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

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

^aThe Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA

^bDepartment of Mathematics, Princeton University, Princeton, NJ, USA

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

September 29, 2018

Abstract

Single-particle cryo-electron microscopy (cryo-EM) has recently joined X-ray crystallography and NMR spectroscopy as a high-resolution structural method for 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, which can be particularly hard to detect. In this paper, we propose a radically different approach: we contend that the structure could, in principle, be reconstructed directly from the micrographs, without intermediate detection. As a result, even small molecules should be within reach for cryo-EM. To support this claim, we setup a simplified mathematical model and demonstrate how our autocorrelation analysis technique allows to go directly from the micrographs to the sought signals. This involves only one pass over the micrographs, which is desirable for large experiments. We show numerical results and discuss challenges that lay ahead to turn this proof-of-concept into a competitive alternative 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 (SPR) 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. The goal is to construct 3-D models of the molecular structure from the micrographs.

The signal to noise ratio (SNR) of the 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 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.

Figure 1: Example of micrographs of size 250×250 with additive white Gaussian noise of variance σ^2 for increasing values of σ . Each micrograph contains the same four occurrences of a 50×50 image of Einstein. In panel (c), the noise level is such that it is very challenging to locate the occurrences of the planted image. In fact, it can be shown that at low SNR, reliable detection of individual image occurrences is impossible, even if the true image is known. By analogy to cryo-EM, this depicts a scenario where particle picking cannot be done.

Figure 2: Recovery of Einstein from micrographs at noise level $\sigma = 3$ (see Figure 1(c)). Averaged autocorrelations of the micrographs allow to estimate the power spectrum of the target image. This does not require particle picking. A phase retrieval algorithm (RRR) produces the estimates here shown, initialized randomly. Estimates are obtained from 2×10^2 , 2×10^3 , 2×10^4 , 2×10^5 micrographs (growing across panels from left to right) of size 4096×4096 , each containing 700 image occurrences on average.

All contemporary methods in the field 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*. 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 1 illustrates how particle picking becomes increasingly challenging as the SNR degrades.

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, Richard Henderson [18] 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., [15]), it was established that particle picking is impossible for molecules below a certain weight (below ~ 50 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 $> 3/[Contrast^2 \times Resolution \text{ (as length)} \times Critical Electron Dose]$* ” [13]. As these two leaders of the cryo-EM community 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. Yet, imaging such molecules is important for structure-guided drug design and other applications [34].

The unique issues raised by small particles have been mitigated by recent technological advances in the field, including the use of Volta phase plates [22, 27] and scaffolding cages [29]. Despite this progress, detecting small molecules in the micrographs remains a challenge. 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.

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 the limits particle picking imposes on molecule size do not necessarily translate into limits on particle reconstruction.

As a proof of concept, we first study a toy model for which it is easier to convey the mathematical principles. In this model, an unknown image appears multiple times at unknown locations in each of several micrographs, each affected by additive Gaussian noise—see Figure 1 for an illustration. The goal is to estimate the planted image. The task is challenging in particular when the SNR is low enough that particle picking (identifying the locations of each image in each micrograph) cannot be done reliably. This problem is interesting on its own as it appears in other scientific applications, including spike sorting [26], passive radar [16] and system identification [30].

In order to recover the image, we use autocorrelation analysis. Specifically, we relate the autocorrelations of the micrographs to the autocorrelations of the image. For any noise level, these autocorrelations can be estimated to any desired accuracy, provided that we observe sufficiently many image occurrences and the latter are separated in the micrograph. Importantly, there is no need to detect individual image occurrences. The autocorrelations of the micrographs are straightforward to compute and require only one pass over the data. After estimation of the density of particles in the micrographs, these directly yield estimates for the autocorrelations of the target image itself. To estimate the image itself from its estimated autocorrelations, we solve a nonlinear inverse problem; see for instance Figure 2.

Beyond this 2-D proof of concept, we look toward 3-D reconstruction as well. Zvi Kam [21] first proposed autocorrelation analysis for 3-D reconstruction, under the assumption of perfect particle picking: his method used autocorrelations of the picked, perfectly centered, particles. In contrast, we derive the mathematical relation between the autocorrelations of the micrographs as a whole and the 3-D volume, under some simplifying conditions. 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 [37, 43, 19, 42]. 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 [37], 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, no templates or human intervention are required, thus significantly reducing concerns about model bias.

2 Proof of concept: A mathematical toy model

In this section, we present a toy model in order to introduce the mathematical principles enabling estimation of a signal in a low SNR regime, even when detection is impossible. Later, we discuss how these principles carry through for the SPR problem. We first formulate the model for 1-D signals to ease exposition.

Let $x \in \mathbb{R}^L$ be the target signal and let $y \in \mathbb{R}^N$ be the observed data, where we assume N is far larger than L . Let $s \in \{0, 1\}^{N-L+1}$ be a binary signal indicating (with 1's) the starting positions of all occurrences of x in y , so that, with additive white Gaussian noise:

$$y = x * s + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I_N), \quad (1)$$

where $*$ denotes linear convolution. While both x and s are unknown, the goal is only to estimate x from y . This is a key difference with other works on *blind deconvolution*, a longstanding problem arising in a variety of engineering and scientific applications such as astronomy, communication, image deblurring, system identification and optics; see [20, 36, 3, 2], just to name a few. The parameters of the signal s (the locations of its nonzero values) are the *nuisance variables* of the problem. As will be shown next, in a low SNR environment, estimating s is impossible, whereas estimating x is tractable under some conditions on s .

We assume the binary signal s obeys a *separation condition*:

$$\text{If } s[i] = 1 \text{ and } s[j] = 1 \text{ for } i \neq j, \text{ then } |i - j| \geq 2L - 1. \quad (2)$$

In words: the starting positions of any two occurrences must be separated by at least $2L - 1$ positions, so that their end points are necessarily separated by at least $L - 1$ signal-free entries in the micrograph.

Our problem can be interpreted as a special case of the *system identification* problem. Similarly to (1), the forward model takes the form $y = x * w + \varepsilon$, where x is the unknown signal (the system's impulse response), w is an unknown (often random) input sequence, and ε is an additive noise. The goal of this problem is to estimate x , usually referred to as “identifying the system.” The question of identifiability of x under this observation model is addressed for certain Gaussian and non-Gaussian w in [6, 23]. In the special case where w is binary and satisfies our separation condition, we recover our model.

Likelihood-based methods estimate x as the maximizer of some function $f(x|y)$, where f is derived from the likelihood function of x given the observed signal y . If some prior is assumed on x , then $f(x|y)$ can be taken to be the posterior distribution of x given the data; this is the simplest form of Bayesian inference. Methods based on such formulations are popular nowadays in cryo-EM; see for instance [38, 35]. Optimizing the function $f(x|y)$ exactly is often intractable, and thus heuristic methods are used instead. One proposed technique is to use Markov Chain Monte Carlo (MCMC) [10]. In special cases, including the case where w is binary, expectation maximization (EM) has been used [10]. The EM method for discrete w is based on a certain “forward-backward” procedure used in hidden Markov models [33]. However, the complexity of this procedure is superlinear in N , and therefore its usage is limited for big data sets.

