligned}$$

The first term is given by (3.7) for $p = 2$. The cross terms have expected value 0, and therefore almost surely vanish in the limit due to the law of large numbers. The fourth term is zero unless $\boldsymbol{\ell} = 0$, in which case it converges to σ^2 . Thus, we conclude

$$a_{\mathcal{I}}^2[\boldsymbol{\ell}] \rightarrow \gamma \langle a_{I_\omega}^2[\boldsymbol{\ell}] \rangle_\omega + \sigma^2 \delta[\boldsymbol{\ell}], \tag{3.8}$$

where the bias term $b_2[\boldsymbol{\ell}] = \sigma^2 \delta[\boldsymbol{\ell}]$ depends only on the variance of the noise σ^2 .

For the third moments, we get 8 terms:

$$\begin{aligned}
a_{\mathcal{I}}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] &= \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_2]}_{(1)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_1] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_2]}_{(2)} \\
&+ \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_1] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_2]}_{(3)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_2]}_{(4)} \\
&+ \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_2]}_{(5)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_1] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_2]}_{(6)} \\
&+ \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_1] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_2]}_{(7)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \varepsilon[\mathbf{i} + \boldsymbol{\ell}_2]}_{(8)}.
\end{aligned}$$

We address these terms one by one:

- Term (1) is treated by (3.7) for $p = 3$;
- Term (2) is the third-order autocorrelation of pure noise which almost surely vanishes in the limit by the law of large numbers;
- Terms (3)-(5) depend linearly on the noise and hence vanish in the limit;
- For term (6), if $\boldsymbol{\ell}_1 \neq \boldsymbol{\ell}_2$ the term vanishes in the limit. If $\boldsymbol{\ell}_1 = \boldsymbol{\ell}_2$ then

$$\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}]^2 = \frac{MP^2}{N^2} \cdot \frac{1}{MP^2} \sum_{t=1}^M \sum_{\mathbf{i} \in \mathcal{S}_t} I_{\omega_t}[\mathbf{i}] \varepsilon[\mathbf{i} + \boldsymbol{\ell}]^2 \rightarrow \gamma \sigma^2 \langle a_{I_\omega}^1 \rangle_\omega,$$

where $\langle a_{I_\omega}^1 \rangle_\omega$ is the mean of the volume.

- Terms (7) and (8) contribute δ functions similar to (6).

Thus, we conclude that

$$a_{\mathcal{I}}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \rightarrow \gamma \langle a_{I_\omega}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \rangle_\omega + \gamma \sigma^2 \langle a_{I_\omega}^1 \rangle_\omega \left(\delta[\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2] + \delta[\boldsymbol{\ell}_1] + \delta[\boldsymbol{\ell}_2] \right), \quad (3.9)$$

where the second term is the bias $b_3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2]$. Note that $\gamma \langle a_{I_\omega}^1 \rangle_\omega$ is approximately the mean of the micrograph since $a_{\mathcal{I}}^1 \approx a_{\mathcal{I}}^2 \approx \gamma \langle a_{I_\omega}^1 \rangle_\omega$. Therefore, we do not need prior knowledge of γ to effectively debias the third-order autocorrelation.

3.4 Accounting for all in-plane rotations

We represent our autocorrelations using Prolate Spheroidal Wave Functions (PSWFs) [57]. As we demonstrate below, this makes it easier to account for the fact that all in-plane rotations of the micrographs are equally likely observations. The PSWFs are given in polar coordinates by¹

$$\psi_{k,q}(r, \varphi) = \begin{cases} \frac{1}{\sqrt{8\pi}} \alpha_{k,q} R_{k,q}(r) e^{i k \varphi}, & r \leq 1, \\ 0, & r > 1, \end{cases} \quad (3.10)$$

¹A different normalization is used in [38].

where k, q are integers with $q \geq 0$, the range of k, q is determined by Eq. (8) in [38], the $R_{k,q}$ are a family of real, one-dimensional functions satisfying $R_{k,q} = R_{|k|,q}$, and the $\alpha_{k,q}$ are scaling factors which will be defined in the next section. The PSWFs are orthogonal on the unit disk. Note that $\psi_{-k,q} = \overline{\psi_{k,q}}$, hence, we only need to consider PSWFs with $k \geq 0$ when expanding real images. As can be seen from the definition of the PSWFs (3.10), the effects of rotations and reflections on expansion coefficients of real images are, respectively, phase modulation and conjugation: this is why the PSWF basis is called *steerable* [38, 67]. Using the steerability property, we will show that the second-order autocorrelation, though two-dimensional, effectively only provides radial information. Similarly, the 4-D third-order autocorrelation truly only carries information along three dimensions. While there exist alternative steerable bases, such as Fourier-Bessel [67], we chose to work with the PSWFs since they are eigenfunctions of the truncated Fourier transform: a fact that we exploit later on.

We start with the second-order autocorrelation. For $\boldsymbol{\ell} \in [-(P-1), P-1]^2$, let us define

$$\mathbf{a}^2[k, q] = \sum_{\boldsymbol{\ell}} a_{\mathcal{I}}^2[\boldsymbol{\ell}] \overline{\psi_{k,q}[\boldsymbol{\ell}]}, \quad (3.11)$$

where $\psi_{k,q}[\boldsymbol{\ell}] := \psi_{k,q}(\boldsymbol{\ell}/(P-1))$ is a discretization of the PSWFs. Knowledge of these coefficients is essentially equivalent to knowledge of the second-order correlations owing to the following approximate identity:

$$a_{\mathcal{I}}^2[\boldsymbol{\ell}] \approx \sum_{k,q} \mathbf{a}^2[k, q] \psi_{k,q}[\boldsymbol{\ell}]. \quad (3.12)$$

This holds because the continuous PSWFs form an orthonormal basis, and their discretized counterparts with indices (k, q) appropriately bounded as in [38] are (empirically) almost orthonormal. As a result, for our purposes, the pair of equations above provides a basis expansion for the autocorrelations.

We now proceed to show that the coefficients $\mathbf{a}^2[k, q]$ can be computed from the micrographs directly. By definition,

$$\begin{aligned} \mathbf{a}^2[k, q] &= \sum_{\boldsymbol{\ell}} a_{\mathcal{I}}^2[\boldsymbol{\ell}] \overline{\psi_{k,q}[\boldsymbol{\ell}]} \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \left(\sum_{\boldsymbol{\ell}} \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}] \overline{\psi_{k,q}[\boldsymbol{\ell}]} \right) \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] a_{k,q}[\mathbf{i}], \end{aligned} \quad (3.13)$$

where we defined,

$$a_{k,q}[\mathbf{i}] = \sum_{\boldsymbol{\ell}} \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}] \overline{\psi_{k,q}[\boldsymbol{\ell}]}. \quad (3.14)$$

These coefficients can be computed efficiently. Indeed, consider a patch of the micrograph \mathcal{I} centered around pixel \mathbf{i} and of size $(2P-1) \times (2P-1)$. This is exactly the patch indexed in the sum above. Hence, using the same approximation as we did in (3.12), a direct expansion of that patch in the discretized PSWFs yields the sought coefficients:

$$\mathcal{I}[\mathbf{i} + \boldsymbol{\ell}] \approx \sum_{k,q} a_{k,q}[\mathbf{i}] \psi_{k,q}[\boldsymbol{\ell}]. \quad (3.15)$$

Thus, we proceed as follows: for each position \mathbf{i} in the micrograph \mathcal{I} , we extract the corresponding patch of size $(2P-1) \times (2P-1)$, expand it in the discretized PSWFs as in (3.15), and collect the $a_{k,q}$ as per (3.13) to constitute the second-order autocorrelation of the micrograph.