Because likelihood methods are computationally expensive, methods based on recovery from moments have also been previously used for system identification. Methods based on the third- and fourth-order moments are described and analyzed in [28, 14, 41]. Building on such ideas, we focus on an autocorrelation analysis for (1).

3 The detection limit

In the low SNR regime—even if x is known—estimating the binary sparse signal s is impossible, that is, one cannot reliably detect occurrences of x in the micrograph y . To support this claim, we consider a strictly simpler problem: suppose an oracle identifies for us one interval of length L 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, that is, to determine whether the corresponding entry of s is 0 or 1. The oracle further provides the signal x , the probability q that the interval contains signal, and the noise variance σ^2 .

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 $X \in \mathbb{R}^L$ (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 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 B.

Proposition 3.1. *For any deterministic estimator $\hat{\eta}$ of η ,*

$$\lim_{\sigma \rightarrow \infty} \text{Prob}[\hat{\eta} = \eta] \leq \max(q, 1 - q); \quad (3)$$

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

Proposition 3.1 implies that in order to estimate the signal at low SNR we must consider methods that aim to estimate the signal x directly, without estimating the nuisance variable s as an intermediate step. In the next sections, we consider autocorrelation analysis for that purpose.

4 Autocorrelation analysis

In general, for a signal z of length m , the autocorrelation of order $q = 1, 2, \dots$ is given for any integer shifts $\ell_1, \dots, \ell_{q-1}$ by

$$a_z^q[\ell_1, \dots, \ell_{q-1}] = \frac{1}{m} \sum_{i=-\infty}^{+\infty} z[i]z[i + \ell_1] \cdots z[i + \ell_{q-1}], \quad (4)$$

where indexing of z out of the range $0, \dots, m - 1$ is zero-padded. Explicitly, the first-, second- and third-order autocorrelations are given by:

$$\begin{aligned} a_z^1 &= \frac{1}{m} \sum_{i=0}^{m-1} z[i], \\ a_z^2[\ell] &= \frac{1}{m} \sum_{i=\max\{0, -\ell\}}^{m-1+\min\{0, -\ell\}} z[i]z[i + \ell], \\ a_z^3[\ell_1, \ell_2] &= \frac{1}{m} \sum_{i=\max\{0, -\ell_1, -\ell_2\}}^{m-1+\min\{0, -\ell_1, -\ell_2\}} z[i]z[i + \ell_1]z[i + \ell_2]. \end{aligned} \quad (5)$$

The autocorrelation functions have symmetries. Specifically, $a_z^2[\ell] = a_z^2[-\ell]$, and $a_z^3[\ell_1, \ell_2] = a_z^3[\ell_2, \ell_1] = a_z^3[-\ell_1, \ell_2 - \ell_1]$. For our purposes, this will be applied both to x (of length L) and to y (of length N).

Under the separation condition, the relation between autocorrelations of the micrograph and those of x is particularly simple, as we now show. It is useful to introduce some notation: let M denote the number of occurrences of x in y (that is, the number of 1's in s), and let

$$\gamma = \frac{ML}{N} \quad (6)$$

denote the density of x in y (that is, the fraction of entries of y occupied by occurrences of x .) The separation condition imposes $\gamma \leq \frac{L}{2L-1} \approx 1/2$.

For shifts in $0, \dots, L-1$, the autocorrelation functions of y depend on the corresponding autocorrelations of x , the noise level σ and the support signal s . Importantly, under the separation condition (2), the dependency on s is only through the density γ . We consider the asymptotic regime where γ remains constant; that is, as N goes to infinity, M also goes to infinity at the same rate (in other words, as we see an increasingly large micrograph, it contains increasingly many signal occurrences, with constant signal density). In that regime, the law of large numbers can be used to show the following statement:

$$\lim_{N \rightarrow \infty} a_y^1 \stackrel{\text{a.s.}}{=} \gamma a_x^1, \quad (7)$$

where equality holds almost surely (a.s.), meaning it holds with probability one. The randomness is over the Gaussian noise ε ; s may be deterministic. Thus, given enough data, if γ is known, we can estimate a_x^1 from y . (We show later how to estimate γ as well.)

We have a similar observation for the second-order autocorrelation: $a_y^2[\ell]$ computes the correlation between y and a copy of y shifted by ℓ entries. Considering ℓ only in the range $0, \dots, L-1$, one can see that any given occurrence of x in y is only ever correlated with itself, and never with another occurrence. As a result,

$$\lim_{N \rightarrow \infty} a_y^2[\ell] \stackrel{\text{a.s.}}{=} \gamma a_x^2[\ell] + \sigma^2 \delta[\ell], \quad (8)$$

for $\ell = 0, \dots, L-1$, where $\delta[\ell]$ equals one for $\ell = 0$ and zero otherwise. The last part captures the autocorrelation of the noise. Notice that, even if σ is unknown, entries $\ell = 1, \dots, L-1$ still provide useful information about a_x^2 .

Along the same lines, one can establish a relation for third-order autocorrelations:

$$\lim_{N \rightarrow \infty} a_y^3[\ell_1, \ell_2] \stackrel{\text{a.s.}}{=} \gamma a_x^3[\ell_1, \ell_2] + \sigma^2 \gamma a_x^1 (\delta[\ell_1, 0] + \delta[0, \ell_2] + \delta[\ell_1, \ell_2]), \quad (9)$$

for $\ell_1, \ell_2 = 0, \dots, L-1$, where $\delta[\ell_1, \ell_2] := \delta[\ell_1 - \ell_2]$. Here too, few entries are affected by σ in the limit. See Appendix C.

Computing the autocorrelations of the micrograph is straightforward. The natural question, treated next, is whether one can recover x from them.

5 Recovering a signal from autocorrelations

A one-dimensional signal is determined uniquely by its second- and third-order autocorrelations. Indeed, assuming $z[0]$ and $z[L-1]$ are nonzero (otherwise, redefine the length of the signal), we can recover z explicitly using this identity for $k = 0, \dots, L-1$:

$$z[k] = \frac{z[0]z[k]z[L-1]}{z[0]z[L-1]} = \frac{a_z^3[k, L-1]}{a_z^2[L-1]}. \quad (10)$$

This proves the following useful fact:

Proposition 5.1. *A signal $z \in \mathbb{R}^L$ is determined uniquely from a_z^2 and a_z^3 .*

A couple of remarks are in order. First, (10) is not numerically stable: if $z[0]$ or $z[L-1]$ are close to 0, recovery of z is sensitive to errors in the autocorrelations. In practice, we recover z by fitting it to its autocorrelations using a nonconvex least-squares (LS) procedure, which is empirically more robust to additive noise; we have observed similar phenomena for related problems [5, 9, 1]. Second, note that the second-order autocorrelation is not by itself sufficient to determine the signal uniquely. However, for dimensions greater than one¹, almost all signals are determined uniquely up to sign (or phase for complex signals) and reflection through the origin (with conjugation in the complex case) [17]. The sign ambiguity can be resolved by the mean of the signal if it is not zero. However, determining the reflection symmetry still requires additional information, beyond the second-order autocorrelation.

The observed moments a_y^1, a_y^2 and a_y^3 of y do not immediately yield the moments of the signal x , as seen by (7), (8) and (9); rather, the two are related by the noise level σ and the ratio γ . We will show, however, that x is still determined by the observed moments of y .

First, we observe that if the noise level σ is known, generally, one can estimate γ from the first two moments of the micrograph. The proof is provided in Appendix D.

Proposition 5.2. *Let $\sigma > 0$ be fixed and assume that the separation condition (2) holds. If the mean of x is nonzero, then*

$$\gamma \stackrel{a.s.}{=} \lim_{N \rightarrow \infty} \frac{L(a_y^1)^2}{a_y^2[0] + 2 \sum_{\ell=1}^{L-1} a_y^2[\ell] - \sigma^2}. \quad (11)$$

Using third-order autocorrelation information of y , both the ratio γ and the noise σ are determined. For the following results, when we say that a result holds for a “generic” signal x , we mean that the set of signals which cannot be determined by these measurements has Lebesgue measure zero. In particular, this means that we can recover almost all signals with the given measurements. The proof is provided in Appendix E.

Proposition 5.3. *Assume $L \geq 3$ and assume that the separation condition (2) holds. In the limit of $N \rightarrow \infty$, the observed autocorrelations a_y^1, a_y^2 and a_y^3 determine the ratio γ and noise level σ uniquely for a generic signal x . If $\gamma > \frac{1}{4}$, then this holds for any signal x with nonzero mean.*

From Propositions 5.1 and 5.3 we deduce the following:

Corollary 5.4. *In the limit of $N \rightarrow \infty$ and under the separation condition (2), the signal x , the ratio γ , and the noise level σ are determined from the first three autocorrelation functions of y if either the signal x is generic or x has nonzero mean and $\gamma > \frac{1}{4}$.*

As a side note, under the separation condition, the length L of the signal can also be determined from the autocorrelations in the asymptotic regime, by inspection of the support of a_y^2 .

6 Numerical experiments

The technique we advocate allows recovery of a signal hidden in noisy micrographs without detecting the location of the signals embedded in these micrographs. To illustrate the underlying principles of the method, we present several numerical examples for the toy model (1) and a simple proof of concept for simulated cryo-EM data. Appendix A provides additional details on the experiments. The code to generate all figures is publicly available in <https://github.com/PrincetonUniversity/BreakingDetectionLimit>.

¹Autocorrelation in d dimensions is defined similarly to (4), where $i, \ell_1, \dots, \ell_{q-1}$ take on values in $\mathbb{Z}^d$ and the summation in i is over all of $\mathbb{Z}^d$.

6.1 Toy model

In the first experiment, we estimated an 50-by-50 pixel image of Einstein with mean zero from a growing number of micrographs, each of size 4096×4096 pixels. Each micrograph contains, on average, 700 occurrences of the target image at random locations. Thus, about 10% of each micrograph contains signal. The micrographs are contaminated with additive white Gaussian noise with standard deviation $\sigma = 3$, corresponding to $\text{SNR} = \frac{M\|x\|_F^2}{\sigma^2 N} \approx 1/370$. This high noise level is illustrated in Figure 1. To simplify the experiment, we assume the number of signal occurrences and the noise standard deviation are known. Micrographs are generated such that any two occurrences are always separated by at least 49 pixels in each direction in accordance with the separation condition (2).

We compute the average second-order autocorrelation of the micrographs. This is a particularly simple computation which can be efficiently executed with a fast Fourier transform (FFT) in parallel. Given the noise level and number of image repetitions, the second-order autocorrelation of the image can be easily deduced from (8). Then, to estimate the target image, we resort to a standard phase retrieval algorithm called relaxed-reflect-reflect (RRR) [12], initialized randomly. Relative error is measured as the ratio of the root mean square error to the norm of the ground truth (square root of the sum of squared pixel intensities).

Figure 2 shows several estimated images for a growing number of micrographs. Figure 3 presents the normalized recovery error as a function of the amount of data available. This is computed after fixing the reflection symmetries (see Section 4). As evidenced by these figures, the ground truth image can be estimated increasingly well from increasingly many micrographs, without particle picking.

Figure 3: Relative error curves for the experiment of Figure 2.

In practice, we do not expect to know γ and maybe not even σ . Figure 4 shows recovery of a 1-D signal from the first three autocorrelations of the data. The autocorrelations are computed from noisy micrographs with $\sigma = 3$ that obey the separation condition of (2); see Figure 5. Both the signal and γ are estimated simultaneously from the observed autocorrelations by LS fitting; more details are given in Appendix A. Crucially, the LS does not require knowledge of σ .

6.2 Cryo-EM

The application of autocorrelation analysis to cryo-EM follows the same mathematical principles. The derivation of the first three autocorrelations of the micrographs and their relations to the volume itself

Figure 4: Recovery of a 1-D signal of length $L = 21$ at noise level $\sigma = 3$ from M repetitions. The length of the micrograph was set to be $N = 410 * M$. The first three panels (left to right) show reconstruction with different M values compared to the ground truth signal (in blue). The true value of γ is 0.0512. The relative errors of γ for the three panels are: 4.8%, 4%, 1.2%. The right panel shows the relative error as a function of M .

Figure 5: An example for a 1-D noisy measurement for the experiment in Figure 4. The signal occurrences appear in blue, and the noisy data in green.

are provided in Appendix F. In particular, numerical evidence suggests that the third-order autocorrelation uniquely determines the 3-D volume. Figure 6 shows recovery of the 3-D volume from the clean autocorrelations and from noisy micrographs. Excerpts from the noisy micrographs are shown in Figure 7. Unfortunately, the mapping between the autocorrelations and volume seems to be ill-conditioned, preventing high-resolution recovery from noisy data. The details of the reconstruction algorithm are given in Appendix A. In the next section, we outline how we suggest to overcome 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 [11] and cryo-EM [19]. 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. Figure 9 presents excerpts of two noisy micrographs, only one of which contains projections. 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.

Figure 6: Reconstructions from the first, second and third order autocorrelations. The ground truth volumes were expanded according to (25) 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 7 (see appendix).