Following this formalism, and using the steerability property, it is now straightforward to account for all in-plane rotations and reflections of the micrograph. We have the following approximate expansions for a patch rotated about its center \mathbf{i} by an angle α :

$$\mathcal{I}^{\alpha,+}[\mathbf{i} + \boldsymbol{\ell}] \approx \sum_{k,q} a_{k,q} e^{-i k \alpha} \psi_{k,q}[\boldsymbol{\ell}],$$

and the reflection followed by a rotation by angle α :

$$\mathcal{I}^{\alpha,-}[\mathbf{i} + \boldsymbol{\ell}] \approx \sum_{k,q} \overline{a_{k,q}} e^{-i k \alpha} \psi_{k,q}[\boldsymbol{\ell}].$$

Averaging over all rotations of the patch $\mathcal{I}(\mathbf{i} + \Delta \mathbf{i})$ and its reflection we get

$$\begin{aligned} \mathbf{a}^2[k, q] &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \left(\frac{1}{4\pi} \int_0^{2\pi} (a_{k,q}[\mathbf{i}] + \overline{a_{k,q}[\mathbf{i}]}) e^{-i k \alpha} d\alpha \right) \\ &= \delta[k] \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] a_{0,q}[\mathbf{i}], \end{aligned} \quad (3.16)$$

where in the last equality we used that $a_{0,q}[\mathbf{i}]$ is real since both \mathcal{I} and $\psi_{0,q}$ are real valued (more generally, $a_{-k,q} = \overline{a_{k,q}}$). Thus, the second-order autocorrelation, though two-dimensional, effectively only provides radial information.

We now follow a similar approach to estimate the bias term b_2 . Introduce the coefficients \mathbf{b}_2 as:

$$\begin{aligned} \mathbf{b}_2[k, q] &= \sigma^2 \sum_{\boldsymbol{\ell}} \delta[\boldsymbol{\ell}] \overline{\psi_{k,q}[\boldsymbol{\ell}]} \\ &= \sigma^2 \overline{\psi_{k,q}[\mathbf{0}]} \\ &= \delta[k] \frac{\sigma^2}{\sqrt{2\pi}} R_{0,q}(0), \end{aligned}$$

where we used the fact that the functions $R_{k,q}$ are zero at the origin for $k \neq 0$. With this definition, we have the usual approximation:

$$b_2[\boldsymbol{\ell}] = \sigma^2 \delta[\boldsymbol{\ell}] \approx \sum_{k,q} \mathbf{b}_2[k, q] \psi_{k,q}[\boldsymbol{\ell}].$$

We now turn our attention to the third-order autocorrelation. Following the same lines, we define the coefficients:

$$\begin{aligned} \mathbf{a}^3[k_1, q_1; k_2, q_2] &= \sum_{\boldsymbol{\ell}_1, \boldsymbol{\ell}_2} a_{\mathcal{I}}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \overline{\psi_{k_1, q_1}[\boldsymbol{\ell}_1]} \psi_{k_2, q_2}[\boldsymbol{\ell}_2] \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \left(\sum_{\boldsymbol{\ell}_1} \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}_1] \overline{\psi_{k_1, q_1}[\boldsymbol{\ell}_1]} \right) \left(\sum_{\boldsymbol{\ell}_2} \mathcal{I}[\mathbf{i} + \boldsymbol{\ell}_2] \psi_{k_2, q_2}[\boldsymbol{\ell}_2] \right) \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] a_{k_1, q_1}[\mathbf{i}] \overline{a_{k_2, q_2}[\mathbf{i}]}, \end{aligned}$$

where the patch expansion coefficients $a_{k,q}$ are as defined in (3.14). The coefficients $\mathbf{a}^3[k_1, q_1; k_2, q_2]$ are related to the third-order autocorrelation via the approximate identity:

$$a_{\mathcal{I}}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \approx \sum_{\substack{k_1, q_1 \\ k_2, q_2}} \mathbf{a}^3[k_1, q_1; k_2, q_2] \psi_{k_1, q_1}[\boldsymbol{\ell}_1] \overline{\psi_{k_2, q_2}[\boldsymbol{\ell}_2]}.$$

Averaging over all rotations of \mathcal{I} and its reflection, we obtain

$$\begin{aligned} \mathbf{a}^3[k_1, q_1; k_2, q_2] &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \frac{1}{4\pi} \int_0^{2\pi} \left(a_{k_1, q_1}[\mathbf{i}] \overline{a_{k_2, q_2}[\mathbf{i}]} + \overline{a_{k_1, q_1}[\mathbf{i}]} a_{k_2, q_2}[\mathbf{i}] \right) e^{-\iota(k_1 - k_2)\alpha} d\alpha \\ &= \delta[k_1 - k_2] \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \Re\{a_{k_1, q_1}[\mathbf{i}] \overline{a_{k_2, q_2}[\mathbf{i}]} \}. \end{aligned} \quad (3.17)$$

Thus, similarly to the second-order autocorrelation, averaging over all in-plane rotations reveals that the 4-D third-order autocorrelation truly only carries information along three dimensions.

Finally, we treat the bias terms:

$$\mathbf{b}_3[k_1, q_1; k_2, q_2] = \gamma\sigma^2 \langle a_{I_\omega}^1 \rangle_\omega \delta[k_1 - k_2] \left[\delta[q_1 - q_2] + \delta[k_1] \frac{1}{2\pi} (\alpha_{0, q_1} + \alpha_{0, q_2}) R_{0, q_1}(0) R_{0, q_2}(0) \right].$$

Thus,

$$\begin{aligned} b_3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] &= \gamma\sigma^2 \langle a_{I_\omega}^1 \rangle_\omega \left(\delta[\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2] + \delta[\boldsymbol{\ell}_1] + \delta[\boldsymbol{\ell}_2] \right) \\ &\approx \sum_{\substack{k_1, q_1 \\ k_2, q_2}} \mathbf{b}_3[k_1, q_1; k_2, q_2] \psi_{k_1, q_1}[\boldsymbol{\ell}_1] \overline{\psi_{k_2, q_2}[\boldsymbol{\ell}_2]}. \end{aligned}$$

3.5 Connection to volume

Until now, we have established simple relations between the autocorrelations of the micrographs and the autocorrelations of the projection images of the 3-D structure, and explained how to compute the former. Now, we complete the picture by relating the latter to the 3-D structure itself.