Figure 7: Excerpts from the noisy micrographs used for the reconstructions in Figure 6 and the corresponding clean excerpts. (a) Excerpt of the clean TRPV1 micrograph; (b) Excerpt of the noisy TRPV1 micrograph with SNR = 1/16; (c) Excerpt from the clean BPTI micrograph; (d) Excerpt from the noisy BPTI micrograph with SNR = 1/64.

Figure 8: Fourier Shell Correlations (FSCs) for the noisy reconstructions presented in Figure 6 using the 0.5 criterion, compared to the (smoothed) ground truth. (a) FSC for the TRPV1 reconstruction, giving resolution of 24Å; (b) FSC for the BPTI reconstruction, giving resolution of 13Å.

7 Discussion

In the simplified mathematical model above, we showed it is possible to estimate a signal without detecting its appearances. Our strategy is to compute autocorrelations of the micrographs and to relate

Figure 9: 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.

these statistics to the unknown signal's parameters. Recovering the parameters from the statistics reduces to solving a set of polynomial equations.

We also showed how this technique can, in principle, be applied to cryo-EM. Crucially, the outlined approach involves no particle picking, hence a fortiori no viewing direction estimation. Concerns for model bias are also greatly reduced since no template matching is involved. Additionally, our technique also allows the use of much lower defocus values. Lower defocus means lower contrast, but will maintain higher frequency information. Consequently, we may be able to get high resolution reconstructions from fewer micrographs.

Looking toward applying the framework to encompass all important features of real cryo-EM experiments, our work implies that 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. We recognize that 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. 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. On this front, 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 output of this data processing stage is a summary in the form of autocorrelation estimates: its size is a function of the resolution, not a function of the number of observed micrographs. Subsequent steps, which involve solving the system of polynomial equations, scale only in the size of that summary. Of course, an important question then is whether the equations can be solved accurately and efficiently in practice.

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. All of these aspects can be handled with the same general strategy: establish a forward model relating the expected autocorrelations of the micrographs to the target volume(s) and all parameters necessary to model the above effects. For instance, for colored noise, the forward model may involve multiple parameters to capture the power spectrum of the noise instead of the single parameter σ^2 . Similarly, instead of the separation condition, we can model the spacing between the projections using a parameterized pair-correlation function. Such a function models the distribution of distances between neighboring projections. The observed autocorrelations depend linearly on these parameters, which would be estimated as part of the inverse problem. All these aspects must be taken into account so the method can be applied on experimental data. We hope to take care of these issues in future research.

Acknowledgment

The authors thank Ayelet Heimowitz, Joe Kileel, Roy Lederman, Amit Moscovich, Nir Sharon and Fred Sigworth for helpful discussions, and Boris Landa and Yoel Shkolnisky for providing the code for the 2-D PSWFs expansion. The research was partially supported by Award Number R01GM090200 from the NIGMS, FA9550-17-1-0291 from AFOSR, Simons Foundation Math+X Investigator Award, and the Moore Foundation Data-Driven Discovery Investigator Award. NB is partially supported by NSF award DMS-1719558.

References

- [1] Emmanuel Abbe, Tamir Bendory, William Leeb, João Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *arXiv preprint arXiv:1710.02793*, 2017.
- [2] Karim Abed-Meraim, Wanzhi Qiu, and Yingbo Hua. Blind system identification. *Proceedings of the IEEE*, 85(8):1310–1322, 1997.
- [3] GR Ayers and J Christopher Dainty. Iterative blind deconvolution method and its applications. *Optics letters*, 13(7):547–549, 1988.
- [4] Afonso S Bandeira, Ben Blum-Smith, Amelia Perry, Jonathan Weed, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *arXiv preprint arXiv:1712.10163*, 2017.
- [5] Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, and Amit Singer. Bispectrum inversion with application to multireference alignment. *IEEE Transactions on Signal Processing*, 66(4):1037–1050, 2017.
- [6] Albert Benveniste, Maurice Goursat, and Gabriel Ruget. Robust identification of a nonminimum phase system: Blind adjustment of a linear equalizer in data communications. *IEEE Transactions on Automatic Control*, 25(3):385–399, 1980.
- [7] Tejal Bhamre, Teng Zhang, and Amit Singer. Anisotropic twicing for single particle reconstruction using autocorrelation analysis. *arXiv preprint arXiv:1704.07969*, 2017.
- [8] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.

- [9] Nicolas Boumal, Tamir Bendory, Roy R Lederman, and Amit Singer. Heterogeneous multireference alignment: A single pass approach. In *Information Sciences and Systems (CISS), 2018 52nd Annual Conference on*, pages 1–6. IEEE, 2018.
- [10] Olivier Cappé, Arnaud Doucet, Marc Lavielle, and Eric Moulines. Simulation-based methods for blind maximum-likelihood filter identification. *Signal processing*, 73(1-2):3–25, 1999.
- [11] David Donoho, Jiashun Jin, et al. Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics*, 32(3):962–994, 2004.
- [12] Veit Elser. Matrix product constraints by projection methods. *Journal of Global Optimization*, 68(2):329–355, 2017.
- [13] Joachim Frank. Single-particle reconstruction of biological molecules-story in a sample (Nobel lecture). *Angewandte Chemie International Edition*, 2018.
- [14] Georgios B Giannakis and Jerry M Mendel. Identification of nonminimum phase systems using higher order statistics. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(3):360–377, 1989.
- [15] Robert M Glaeser. Electron crystallography: present excitement, a nod to the past, anticipating the future. *Journal of structural biology*, 128(1):3–14, 1999.
- [16] Sandeep Gogineni, Pawan Setlur, Muralidhar Rangaswamy, and Raj Rao Nadakuditi. Passive radar detection with noisy reference channel using principal subspace similarity. *IEEE Transactions on Aerospace and Electronic Systems*, 54(1):18–36, 2018.
- [17] MHMH Hayes. The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(2):140–154, 1982.
- [18] 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.
- [19] 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.
- [20] Stuart M Jefferies and Julian C Christou. Restoration of astronomical images by iterative blind deconvolution. *The Astrophysical Journal*, 415:862, 1993.
- [21] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [22] 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*, 8:16099, 2017.
- [23] John Kormylo and J Mendel. Identifiability of nonminimum phase linear stochastic systems. *IEEE transactions on automatic control*, 28(12):1081–1090, 1983.
- [24] Boris Landa and Yoel Shkolnisky. Steerable principal components for space-frequency localized images. *SIAM journal on imaging sciences*, 10(2):508–534, 2017.

- [25] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. In *Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on*, pages 1569–1573. IEEE, 2018.
- [26] Michael S Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9(4):R53–R78, 1998.
- [27] 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, 2017.
- [28] KS Lii, M Rosenblatt, et al. Deconvolution and estimation of transfer function phase and coefficients for nongaussian linear processes. *The annals of statistics*, 10(4):1195–1208, 1982.
- [29] 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.
- [30] Lennart Ljung. System identification. In *Signal analysis and prediction*, pages 163–173. Springer, 1998.
- [31] Frank Natterer. *The mathematics of computerized tomography*, volume 32. SIAM, 1986.
- [32] 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.
- [33] Lawrence R Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [34] Giovanna Scapin, Clinton S Potter, and Bridget Carragher. Cryo-EM for small molecules discovery, design, understanding, and application. *Cell Chemical Biology*, 2018.
- [35] Sjors HW Scheres. RELION: implementation of a bayesian approach to cryo-EM structure determination. *Journal of structural biology*, 180(3):519–530, 2012.
- [36] Ofir Shalvi and Ehud Weinstein. New criteria for blind deconvolution of nonminimum phase systems (channels). *IEEE Transactions on information theory*, 36(2):312–321, 1990.
- [37] 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.
- [38] FJ Sigworth. A maximum-likelihood approach to single-particle image refinement. *Journal of structural biology*, 122(3):328–339, 1998.
- [39] David Slepian. Prolate spheroidal wave functions, Fourier analysis and uncertainty—iv: Extensions to many dimensions; generalized prolate spheroidal functions. *The Bell System Technical Journal*, 43(6):3009–3057, Nov 1964.
- [40] 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.

- [41] Jitendra Tugnait. Identification of nonminimum phase linear stochastic systems. In *The 23rd IEEE Conference on Decision and Control*, number 23, pages 342–347, 1984.
- [42] Marin van Heel. Finding trimeric HIV-1 envelope glycoproteins in random noise. *Proceedings of the National Academy of Sciences*, 110(45):E4175–E4177, 2013.
- [43] Marin van Heel, Michael Schatz, and Elena Orlova. Correlation functions revisited. *Ultramicroscopy*, 46(1–4):307–316, 1992.
- [44] Yuqian Zhang, Han-Wen Kuo, and John Wright. Structured local optima in sparse blind deconvolution. *arXiv preprint arXiv:1806.00338*, 2018.
- [45] Zhizhen Zhao, Yoel Shkolnisky, and Amit Singer. Fast steerable principal component analysis. *IEEE transactions on computational imaging*, 2(1):1–12, 2016.

A Numerical experiments details

We run the toy model experiments on a shared computer with 144 logical CPUs of type Intel(R) Xeon(R) CPU E7-8880 v3 @ 2.30GHz and 792 Gb of RAM; we use at most 72 of these CPUs.

2-D experiment. For the 2-D experiment shown in Figures 2 and 3, we generated micrographs of size 4096×4096 pixels. In each micrograph, we placed Einstein’s image (of zero mean) of size 50×50 in random locations, while preserving the separation condition (2). This is done by randomly selecting 4000 placements in the micrograph, one at a time with an accept/reject rule based on the separation condition and locations picked so far. On average, 700 images are placed in each micrograph. Then, i.i.d. Gaussian noise with standard deviation $\sigma = 3$ is added, inducing an SNR of approximately 1/370. An example of a micrograph’s excerpt is presented in the right panel of Figure 1.

In this experiment, we assume to know the noise level σ and the total number of occurrences of the target image across all micrographs. In stark contrast with the 1-D setup, the second-order autocorrelation determines almost any target image uniquely, up to reflection through the origin. This is because the second-order autocorrelations correspond to the Fourier magnitudes of the signal through the 2-D Fourier transform. Therefore, we estimate the signal’s Fourier magnitudes (or power spectrum) from the Fourier magnitudes of the micrographs, at the cost of one 2-D fast Fourier transform (FFT) per micrograph. These can be computed highly efficiently and in parallel. Computing the power spectrum of 200 micrographs took less than 23 seconds.

To recover the target image from the estimated power spectrum, we use a standard phase retrieval algorithm called relaxed-reflect-reflect (RRR). This algorithm iterates the map

$$z \leftarrow z + \beta(P_2(2P_1(z) - z) - P_1(z)),$$

on an image z of size $2L \times 2L$. We set the parameter β to 1. The map is designed so that, if the estimated power spectrum is exact, then fixed points contain Einstein’s image in the upper-left corner of size $L \times L$, possibly reflected through its origin, and zeros elsewhere. The operator $P_2(z)$ combines the Fourier phases of the current estimation z with the estimated Fourier magnitudes. The operator $P_1(z)$ zeros out all entries of z outside the $L \times L$ upper-left corner. In all experiments, the algorithm halted after a fixed number of 2000 iterations. Running the RRR took about 3 seconds.

The computational complexity scales as $O(N^2L^2 + F(L))$, where F is the complexity of the RRR algorithm and N^2 stands for the total number of pixels across micrographs.

1-D experiment. For the 1-D experiment depicted in Figure 4, we used a signal of length $L = 21$ and generated an observation y of length $N = 10(2L - 1)M$, where M is the number of repetitions. The observation was generated by randomly selecting placements in y , one at a time with an accept/reject rule based on the separation condition and locations picked so far until reaching the desired number of signal occurrences. Then, i.i.d. Gaussian noise with mean zero and standard deviation $\sigma = 3$ is added, to form the observed y . The resulting SNR of y is about $1/175$. Figure 5 shows an example of the observed data.

Given the observation y , we proceed to compute the autocorrelations. We excluded biased elements of the autocorrelations; that is $a_y^2[\ell_1]$ and $a_y^3[\ell_1, \ell_2]$ such that ℓ_1, ℓ_2 or $\ell_1 - \ell_2$ are zero. This has the practical effect that we need not know σ . By also taking symmetries into account, the third-order autocorrelation contains $\frac{(L-1)(L-2)}{2}$ remaining entries. Thus, in total we have

$$1 + (L - 1) + \frac{(L - 1)(L - 2)}{2} = \frac{1}{2}L(L - 1) + 1$$

coefficients. This redundancy suggests that it might be possible to estimate several signals simultaneously (compare with [9, 4]). In practice, the autocorrelations are computed on disjoint segments of y of length $100 \cdot 10^6$ (if the length of the measurement is larger than that) and added up, without correction for the junction points.

Having computed the moments of interest, we now estimate the signal x and coefficient γ which agree with the data. We choose to do so by running an optimization algorithm on the following nonlinear LS problem:

$$\min_{\substack{\hat{x} \in \mathbb{R}^W \\ \hat{\gamma} > 0}} w_1 (a_y^1 - \hat{\gamma} a_{\hat{x}}^1)^2 + w_2 \sum_{\ell=1}^{L-1} (a_y^2[\ell] - \hat{\gamma} a_{\hat{x}}^2[\ell])^2 + w_3 \sum_{\substack{2 \leq \ell_1 \leq L-1 \\ 1 \leq \ell_2 \leq \ell_1-1}} (a_y^3[\ell_1, \ell_2] - \hat{\gamma} a_{\hat{x}}^3[\ell_1, \ell_2])^2, \quad (12)$$

where $W \geq L$ is the length of the sought signal and the weights are set to $w_1 = 1/2, w_2 = 1/2n_2, w_3 = 1/2n_3$, where n_2, n_3 are the number of moments used: $n_2 = L - 1, n_3 = \frac{(L-1)(L-2)}{2}$ (weights could also be set in accordance with variance estimates as in [9]).