Using the 2-D PSWFs and the Fourier slice theorem, we can express each projection (3.3) as

$$\widehat{I}_\omega(ck, \theta) = \sum_{N, n} b_{N, n}(\omega) \psi_{N, n}(k, \theta),$$

where

$$\begin{aligned} b_{N, n}(\omega) &= \frac{4}{\sqrt{2\pi} |\alpha_{N, n}|^2} \int_0^{2\pi} \int_0^1 \widehat{I}_\omega(ck, \theta) R_{N, n}(k) e^{-\iota N \theta} k dk d\theta, \\ &= \sum_{\ell, m, m', s} x_{\ell, m, s} \left[\frac{\sqrt{8\pi}}{\alpha_{N, n}} Y_\ell^{m'}(\pi/2, 0) \right] D_{m', m}^\ell(\omega) \left(\int_0^1 j_{\ell, s}(k) R_{N, n}(k) k dk \right) \left(\frac{1}{2\pi} \int_0^{2\pi} e^{\iota(m' - N)\theta} d\theta \right) \\ &= \sum_{|N| \leq \ell} \sum_{m, s} x_{\ell, m, s} D_{N, m}^\ell(\omega) \beta_{\ell, s; N, n}, \end{aligned}$$

and the coefficients

$$\beta_{\ell, s; N, n} := \begin{cases} \frac{\sqrt{8\pi}}{\alpha_{N, n}} Y_\ell^N(\pi/2, 0) \int_0^1 j_{\ell, s}(k) R_{N, n}(k) k dk, & \ell \geq |N| \\ 0, & \ell < |N| \end{cases}, \quad (3.18)$$

can be precomputed. The PSWFs are eigenfunctions of the truncated Fourier transform [38] and hence satisfy

$$\alpha_{N, n} \psi_{N, n}(\mathbf{k}) = \int_{\|\mathbf{r}\|_2 \leq 1} \psi_{N, n}(\mathbf{r}) e^{\iota c(\mathbf{r} \cdot \mathbf{k})} d\mathbf{r}. \quad (3.19)$$

We can now express the projection in real space as

$$\begin{aligned}
I_\omega(r, \varphi) &= \sum_{N,n} \widehat{\alpha}_{N,n} b_{N,n}(\omega) \psi_{N,n}(r, \varphi) \\
&= \sum_{\ell=0}^{L_{\max}} \sum_{N,m=-\ell}^{\ell} \sum_{n=0}^{n_{\max}(N)} \sum_{s=1}^{S(\ell)} x_{\ell,m,s} \widehat{\beta}_{\ell,s;N,n} D_{N,m}^\ell(\omega) \psi_{N,n}(r, \varphi),
\end{aligned} \tag{3.20}$$

where $n_{\max}(N)$ is chosen according to Eq. (8) in [38], $\alpha_{N,n}$ is the eigenvalue corresponding to the (N, n) th PSWF, $\widehat{\alpha}_{N,n} = (c/2\pi)^2 \alpha_{N,n}$, and $\widehat{\beta}_{\ell,s;N,n} = \widehat{\alpha}_{N,n} \beta_{\ell,s;N,n}$.

3.5.1 First-order autocorrelation (the mean)

Since $j_{\ell,s}(0) = 0$ unless $\ell = 0$, and since $Y_{0,0}(\theta, \varphi) = \frac{1}{\sqrt{4\pi}}$, we conclude from (3.2) that

$$\mathbf{a}_x^1 = \langle a_{I_\omega}^1 \rangle_\omega = \widehat{\phi}(\mathbf{0}) = \frac{1}{\sqrt{4\pi}} \sum_s x_{0,0,s} j_{0,s}(0). \tag{3.21}$$

3.5.2 Second-order autocorrelation

The second-order autocorrelation is easier to derive directly in Fourier space, to avoid integration of shifted PSWFs against centered ones. The relation between the second-order autocorrelation of the micrographs and projection images of the volume is given in (3.8) and (3.16). The connection with the expansion coefficients of the volume can be derived in Fourier space directly from Kam's original formula [30, Eq. (11)] by setting $\mathbf{k}_1 = \mathbf{k}_2$ to obtain,

$$\begin{aligned}
\langle a_{I_\omega}^2(k, \theta) \rangle_\omega &= \frac{1}{4\pi} \sum_{\ell,m} \left| \sum_s x_{\ell,m,s} j_{\ell,s}(k) \right|^2 \\
&= \frac{1}{4\pi} \sum_{\substack{\ell,m \\ s_1, s_2}} x_{\ell,m,s_1} \overline{x_{\ell,m,s_2}} j_{\ell,s_1}(k) j_{\ell,s_2}(k),
\end{aligned}$$

where we used the fact that the normalized spherical Bessel functions $j_{\ell,s}$ are real.

As before, we want to derive the relation with respect to the PSWF coefficients of the autocorrelation. Hence, we expand the above in 2-D PSWFs by

$$\langle a_{I_\omega}^2(k, \theta) \rangle_\omega = \sum_q \mathbf{a}_x^2(q) \psi_{0,q}(k),$$

and conclude that

$$\mathbf{a}_x^2[q] = \frac{1}{\sqrt{8\pi}} \sum_{\substack{\ell,m \\ s_1, s_2}} x_{\ell,m,s_1} \overline{x_{\ell,m,s_2}} \int_0^1 j_{\ell,s_1}(k) j_{\ell,s_2}(k) R_{0,q}(k) k dk. \tag{3.22}$$

The integral on k is precomputed.

3.5.3 Third-order autocorrelation

In (3.9) and (3.17) we have shown how the third-order autocorrelations of the micrographs and the projection images of the volume are related, and how we can represent them in PSWFs. Now, we relate these expressions to the expansion coefficients of the volume itself.

The third-order autocorrelation of the volume can be expressed in terms of (3.20):

$$\begin{aligned} \langle a_{I_\omega}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \rangle_\omega &= \frac{1}{L^2} \sum_{\mathbf{i}} \langle I_\omega[\mathbf{i}] I_\omega[\mathbf{i} + \boldsymbol{\ell}_1] \overline{I_\omega([\mathbf{i} + \boldsymbol{\ell}_2])} \rangle_\omega, \\ &\approx \frac{1}{L^2} \sum_{\substack{N_1, n_1 \\ N_2, n_2 \\ N_3, n_3}} \langle b_{N_1, n_1}(\omega) b_{N_2, n_2}(\omega) \overline{b_{N_3, n_3}(\omega)} \rangle_\omega \sum_{\mathbf{i}} \psi_{N_1, n_1}[\mathbf{i}] \psi_{N_2, n_2}[\mathbf{i} + \boldsymbol{\ell}_1] \overline{\psi_{N_3, n_3}([\mathbf{i} + \boldsymbol{\ell}_2])}, \end{aligned} \quad (3.23)$$

where the approximation is due to discretization. Now,

$$\begin{aligned} \langle b_{N_1, n_1}(\omega) b_{N_2, n_2}(\omega) \overline{b_{N_3, n_3}(\omega)} \rangle_\omega &= \sum_{\substack{\ell_1, m_1, s_1 \\ \ell_2, m_2, s_2 \\ \ell_3, m_3, s_3}} x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} \overline{x_{\ell_3, m_3, s_3}} \\ &\quad \times \langle D_{N_1, m_1}^{\ell_1}(\omega) D_{N_2, m_2}^{\ell_2}(\omega) \overline{D_{N_3, m_3}^{\ell_3}(\omega)} \rangle_\omega \widehat{\beta}_{\ell_1, s_1; N_1, n_1} \widehat{\beta}_{\ell_2, s_2; N_2, n_2} \overline{\widehat{\beta}_{\ell_3, s_3; N_3, n_3}}, \end{aligned}$$

where the latter coefficients are given explicitly in (3.18). Using standard properties of Wigner-D functions, we obtain

$$\left\langle D_{N_1, m_1}^{\ell_1}(\omega) D_{N_2, m_2}^{\ell_2}(\omega) \overline{D_{N_3, m_3}^{\ell_3}(\omega)} \right\rangle_\omega = (-1)^{N_3 + m_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ N_1 & N_2 & -N_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & -m_3 \end{pmatrix},$$

where $\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}$ are called Wigner 3-j symbols. Notably, these terms are zero unless $m_1 + m_2 + m_3 = 0$ and $|\ell_1 - \ell_2| \leq \ell_3 \leq \ell_1 + \ell_2$. Thus, we conclude that

$$\begin{aligned} \langle b_{N_1, n_1}(\omega) b_{N_2, n_2}(\omega) \overline{b_{N_3, n_3}(\omega)} \rangle_\omega &= \delta_{N_3, N_1 + N_2} \sum_{\substack{\ell_1, m_1, s_1 \\ \ell_2, m_2, s_2 \\ s_3}}^{\min(L, \ell_1 + \ell_2)} \sum_{\ell_3 = |\ell_1 - \ell_2|} x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} \overline{x_{\ell_3, m_1 + m_2, s_3}} \\ &\quad \times (-1)^{N_1 + N_2 + m_1 + m_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ N_1 & N_2 & -N_1 - N_2 \end{pmatrix} \\ &\quad \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & -m_1 - m_2 \end{pmatrix} \widehat{\beta}_{\ell_1, s_1; N_1, n_1} \widehat{\beta}_{\ell_2, s_2; N_2, n_2} \overline{\widehat{\beta}_{\ell_3, s_3; N_1 + N_2, n_3}}. \end{aligned} \quad (3.24)$$

Combining (3.24) with (3.23) provides the explicit relation between the third-order autocorrelation and the volume.

Recall that we obtain the autocorrelations of the volume in PSWFs coefficients; see (3.17). Hence, to conclude the derivation we expand

$$\langle a_{I_\omega}^3[\boldsymbol{\ell}_1, \boldsymbol{\ell}_2] \rangle_\omega = \sum_{k, q_1, q_2} \mathbf{a}_x^3[k, q_1, q_2] \psi_{k, q_1}[\boldsymbol{\ell}_1] \overline{\psi_{k, q_2}[\boldsymbol{\ell}_2]},$$

where we only include the block-diagonal terms in the expansion; the rest are equal to zero. Let,

$$\begin{aligned} \Psi_{\ell, N, s}[\boldsymbol{\ell}] &:= \sum_{n=0}^{n_{\max}(N)} \widehat{\beta}_{\ell, s; N, n} \psi_{N, n}[\boldsymbol{\ell}], \\ \rho_{\ell, N, s}^{(k, q)} &:= \int_{\boldsymbol{\ell}} \Psi_{\ell, N, s}([\mathbf{i} + \boldsymbol{\ell}]) \overline{\psi_{k, q}(\boldsymbol{\ell})}. \end{aligned}$$

Then, the final formula reads

$$\begin{aligned}
\mathbf{a}_x^3[k, q_1, q_2] &= \sum_{\substack{\ell_1, m_1, s_1 \\ \ell_2, m_2, s_2 \\ s_3}}^{\min(L, \ell_1 + \ell_2)} \sum_{\ell_3 = |\ell_1 - \ell_2|} x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} \overline{x_{\ell_3, m_1 + m_2, s_3}} \\
&\times (-1)^{m_1 + m_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & -m_1 - m_2 \end{pmatrix} \\
&\times \sum_{N_1 = -\ell_1}^{\ell_1} \sum_{N_2 = -\ell_2}^{\ell_2} (-1)^{N_1 + N_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ N_1 & N_2 & -N_1 - N_2 \end{pmatrix} \\
&\times \frac{1}{L^2} \sum_{\mathbf{i}} \Psi_{\ell_1, N_1, s_1}[\mathbf{i}] \rho_{\ell_2, N_2, s_2}^{(k, q_1)}[\mathbf{i}] \overline{\rho_{\ell_3, N_1 + N_2, s_3}^{(k, q_2)}[\mathbf{i}]}.
\end{aligned} \tag{3.25}$$

In practice, the last two lines of the above expression for $\mathbf{a}_x^3[k, q_1, q_2]$ are precomputed, and both the integration over \mathbf{i} and over ℓ are performed on the grid of the images in the data set, to match the integration performed on the actual images.

3.6 Recovering the volume from the autocorrelations of the micrographs

To estimate the coefficients of the volume itself, we solve the least squares problem

$$\min_{\tilde{x}, \tilde{\gamma}} w_1 |\mathbf{a}_x^1 - \hat{\gamma} \mathbf{a}_{\tilde{x}}^1|^2 + w_2 \|\mathbf{a}_x^2 - \tilde{\gamma} \mathbf{a}_{\tilde{x}}^2\|_2^2 + w_3 \|\mathbf{a}_x^3 - \tilde{\gamma} \mathbf{a}_{\tilde{x}}^3\|_F^2, \tag{3.26}$$

where the explicit expressions of \mathbf{a}_x^1 , \mathbf{a}_x^2 and \mathbf{a}_x^3 are given in (3.21), (3.22) and (3.25), respectively. In the experiments, we set $w_1 = w_2 = w_3 = 1$. We solve the least squares problem using Matlab's `lsqnonlin` solver for nonlinear least squares problems with the trust-regions algorithm. The algorithm is initialized by sampling the coefficients from i.i.d. random Gaussians.

3.7 Complexity analysis

The computational complexity for the computation of the moments from the micrograph is $O(KN^2P^6)$, since we have N^2 patches that we extract from each of K micrographs. For each patch, we perform an expansion in PSWFs with complexity $O(P^3)$ [38], and then compute the autocorrelations dominated by computation of the third-order autocorrelation as in (3.17) costing $O(P^6)$. We note that computing autocorrelations can be executed efficiently on CPUs and GPUs, and in parallel across micrographs. It can even be done in a streaming mode, as only one pass through each micrograph is necessary. The complexity of evaluating the third-order autocorrelation using (3.25) after the precomputation takes $O(BP^6)$, where B is the number of entries in the third-order autocorrelation. Since $B = O(P^3)$, the total complexity for evaluating the third-order autocorrelation is $O(P^9)$. That is because, if the number of volume expansion coefficients is $V = O(P^3)$, the dominant step can be written as a matrix-vector multiplication with a matrix of size $BV \times V$ and a vector $V \times 1$, so the cost is $BV^2 = O(P^9)$. The complexity of solving the least squares problem (3.26) is $O(TP^9)$, where T is the number of iterations required for the optimizer to converge, since the third-order autocorrelation evaluation dominates the cost of each iteration. The optimizer terminates when the norm of the gradient of the cost function in (3.26) decreases below 10^{-6} or the number of iterations exceeds 10^4 .

4 Numerical experiments

The technique we advocate allows recovery of a 3-D structure from its tomographic projections hidden in noisy micrographs without detecting their locations. To illustrate the underlying principles of the method, we present several simple proof-of-concept results for simulated cryo-EM data.² We first present numerical results, and then provide additional technical details.

Numerical results. Figure 2 shows recoveries of the 3-D volumes from the clean autocorrelations and from 300 noisy micrographs of size 7420^2 . The experiments were conducted on the Bovine Pancreatic Trypsin Inhibitor (BPTI) and the TRPV1 molecules (see technical details at the end of this section). In the experiments, we define SNR as $\text{SNR} = \frac{\text{var}(\mathcal{I})}{\sigma^2}$, where $\text{var}(\mathcal{I})$ is the variance of our stack of micrographs and σ^2 is the variance of noise. The noise level in Figure 3 was $\text{SNR} = 1/16$ for the TRPV1 micrographs and $\text{SNR} = 1/64$ for the BPTI micrographs. We present excerpts from noisy micrographs side-by-side with the corresponding clean ones in Figure 3. Figure 4 presents the Fourier shell correlation curves.