Setting $W = L$ (as is a priori desired) is problematic because the above optimization problems appears to have numerous poor local optimizers. A similar phenomenon was recently observed in a related problem [44]. Thus, we first run the optimization with $W = 2L - 1$. This problem appears to have few poor local optima, perhaps because the additional degrees of freedom allow for more escape directions. Since we hope the signals estimated this way correspond to the true signals zero-padded to length W , we extract from each one a subsignal of length L that has largest ℓ_2 -norm. This estimator is then used as initial iterate for (12), this time with $W = L$. We find that this procedure is reliable for a wide range of experimental parameters. To solve (12), we run the trust-region method implemented in Manopt [8], which allows to treat the positivity constraints on coefficients $\hat{\gamma}$. Notice that the cost function is a polynomial in the variables, so that it is straightforward to compute it and its derivatives.

The computational complexity scales as $O(NL^2 + F(L))$, where F is the complexity of the solving the LS. For example, computing the moments for the experiment in the left panel of Figure 4 took 38 seconds, while the optimization took 5.5 seconds in total.

3-D experiments. The experiments in Figure 6, and the recoveries of γ for the experiment in Figure 9, 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 in Figure 9 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 to compute the first three autocorrelations.

For the clean experiments in Figure 6, we computed the analytical first three autocorrelations of the volume as explained in Appendix F. To estimate the coefficients of the volume itself, we solve the LS problem

$$\min_{\hat{x}, \hat{\gamma}} w_1 \|\mathbf{a}_x^1 - \hat{\gamma} \mathbf{a}_{\hat{x}}^1\|^2 + w_2 \|\mathbf{a}_x^2 - \hat{\gamma} \mathbf{a}_{\hat{x}}^2\|_2^2 + w_3 \|\mathbf{a}_x^3 - \hat{\gamma} \mathbf{a}_{\hat{x}}^3\|_F^2, \quad (13)$$

where the explicit expressions of $\mathbf{a}_x^1$, $\mathbf{a}_x^2$ and $\mathbf{a}_x^3$ are given in (46), (45) and (44), respectively. In the experiments, we set $w_1 = w_2 = w_3 = 1$. The LS problem was solved using Matlab's `lsqnonlin` solver for nonlinear LS problems. The expected $\hat{\gamma}$ in this experiment is one.

The computational complexity for the computation of the moments from the micrograph is $O(N^2 P^6)$, since we have N^2 patches that we extract from the micrograph. For each patch we perform an expansion in PSWFs with complexity $O(P^3)$ [24], and then compute the autocorrelations dominated by computation of the bispectrum as in (38) costing $O(P^6)$. The complexity of evaluating the bispectrum using (44) after the precomputation takes $O(BP^6)$, where B is the number of entries in the bispectrum. Since $B = O(P^3)$, the total complexity for evaluating the bispectrum 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 $BV^2 = O(P^9)$. Thus, the complexity of solving the LS problem (13) is $O(KP^9)$ where K is the number of iterations required for the optimizer to converge, since the bispectrum evaluation dominates the cost of each iteration.

The true volume used in the experiments in Figures 9 and 6b 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 1QLQ². We generated an EM map from this atomic model in UCSF Chimera [32] 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 6b, the volume was downsampled to size 20^3 . For the experiment in Figure 6a, we used the TRPV1 in complex with DkTx and RTX, whose EM map is available in the Electron Microscopy Data Bank (EMDB) as EMD-8117³. The original map 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 (25) with cutoff $L_{\max} = 5$ for TRPV1 and $L_{\max} = 2$ for BPTI. For the TRPV1 reconstruction, the optimizer converged to a point giving 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.

The micrographs for the experiments presented in Figures 9 and 7 were generated as follows. We sample rotation matrices from $SO(3)$ uniformly at random using the QR-based algorithm described in [40], and generate the projection of the volume corresponding to that rotation matrix using `cryo_project` in ASPIRE⁴. The projections for the experiments in Figure 7 were obtained from the smoothed volumes, not the original ones. 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 $L - 1$ pixels in each dimension, where the projections are contained in a box of size $L \times L$. 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. In the experiments, we define SNR slightly differently than in the toy model, namely, 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 the noisy micrograph in Figure 9 was $\text{SNR} = 1/1024$, the

²<https://www.rcsb.org>

³<http://www.ebi.ac.uk/pdbe/emdb>

⁴<http://spr.math.princeton.edu>

noise level in Figure 7 was $\text{SNR} = 1/16$ for the TRPV1 micrographs and $\text{SNR} = 1/64$ for the BPTI micrographs. We used 300 micrographs for each of the noisy reconstructions presented in Figure 6.

For the noisy reconstructions presented in Figure 6, the LS problem (13) was solved assuming spherical harmonic cutoffs of $L_{\max} = 5$ for TRPV1 and $L_{\max} = 2$ for BPTI, same as for the ground truth. For the experiment in Figure 9, the LS problem (13) 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.

B Proof of Proposition 3.1

The proof is based on a variant of the Neyman–Pearson Lemma to derive the best (deterministic) estimator $\hat{\eta}$. Take any estimator $\hat{\eta}$; it is characterized by S : the set of X 's where $\hat{\eta} = 1$, where X is a random vector as defined above Proposition 3.1. 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:

$$\begin{aligned} \text{Prob}[\hat{\eta} \neq \eta] &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q)\text{Prob}_1[\hat{\eta} = 0] \\ &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q)(1 - \text{Prob}_1[\hat{\eta} = 1]) \\ &= q\text{Prob}_0[\hat{\eta} = 1] + (1 - q) - (1 - q)\text{Prob}_1[\hat{\eta} = 1] \\ &= (1 - q) + \int_S (qf_0(x) - (1 - q)f_1(x))dx, \end{aligned} \tag{14}$$

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 (14) 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 $\max(q, 1 - q)$.

C Derivation of the identities in Section 4

We consider the asymptotic regime $N, M \rightarrow \infty$ and assume that $M = \Omega(N)$, so that

$$\gamma := \frac{ML}{N} > 0.$$

Any vectors with indices out of range are given value 0.

We start by considering the mean of the data:

$$a_y^1 = \frac{1}{N} \sum_{i=0}^{N-1} y[i] = \frac{1}{N/L} \sum_{j=0}^{M-1} \frac{1}{L} \sum_{i=0}^{L-1} x[i] + \underbrace{\frac{1}{N} \sum_{i=0}^{N-1} \varepsilon[i]}_{\text{noise term}} \xrightarrow{a.s.} \gamma a_x^1,$$

where the noise term converges to zero almost surely (a.s.) by the strong law of large numbers.