Numerical evidence suggests that autocorrelations up to order three, together, uniquely determine the 3-D volume (see also [6, 18] for statistical analysis of closely-related models). Unfortunately, the mapping between the autocorrelations and volume seems to be ill-conditioned, preventing high-resolution recovery from noisy data. For the TRPV1 reconstruction, the optimizer converged to an estimator with relative ℓ_2 error of 10^{-6} in the first three autocorrelations and an error of 10^{-1} in the expansion coefficients of the volume. For the BPTI reconstruction, the errors in the autocorrelations were 10^{-6} , 10^{-7} and 10^{-8} for the third, second and first autocorrelations, respectively, while the error in the expansion coefficients of the volume was 5×10^{-2} . This illustrates the ill-conditioning of the map between the volume and its first three autocorrelations that prevents us from obtaining high-resolution results from noisy data. In the next section, we outline how we suggest overcoming the ill-conditioning in future work.

While we cannot provide a high-resolution 3-D reconstruction from noisy data with the current algorithm, our method can be easily applied to the problem of deciding whether a micrograph contains projections or merely pure noise—a problem considered in classical works in statistics [16] and cryo-EM [27]. This task can be performed by considering solely the recovered γ (the fraction of pixels occupied by projections in the micrograph), estimated as part of the recovery algorithm. Specifically, for this experiment, we used 25 micrographs of size 7420^2 , and the LS problem (3.26) was solved assuming the spherical harmonic cutoff for the volume is $L_{\max} = 0$, which is sufficient to recover a significant $\hat{\gamma}$ in the presence of projections in the micrograph.

Figure 5 presents excerpts of two noisy micrographs, only one of which contains projections. The noise level corresponds to $\text{SNR} = 1/1024$. In the presence of projections, the estimated γ was 0.12, corresponding to approximately 6784 projections. On the other hand, the estimated γ drops to 10^{-5} for the pure noise micrograph, corresponding to less than one projection.

More technical details. The true volume used in the experiments in Figures 5 and 2b was the Bovine Pancreatic Trypsin Inhibitor (BPTI) mutant with altered binding loop sequence, whose atomic model is available in the Protein Data Bank (PDB) as 1QLQ.³ We generated an EM map from this atomic model in UCSF Chimera [47] at a resolution of 5 Å, and cropped it to remove zeros at the boundary to obtain a volume of size 31^3 . For the experiment in Figure 2b, the volume was downsampled to size 20^3 . For the experiment in Figure 2a, we used the TRPV1 in complex with DkTx and RTX, whose EM map is available in the Electron Microscopy Data Bank (EMDB) as EMD-8117.⁴ The original map

²The code to generate all figures is publicly available in <https://github.com/PrincetonUniversity/BreakingDetectionLimit>.

³<https://www.rcsb.org>

⁴<http://www.ebi.ac.uk/pdbe/emdb>

has size 192^3 , and was downsampled to size 20^3 . To generate the ground truth for our reconstructions, we expanded both volumes as in (3.2) with cutoff $L_{\max} = 5$ for TRPV1 and $L_{\max} = 2$ for BPTI. Each volume is described by $\sum_{\ell=0}^{L_{\max}} (2\ell + 1)S(\ell)$ coefficients using the expansion (3.2), where the cutoffs $S(\ell)$ are determined using the Nyquist criterion as in [15].

The micrographs for the experiments were generated as follows. We sample rotation matrices from $SO(3)$ uniformly at random using the QR-based algorithm described in [59], and generate the projection of the volume corresponding to that rotation matrix using ASPIRE [1]. The projections for the experiments were obtained from the smoothed volumes, not the original ones, to ensure that the only sources of error are the noise and our ability to invert the moments via (3.26). We keep track of the indices at which the upper left corner of a projection can be placed without violating the separation condition, so all projections are separated by at least $P - 1$ pixels in each dimension, where the projections are contained in a box of size $P \times P$. The location of the upper left corner of each new projection is picked uniformly at random from the set of available indices. We continue adding projections to the micrograph until no more projections can be added without violating the separation condition.

The experiments were performed on a machine with 40 cores of Intel Xeon E5-2698 v4 @ 2.20GHz with 100 GB of RAM, and took 2 hours per reconstruction. The computation of the moments from each micrograph was performed on a machine with 4 nVidia P100 GPUs with 16 GB of memory each and with 100 GB of RAM. It took 3 minutes per micrograph of size 7420^2 to compute the first three autocorrelations.




Figure 2: Reconstructions from the first, second and third order autocorrelations. The ground truth volumes were expanded according to (3.2) with cutoff L_{\max} . The original molecules are shown in purple and the (smoothed) ground truths in blue to illustrate the smoothing effect of our downsampling and truncation of the spherical harmonics expansion. The reconstructions from clean autocorrelations are shown in yellow, and recoveries from autocorrelations estimated from noisy data in gray. For the noisy experiments, we used 300 micrographs with SNRs of $1/16$ for TRPV1 and $1/64$ for BPTI. We present excerpts from noisy micrographs side-by-side with the corresponding clean ones in Figure 3.

5 Discussion

In this paper, we showed that, in principle, it is possible to estimate a 3-D structure from micrographs, without detecting its projections. Our strategy is to compute autocorrelations of the micrographs and to relate these statistics to the unknown volume’s parameters. Recovering the parameters from the statistics reduces to solving a set of polynomial equations. Crucially, the outlined approach involves no particle picking, hence it might be possible to reconstruct small molecules, particularly, molecules that are too small to be detected in micrographs. In pursuing this research direction, our goal is to significantly increase the range of molecules to which cryo-EM can be successfully applied. Concerns for model bias are also greatly reduced since no template matching is involved. We recognize that




Figure 3: Excerpts of size 250×250 from the noisy micrographs used for the reconstructions in Figure 2 and the corresponding clean excerpts. (a) Excerpt of the clean TRPV1 micrograph; (b) Excerpt of the noisy TRPV1 micrograph with $\text{SNR} = 1/16$; (c) Excerpt from the clean BPTI micrograph; (d) Excerpt from the noisy BPTI micrograph with $\text{SNR} = 1/64$.




Figure 4: Fourier Shell Correlations (FSCs) for the noisy reconstructions presented in Figure 2. The FSC measures the normalized cross-correlation between the recovered 3-D structure and the (smoothed) ground truth 3-D structure over corresponding shells in Fourier space (i.e., as a function of spatial frequency). The resolution is determined as the frequency where the FSC curve drops below 0.5. This is the standard resolution measure in the cryo-EM literature. (a) FSC for the TRPV1 reconstruction, giving resolution of 24\AA ; (b) FSC for the BPTI reconstruction, giving resolution of 13\AA .

significant challenges lay ahead for the implementation of the proposed approach to high-resolution 3-D reconstruction directly from the micrographs. We discuss a few now.

The numerical experiments we have performed reveal that the third-order autocorrelation may not be enough for 3-D reconstruction in practice, due to high sensitivity. One possible remedy might be adding priors to the least squares problem (3.26). For example, one can use data-driven or sparsity aware priors [32, 52, 62, 13, 11], or priors based on the statistical properties of typical proteins [55, 21]. Alternatively, this suggests that fourth-order autocorrelation may be necessary. This in turn would imply that the procedure might require a large amount of data. Recent trends in high-throughput cryo-EM technology give hope that this may be a lesser concern in the long term. Still, large amounts of data also require large amounts of computation.

To reach high-resolution reconstruction, beyond data acquisition and computational challenges, there are modeling issues to consider. In contrast to the simplifying assumptions we have made above, the noise might be colored; the viewing directions of the particles may be distributed non-uniformly; there may be conformational heterogeneity; particles generally do not satisfy our separation condition; and micrographs undergo a contrast transfer function which we have omitted. We believe that these aspects can be handled with the same general strategy: establish a forward model relating the expected



(a) An excerpt of the pure noise micrograph

(b) An excerpt of the clean micrograph in the right panel perturbed by the noise in the left panel.

(c) An excerpt of the clean micrograph

Figure 5: All micrographs are of size 7420^2 pixels and the projections are taken from the BPTI molecule of size 31^3 . The added noise was drawn from i.i.d. Gaussian distribution with zero mean and standard deviation 25, corresponding to an SNR below $1/1024$. The noise realization is identical in both micrographs.

autocorrelations of the micrographs to the target volume(s) and all parameters necessary to model the above effects. The effect of many of these aspects was recently studied in the context of the multi-target detection problem, which can be interpreted as a simplified model of the problem considered in this paper [10, 44, 12, 37]. For example, it was shown that the strict separation condition can be replaced by a parameterized pair-correlation function [36, 34]. Such a function models the distribution of distances between neighboring projections. The observed autocorrelations depend linearly on these parameters, which could be estimated as part of the inverse problem. In addition, alternative computational strategies to autocorrelation analysis, such as approximate expectation-maximization [36, 35], GAN-based techniques [23], and generalized autocorrelation analysis [2, 63], were developed. We believe that these techniques can be adapted to recovering molecular structures directly from the micrographs, and perhaps to alleviate some of the computational challenges. We hope to take care of these issues in future research.

In addition, our technique allows the use of much lower defocus values (a parameter controlled by the user, and which affects the microscope’s point spread function). Lower defocus means lower contrast, but also maintains higher frequency information. Consequently, we may be able to get high resolution reconstructions from fewer micrographs.

From the information-theoretic perspective, it is essential to understand the sample complexity of the problem: how many particle projections are required to estimate the 3-D structure to a desired accuracy? In the closely-related model of multi-reference alignment, it was shown that in the finite-dimensional, low SNR regime, the sample complexity is governed by the moments [46, 6, 4, 3]. If the same is true for our problem, and since we believe that the third-order autocorrelation suffices to determine uniquely the 3-D structure, this will imply that $N = \omega(\sigma^6)$ is a necessary condition for accurate recovery (for a fixed $\gamma > 0$). If the dimension of the problem is high (many parameters are required to describe the 3-D structure), the sample complexity of multi-reference alignment is no longer controlled by the moments, but by the ratio between the dimension and the noise level [48]; we believe a similar phenomenon will hold true for the problem studied in this paper.

Acknowledgment

This work was primarily performed when all authors were at Princeton University. The authors thank Ayelet Heimowitz, Joe Kileel, Ti-Yen Lan, Roy Lederman, Amit Moscovich, Nir Sharon, and Fred Sigworth for helpful discussions, Boris Landa and Yoel Shkolnisky for providing the code for the 2-D PSWFs expansion, and Alberto Bartesaghi, Amitay Eldar, and Yoel Shkolnisky for helping to produce Figure 1. The authors also thank the reviewers for the thoughtful and helpful comments on the manuscript. AS was partially supported by AFOSR FA9550-20-1-0266, the Simons Foundation Math+X Investigator Award, NSF BIGDATA Award IIS-1837992, NSF DMS-2009753, and NIH/NIGMS 1R01GM136780-01. NB was partially supported by NSF award DMS-1719558. TB is partially supported by the BSF grant no. 2020159, the NSF-BSF grant no. 2019752, and the ISF grant no. 1924/21. WL was partially supported by BSF grant no. 2018230 and NSF BIGDATA Award IIS-1837992.

References

- [1] ComputationalCryoEM/ASPIRE-Python: v0.9.2 . <https://doi.org/10.5281/zenodo.5657281>.
- [2] Asaf Abas, Tamir Bendory, and Nir Sharon. The generalized method of moments for multi-reference alignment. *IEEE Transactions on Signal Processing*, 70:1377–1388, 2022.
- [3] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *IEEE Transactions on Information Theory*, 65(6):3565–3584, 2018.
- [4] Emmanuel Abbe, Joao M Pereira, and Amit Singer. Estimation in the group action channel. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 561–565. IEEE, 2018.
- [5] Xiao-chen Bai. Seeing atoms by single-particle cryo-EM. *Trends in biochemical sciences*, 46(4):253–254, 2021.
- [6] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Amelia Perry, Jonathan Weed, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *arXiv preprint arXiv:1712.10163*, 2017.
- [7] Afonso S Bandeira, Yutong Chen, Roy R Lederman, and Amit Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *Inverse Problems*, 36(6):064002, 2020.
- [8] Alberto Bartesaghi, Alan Merk, Soojay Banerjee, Doreen Matthies, Xiongwu Wu, Jacqueline LS Milne, and Sriram Subramaniam. 2.2 Å resolution cryo-EM structure of β -galactosidase in complex with a cell-permeant inhibitor. *Science*, 348(6239):1147–1151, 2015.
- [9] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE Signal Processing Magazine*, 37(2):58–76, 2020.
- [10] Tamir Bendory, Nicolas Boumal, William Leeb, Eitan Levin, and Amit Singer. Multi-target detection with application to cryo-electron microscopy. *Inverse Problems*, 35(10):104003, 2019.
- [11] Tamir Bendory, Yuehaw Khoo, Joe Kileel, Oscar Mickelin, and Amit Singer. Autocorrelation analysis for cryo-EM with sparsity constraints: Improved sample complexity and projection-based algorithms. *arXiv preprint arXiv:2209.10531*, 2022.

- [12] Tamir Bendory, Ti-Yen Lan, Nicholas F Marshall, Iris Rukshin, and Amit Singer. Multi-target detection with rotations. *arXiv preprint arXiv:2101.07709*, 2021.
- [13] Tamir Bendory, Oscar Michelin, and Amit Singer. Sparse multi-reference alignment: Sample complexity and computational hardness. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8977–8981. IEEE, 2022.
- [14] Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J Noble, and Bonnie Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nature methods*, 16(11):1153–1160, 2019.
- [15] Tejal Bhamre, Teng Zhang, and Amit Singer. Anisotropic twicing for single particle reconstruction using autocorrelation analysis. *arXiv preprint arXiv:1704.07969*, 2017.
- [16] David Donoho and Jiashun Jin. Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics*, 32(3):962–994, 2004.
- [17] Amitay Eldar, Boris Landa, and Yoel Shkolnisky. KLT picker: particle picking using data-driven optimal templates. *Journal of structural biology*, 210(2):107473, 2020.
- [18] Zhou Fan, Roy R Lederman, Yi Sun, Tianhao Wang, and Sheng Xu. Maximum likelihood for high-noise group orbit estimation and single-particle cryo-EM. *arXiv preprint arXiv:2107.01305*, 2021.
- [19] Joachim Frank. *Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state*. Oxford university press, 2006.
- [20] Joachim Frank. Single-particle reconstruction of biological molecules—story in a sample (Nobel lecture). *Angewandte Chemie International Edition*, 57(34):10826–10841, 2018.
- [21] Marc Aurèle Gilles and Amit Singer. A molecular prior distribution for Bayesian inference based on Wilson statistics. *Computer Methods and Programs in Biomedicine*, page 106830, 2022.
- [22] Robert M Glaeser. Electron crystallography: present excitement, a nod to the past, anticipating the future. *Journal of structural biology*, 128(1):3–14, 1999.
- [23] Harshit Gupta, Michael T McCann, Laurene Donati, and Michael Unser. CryoGAN: a new reconstruction paradigm for single-particle cryo-EM via deep adversarial learning. *IEEE Transactions on Computational Imaging*, 7:759–774, 2021.
- [24] Ayelet Heimowitz, Joakim Andén, and Amit Singer. APPLE picker: Automatic particle picking, a low-effort cryo-EM framework. *Journal of structural biology*, 204(2):215–227, 2018.
- [25] Ayelet Heimowitz, Joakim Andén, and Amit Singer. Reducing bias and variance for CTF estimation in single particle cryo-EM. *Ultramicroscopy*, 212:112950, 2020.
- [26] Richard Henderson. The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological molecules. *Quarterly reviews of biophysics*, 28(2):171–193, 1995.
- [27] Richard Henderson. Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise. *Proceedings of the National Academy of Sciences*, 110(45):18037–18041, 2013.