We proceed with the (second-order) autocorrelation for fixed $\ell \in [0, \dots, L-1]$. We can compute:

$$\begin{aligned} a_y^2[\ell] &= \frac{1}{N} \sum_{i=0}^{N-1-\ell} y[i]y[i+\ell] \\ &= \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-\ell-1} x[i]x[i+\ell]}_{\text{signal term}} + \underbrace{\frac{1}{N} \sum_{i=0}^{N-1-\ell} \varepsilon[i]\varepsilon[i+\ell]}_{\text{noise term}} \\ &\quad + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-1} x[i](\varepsilon[s_j+i+\ell] + \varepsilon[s_j+i-\ell])}_{\text{cross-term}}. \end{aligned}$$

The cross-term is linear in the noise, and is easily shown to vanish almost surely in the limit $N \rightarrow \infty$, by the strong law of large numbers. We break the signal term into M different sums, each containing one copy of the signal. This gives:

$$\begin{aligned} \frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-\ell-1} x[i]x[i+\ell] &= \frac{ML}{N} \frac{1}{L} \sum_{i=0}^{L-\ell-1} x[i]x[i+\ell] \\ &\xrightarrow{N \rightarrow \infty} \gamma a_x^2[\ell]. \end{aligned} \quad (15)$$

We next analyze the pure noise term. When $\ell \neq 0$, we can break the noise term into a sum of independent terms:

$$\frac{1}{N} \sum_{i=0}^{N-1-\ell} \varepsilon[i]\varepsilon[i+\ell] = \frac{1}{\ell} \sum_{i=0}^{\ell-1} \frac{1}{N/\ell} \sum_{j=0}^{N/\ell-2} \varepsilon[j\ell+i]\varepsilon[(j+1)\ell+i].$$

Each sum $\frac{1}{N/\ell} \sum_{j=0}^{N/\ell-2} \varepsilon[j\ell+i]\varepsilon[(j+1)\ell+i]$ is an average of N/ℓ independent terms with expectation zero, hence converges to zero almost surely as $N \rightarrow \infty$. If $\ell = 0$, then we have:

$$\frac{1}{N} \sum_{i=0}^{N-1} \varepsilon^2[i] \xrightarrow{a.s.} \sigma^2.$$

We now analyze the third-order autocorrelation. Let us fix $\ell_1 \geq \ell_2 \geq 0$. We have:

$$\begin{aligned} a_y^3[\ell_1, \ell_2] &= \frac{1}{N} \sum_{i=0}^{N-1-\ell_1} y[i]y[i+\ell_1]y[i+\ell_2] \\ &= \underbrace{\frac{ML}{N} \frac{1}{M} \sum_{j=1}^M \frac{1}{L} \sum_{i=0}^{L-1-\ell_1} x[i]x[i+\ell_1]x[i+\ell_2]}_{(1)} + \underbrace{\frac{1}{N} \sum_{i=0}^{N-1-\ell_1} \varepsilon[i]\varepsilon[i+\ell_1]\varepsilon[i+\ell_2]}_{(2)} \\ &\quad + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-1} x[i]\varepsilon[s_j+i+\ell_1]\varepsilon[s_j+i+\ell_2]}_{(3)} + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-1} \varepsilon[s_j+i-\ell_1]x[i]\varepsilon[s_j+i+\ell_2-\ell_1]}_{(4)} \\ &\quad + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-1} \varepsilon[s_j+i-\ell_2]\varepsilon[s_j+i+\ell_1-\ell_2]x[i]}_{(5)} + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-\ell_1+\ell_2-1} \varepsilon[s_j+i-\ell_2]x[i+\ell_1-\ell_2]x[i]}_{(6)} \end{aligned}$$

$$+ \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-\ell_2-1} x[i] \varepsilon[s_j + i + \ell_1] x[i + \ell_2]}_{(7)} + \underbrace{\frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-\ell_1-1} x[i] x[i + \ell_1] \varepsilon[s_j + i + \ell_2]}_{(8)}.$$

Terms (6), (7) and (8) are linear in ε , and can easily be shown to converge to 0 almost surely by the law of large numbers, by similar arguments as used previously. Term (1) converges to $\gamma a_x^3[\ell_1, \ell_2]$ almost surely, for the same reasons as (15). To deal with terms (2)–(5), we must distinguish between different values of ℓ_1 and ℓ_2 .

Case 1: $0 < \ell_2 < \ell_1$. Here, all summands with elements of ε involve products of distinct entries, which have expected value 0. Consequently, the usual argument shows that terms (2)–(5) all converge to 0 almost surely as $N \rightarrow \infty$.

Case 2: $0 = \ell_2 < \ell_1$. Term (2) is an average of products of the form $\varepsilon[i]^2 \varepsilon[i + \ell_1]$, which have mean zero; consequently, term (2) converges to 0 almost surely. The same argument as for Case 1 shows that (3) and (5) also converge to 0. For term (4), we write:

$$\begin{aligned} & \frac{1}{N} \sum_{j=1}^M \sum_{i=0}^{L-1} \varepsilon[s_j + i - \ell_1] x[i] \varepsilon[s_j + i + \ell_2 - \ell_1] \\ &= \frac{ML}{N} \frac{1}{L} \sum_{i=0}^{L-1} x[i] \frac{1}{M} \sum_{j=1}^M \varepsilon[s_j + i - \ell_1]^2 \\ &\xrightarrow{a.s.} \gamma \frac{1}{L} \sum_{i=0}^{L-1} x[i]^2 \sigma^2 = \gamma a_x^1 \sigma^2. \end{aligned}$$

Case 3: $0 < \ell_2 = \ell_1$. An argument nearly identical to that for Case 2 shows that terms (2), (4) and (5) converge to 0, while term (3) converges to $\gamma a_x^1 \sigma^2$.

Case 4: $0 = \ell_2 = \ell_1$. The same argument as for term (4) in Case 2 shows that terms (3), (4) and (5) all converge to $\gamma a_x^1 \sigma^2$. Term (2) is an average of $\varepsilon[i]^3$, which is mean zero; consequently, it converges to 0. This completes the proof.

D Proof of Proposition 5.2

In the limit,

$$(a_y^1)^2 = \frac{\gamma^2}{L^2} \sum_{i=0}^{L-1} \sum_{j=0}^{L-1} x[i] x[j].$$

Similarly,

$$\sum_{\ell=1}^{L-1} a_y^2[\ell] = \frac{\gamma}{L} \sum_{\ell=1}^{L-1} \sum_{i=0}^{L-1-\ell} x[i] x[i + \ell],$$

and $a_y^2[0] = \frac{\gamma}{L} \sum_{i=0}^{L-1} x^2[i] + \sigma^2$. The proof is concluded by noting that $a_x^2[-\ell] = a_x^2[\ell]$.