- [28] Mark A Herzik, Mengyu Wu, and Gabriel C Lander. High-resolution structure determination of sub-100 kDa complexes using conventional cryo-EM. *Nature communications*, 10(1):1–9, 2019.
- [29] Andrii Iudin, Paul K Korir, José Salavert-Torres, Gerard J Kleywegt, and Ardan Patwardhan. EMPIAR: a public archive for raw electron microscopy image data. *Nature methods*, 13(5):387–388, 2016.
- [30] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [31] Maryam Khoshouei, Mazdak Radjainia, Wolfgang Baumeister, and Radostin Danev. Cryo-EM structure of haemoglobin at 3.2 Å determined with the Volta phase plate. *Nature communications*, 8(1):1–6, 2017.
- [32] Dari Kimanius, Gustav Zickert, Takanori Nakane, Jonas Adler, Sebastian Lunz, C-B Schönlieb, Ozan Öktem, and Sjors HW Scheres. Exploiting prior knowledge about biological macromolecules in cryo-EM structure determination. *IUCrJ*, 8(1):60–75, 2021.
- [33] A Klug and RA Crowther. Three-dimensional image reconstruction from the viewpoint of information theory. *Nature*, 238(5365):435–440, 1972.
- [34] Shay Kreymer and Tamir Bendory. Two-dimensional multi-target detection: An autocorrelation analysis approach. *IEEE Transactions on Signal Processing*, 70:835–849, 2022.
- [35] Shay Kreymer, Amit Singer, and Tamir Bendory. An approximate expectation-maximization for two-dimensional multi-target detection. *IEEE Signal Processing Letters*, 2022.
- [36] Ti-Yen Lan, Tamir Bendory, Nicolas Boumal, and Amit Singer. Multi-target detection with an arbitrary spacing distribution. *IEEE Transactions on Signal Processing*, 68:1589–1601, 2020.
- [37] Ti-Yen Lan, Nicolas Boumal, and Amit Singer. Random conical tilt reconstruction without particle picking in cryo-electron microscopy. *Acta Crystallographica Section A*, 78(4):294–301, 2022.
- [38] Boris Landa and Yoel Shkolnisky. Steerable principal components for space-frequency localized images. *SIAM journal on imaging sciences*, 10(2):508–534, 2017.
- [39] Boris Landa and Yoel Shkolnisky. The steerable graph laplacian and its application to filtering image datasets. *SIAM Journal on Imaging Sciences*, 11(4):2254–2304, 2018.
- [40] Gabriel C Lander and Robert M Glaeser. Conquer by cryo-EM without physically dividing. *Biochemical Society Transactions*, 49(5):2287–2298, 2021.
- [41] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1569–1573. IEEE, 2018.
- [42] Yi-Lynn Liang, Maryam Khoshouei, Mazdak Radjainia, Yan Zhang, Alisa Glukhova, Jeffrey Tarasch, David M Thal, Sebastian GB Furness, George Christopoulos, Thomas Coudrat, et al. Phase-plate cryo-EM structure of a class b GPCR–G-protein complex. *Nature*, 546(7656):118–123, 2017.
- [43] Yuxi Liu, Shane Gonen, Tamir Gonen, and Todd O Yeates. Near-atomic cryo-EM imaging of a small protein displayed on a designed scaffolding system. *Proceedings of the National Academy of Sciences*, 115(13):3362–3367, 2018.

- [44] Nicholas F Marshall, Ti-Yen Lan, Tamir Bendory, and Amit Singer. Image recovery from rotational and translational invariants. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5780–5784. IEEE, 2020.
- [45] Frank Natterer. *The mathematics of computerized tomography*. SIAM, 2001.
- [46] Amelia Perry, Jonathan Weed, Afonso S Bandeira, Philippe Rigollet, and Amit Singer. The sample complexity of multireference alignment. *SIAM Journal on Mathematics of Data Science*, 1(3):497–517, 2019.
- [47] Eric F Pettersen, Thomas D Goddard, Conrad C Huang, Gregory S Couch, Daniel M Greenblatt, Elaine C Meng, and Thomas E Ferrin. UCSF Chimera—a visualization system for exploratory research and analysis. *Journal of computational chemistry*, 25(13):1605–1612, 2004.
- [48] Elad Romanov, Tamir Bendory, and Or Ordentlich. Multi-reference alignment in high dimensions: sample complexity and phase transition. *SIAM Journal on Mathematics of Data Science*, 3(2):494–523, 2021.
- [49] Giovanna Scapin, Clinton S Potter, and Bridget Carragher. Cryo-EM for small molecules discovery, design, understanding, and application. *Cell chemical biology*, 25(11):1318–1325, 2018.
- [50] Sjors HW Scheres. Semi-automated selection of cryo-EM particles in RELION-1.3. *Journal of structural biology*, 189(2):114–122, 2015.
- [51] Osip Schwartz, Jeremy J Axelrod, Sara L Campbell, Carter Turnbaugh, Robert M Glaeser, and Holger Müller. Laser phase plate for transmission electron microscopy. *Nature methods*, 16(10):1016–1020, 2019.
- [52] Jonathan Shani, Raja Giryes, and Tamir Bendory. Denoiser-based projections for 2-D super-resolution multi-reference alignment. *arXiv preprint arXiv:2204.04754*, 2022.
- [53] Nir Sharon, Joe Kileel, Yuehaw Khoo, Boris Landa, and Amit Singer. Method of moments for 3D single particle ab initio modeling with non-uniform distribution of viewing angles. *Inverse Problems*, 36(4):044003, 2020.
- [54] Maxim Shatsky, Richard J Hall, Steven E Brenner, and Robert M Glaeser. A method for the alignment of heterogeneous macromolecules from electron microscopy. *Journal of structural biology*, 166(1):67–78, 2009.
- [55] Amit Singer. Wilson statistics: derivation, generalization and applications to electron cryomicroscopy. *Acta Crystallographica Section A: Foundations and Advances*, 77(5), 2021.
- [56] Amit Singer and Fred J Sigworth. Computational methods for single-particle electron cryomicroscopy. *Annual Review of Biomedical Data Science*, 3:163–190, 2020.
- [57] David Slepian. Prolate spheroidal wave functions, Fourier analysis and uncertainty—IV: extensions to many dimensions; generalized prolate spheroidal functions. *Bell System Technical Journal*, 43(6):3009–3057, 1964.
- [58] Carlos Oscar S Sorzano, A Jiménez, Javier Mota, José Luis Vilas, David Maluenda, Marta Martínez, Erney Ramírez-Aportela, Tomas Majtner, J Segura, Ruben Sánchez-García, et al. Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. *Acta Crystallographica Section F: Structural Biology Communications*, 75(1):19–32, 2019.

- [59] Gilbert W Stewart. The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM Journal on Numerical Analysis*, 17(3):403–409, 1980.
- [60] Marin van Heel. Finding trimeric HIV-1 envelope glycoproteins in random noise. *Proceedings of the National Academy of Sciences*, 110(45):E4175–E4177, 2013.
- [61] Marin van Heel, Michael Schatz, and Elena Orlova. Correlation functions revisited. *Ultramicroscopy*, 46(1-4):307–316, 1992.
- [62] Cédric Vonesch, Lanhui Wang, Yoel Shkolnisky, and Amit Singer. Fast wavelet-based single-particle reconstruction in cryo-EM. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1950–1953. IEEE, 2011.
- [63] Ran Weber, Asaf Abas, Shay Kreymer, Tamir Bendory, et al. Generalized autocorrelation analysis for multi-target detection. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5907–5911. IEEE, 2022.
- [64] Wilson Wong, Xiao-chen Bai, Alan Brown, Israel S Fernandez, Eric Hanssen, Melanie Condron, Yan Hong Tan, Jake Baum, and Sjors HW Scheres. Cryo-EM structure of the Plasmodium falciparum 80S ribosome bound to the anti-protozoan drug emetine. *Elife*, 3:e03080, 2014.
- [65] Mengyu Wu and Gabriel C Lander. How low can we go? structure determination of small biological complexes using single-particle cryo-EM. *Current opinion in structural biology*, 64:9–16, 2020.
- [66] Kaiming Zhang, Shanshan Li, Kalli Kappel, Grigore Pintilie, Zhaoming Su, Tung-Chung Mou, Michael F Schmid, Rhiju Das, and Wah Chiu. Cryo-EM structure of a 40 kDa SAM-IV riboswitch RNA at 3.7 Å resolution. *Nature communications*, 10(1):1–6, 2019.
- [67] Zhizhen Zhao, Yoel Shkolnisky, and Amit Singer. Fast steerable principal component analysis. *IEEE Transactions on Computational Imaging*, 2(1):1–12, 2016.

A Proof of Proposition 2.1

The proof is based on a variant of the Neyman-Pearson Lemma to derive the best (deterministic) estimator $\hat{\eta}$. Take any deterministic estimator $\hat{\eta} = \hat{\eta}(X)$, with values in $\{0, 1\}$; then $\hat{\eta}$ is characterized by S , defined to be the set of X 's where $\hat{\eta} = 1$. We write Prob_i to mean the probability conditional on the event $\eta = i$; that is, $\text{Prob}_i[A] = \text{Prob}[A|\eta = i]$. Then, the probability that $\hat{\eta}$ fails is:

$$\begin{aligned}
 \text{Prob}[\hat{\eta} \neq \eta] &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q)\text{Prob}_1[\hat{\eta} = 0] \\
 &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q)(1 - \text{Prob}_1[\hat{\eta} = 1]) \\
 &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q) - (1 - q)\text{Prob}_1[\hat{\eta} = 1] \\
 &= (1 - q) + \int_S (qf_0(x) - (1 - q)f_1(x))dx,
 \end{aligned} \tag{A.1}$$

where $f_i(x)$ is the normal density with mean θ_i and variance σ^2 . The best estimator of η based on X minimizes the failure probability; hence, it minimizes the integral in (A.1) through an appropriate choice of the set S . This is achieved by picking all x 's such that the integrand is nonpositive:

$$S = \{x : qf_0(x) - (1 - q)f_1(x) \leq 0\}.$$

With $\Lambda(x) = f_0(x)/f_1(x)$ and $b = (1 - q)/q$, the corresponding estimator is:

$$\hat{\eta} = \begin{cases} 1 & \text{if } \Lambda(x) \leq b, \\ 0 & \text{if } \Lambda(x) > b. \end{cases}$$

Taking logarithms, the set S can be rewritten as the set of x 's where:

$$-\|x - \theta_0\|^2 \leq -\|x - \theta_1\|^2 + 2\sigma^2 \log(b),$$

or equivalently

$$\langle x, \theta_1 - \theta_0 \rangle \geq \frac{\|\theta_1\|^2 - \|\theta_0\|^2}{2} - \sigma^2 \log(b).$$

Now let us compute the probability of failure conditional on the event $\eta = 0$. In this case, failure occurs when $X \in S$. Since $X|(\eta = 0) \sim N(\theta_0, \sigma^2)$, we can write $X|(\eta = 0) = \sigma Z + \theta_0$, where $Z \sim N(0, I)$. On that condition,

$$\begin{aligned} \langle X, \theta_1 - \theta_0 \rangle &= \sigma \langle Z, \theta_1 - \theta_0 \rangle + \langle \theta_0, \theta_1 - \theta_0 \rangle \\ &= \sigma \langle Z, \theta_1 - \theta_0 \rangle + \langle \theta_0, \theta_1 \rangle - \|\theta_0\|^2, \end{aligned}$$

and failure occurs when

$$\sigma \langle Z, \theta_1 - \theta_0 \rangle \geq \frac{\|\theta_1\|^2 + \|\theta_0\|^2}{2} - \langle \theta_0, \theta_1 \rangle - \sigma^2 \log(b) = \frac{1}{2} \|\theta_1 - \theta_0\|^2 - \sigma^2 \log(b).$$

Define $Y = \langle Z, \theta_1 - \theta_0 \rangle \sim N(0, \|\theta_1 - \theta_0\|^2)$ and divide through by σ . The above event is equivalent to:

$$Y \geq \frac{c}{\sigma} - \sigma \log(b),$$

where $c = \|\theta_1 - \theta_0\|^2/2$. For simplicity, let us assume $\|\theta_1 - \theta_0\| = 1$, so that $Y \sim N(0, 1)$. Then,

$$\text{Prob}_0[\hat{\eta} = 1] = \text{Prob} \left[Y \geq \frac{c}{\sigma} - \sigma \log(b) \right], \quad Y \sim N(0, 1).$$

Similarly,

$$\text{Prob}_1[\hat{\eta} = 0] = \text{Prob} \left[Y \geq \frac{c}{\sigma} + \sigma \log(b) \right], \quad Y \sim N(0, 1).$$

Thus, the overall probability of failure is:

$$\text{Prob}[\hat{\eta} \neq \eta] = q \text{Prob} \left[Y \geq \frac{c}{\sigma} - \sigma \log(b) \right] + (1 - q) \text{Prob} \left[Y \geq \frac{c}{\sigma} + \sigma \log(b) \right].$$

Now, if $q = 1/2$, then $\log(b) = 0$. Hence the probability of failure is simply:

$$\text{Prob} \left[Y \geq \frac{c}{\sigma} \right] \longrightarrow \frac{1}{2} = q \quad \text{as } \sigma \rightarrow \infty.$$

If $q > 1/2$, then $q > 1 - q$ and $\log(b) < 0$. Consequently,

$$\text{Prob} \left[Y \geq \frac{c}{\sigma} - \sigma \log(b) \right] \longrightarrow 0,$$

while

$$\text{Prob} \left[Y \geq \frac{c}{\sigma} + \sigma \log(b) \right] \longrightarrow 1,$$

as $\sigma \rightarrow \infty$. Hence,

$$\text{Prob}[\hat{\eta} \neq \eta] \longrightarrow 1 - q \quad \text{as } \sigma \rightarrow \infty.$$

That is, the probability of success converges to q . Finally, if $q < 1/2$, then $\log(b) > 0$ and a similar reasoning shows the probability of success converges to $1 - q$. In all cases, the probability of success of the best possible deterministic estimator converges to $\max(q, 1 - q)$.