E Proof of Proposition 5.3

We prove that both σ and γ are identifiable from the observed first three moments of y . For convenience, we work with $\beta = \gamma/L$ rather than γ itself. To this end, we construct two quadratic equations satisfied by β and whose coefficients can be computed from observable quantities (in the limit). Then, we show

that these equations are independent, and hence that β is uniquely defined. Given β , we can estimate σ using Proposition 5.2.

Throughout the proof, it is important to distinguish between observed and unobserved values. We denote the observed values by E_i or a_y^1, a_y^2, a_y^3 . We use F_i to denote functions of the signal's autocorrelations (which are not directly observable).

In the limit $N \rightarrow \infty$, almost surely, $a_y^1 = \beta(\mathbf{1}^T x)$ and $a_y^2[0] = \beta\|x\|^2 + \sigma^2$, where $\mathbf{1} \in \mathbb{R}^L$ is the vector of all-ones. (In this whole section, for clarity, we now omit to specify that identities hold almost surely in the limit.) Consider the product:

$$E_1 := a_y^1 a_y^2[0] = (\beta(\mathbf{1}^T x))(\beta\|x\|^2 + \sigma^2) = \sigma^2 a_y^1 + L\beta^2 F_1, \quad (16)$$

where $F_1 := a_x^3[0, 0] + \sum_{j=1}^{L-1} (a_x^3[j, j] + a_x^3[0, j])$. The terms of F_1 can also be estimated from a_y^3 , while taking the scaling and bias terms into account. This yields another observable:

$$\begin{aligned} E_2 &:= a_y^3[0, 0] + \sum_{j=1}^{L-1} (a_y^3[j, j] + a_y^3[0, j]) \\ &= L\beta F_1 + (2L + 1)\sigma^2 a_y^1. \end{aligned} \quad (17)$$

Therefore, from (16) and (17) we get:

$$E_2 \beta - (2L + 1)\sigma^2 \beta a_y^1 = E_1 - \sigma^2 a_y^1. \quad (18)$$

Let $E_3 := a_y^2[0] + 2 \sum_{j=1}^{L-1} a_y^2[j]$; recall from Proposition 5.2:

$$\sigma^2 = E_3 - (a_y^1)^2 / \beta. \quad (19)$$

Plugging into (18) and rearranging, we get a first quadratic equation in β ,

$$\mathcal{A}\beta^2 + \mathcal{B}\beta + \mathcal{C} = 0, \quad (20)$$

where

$$\begin{aligned} \mathcal{A} &= E_2 - (2L + 1)a_y^1 E_3, \\ \mathcal{B} &= -E_1 + (2L + 1)(a_y^1)^3 + a_y^1 E_3, \\ \mathcal{C} &= -(a_y^1)^3. \end{aligned}$$

Importantly, these coefficients are observable quantities. As we assume throughout this proof that x has nonzero mean, $a_y^1 \neq 0$ and we conclude that this equation is non-trivial.

Next, we derive the second quadratic equation for β . We notice that

$$E_4 := \frac{1}{L}(a_y^1)^3 = \frac{1}{L}\beta^3(\mathbf{1}^T x)^3 = \beta^3 F_2, \quad (21)$$

where $F_2 = \frac{1}{L}(\mathbf{1}^T x)^3$, and we can work out that:

$$F_2 = a_x^3[0, 0] + 3 \sum_{j=1}^{L-1} (a_x^3[j, j] + a_x^3[0, j]) + 6 \sum_{1 \leq i < j \leq L-1} a_x^3[i, j].$$

Once again, F_2 can be estimated from a_y^3 , taking bias and scaling into account:

$$E_5 := a_y^3[0, 0] + 3 \sum_{j=1}^{L-1} (a_y^3[j, j] + a_y^3[0, j]) + 6 \sum_{1 \leq i < j \leq L-1} a_y^3[i, j] = L\beta F_2 + (6L - 3)\sigma^2 a_y^1. \quad (22)$$

Consider the following ratio:

$$\frac{E_5}{E_4} = \frac{L}{\beta^2} + \frac{(6L-3)\sigma^2 a_y^1}{E_4}.$$

From the latter, we deduce:

$$\sigma^2 = \frac{E_5}{a_y^1(6L-3)} - \frac{LE_4}{\beta^2 a_y^1(6L-3)}.$$

Using (19) and rearranging, we get the second quadratic:

$$\mathcal{D}\beta^2 + \mathcal{E}\beta + \mathcal{F} = 0, \quad (23)$$

where

$$\begin{aligned} \mathcal{D} &= E_3 - \frac{E_5}{a_y^1(6L-3)}, \\ \mathcal{E} &= -(a_y^1)^2, \\ \mathcal{F} &= \frac{LE_4}{a_y^1(6L-3)}. \end{aligned}$$

It is also non-trivial since $E_4 \neq 0$.

To complete the proof, we need to show that the two quadratic equations (20) and (23) are independent. To this end, it is enough to show that the ratios between coefficients differ. From (20) and (16), we have:

$$\frac{\mathcal{B}}{\mathcal{C}} = \frac{E_1 - (2L+1)(a_y^1)^3 - a_y^1 E_3}{(a_y^1)^3} = \frac{a_y^2[0] - (2L+1)(a_y^1)^2 - E_3}{(a_y^1)^2}.$$

In addition, using (21),

$$\frac{\mathcal{E}}{\mathcal{F}} = \frac{(3-6L)(a_y^1)^3}{LE_4} = 3-6L.$$

For contradiction, suppose that the quadratics are dependent. Then, $\frac{\mathcal{B}}{\mathcal{C}} = \frac{\mathcal{E}}{\mathcal{F}}$, that is,

$$a_y^2[0] - (2L+1)(a_y^1)^2 - E_3 = (a_y^1)^2(3-6L).$$

Rewriting the identity in terms of x and dividing by β we get:

$$4(L-1)\beta(\mathbf{1}^\top x)^2 - (\mathbf{1}^\top x)^2 + \|x\|^2 = 0. \quad (24)$$

For generic x , this polynomial equation is not satisfied so that the quadratic equations are independent. Furthermore, from the inequality $L\|x\|^2 \geq (\mathbf{1}^\top x)^2$ it follows immediately that the equations must be independent so long as

$$\beta > \frac{1}{4L}.$$

F Autocorrelations for the cryo-EM problem

F.1 Model and autocorrelation functions

Let ϕ be the Coulomb potential representing the molecule we aim to recover. We assume that molecule is real-valued and smooth. In spherical coordinates, its 3-D Fourier transform $\hat{\phi}$ admits a finite expansion of the form

$$\hat{\phi}(ck, \theta, \varphi) = \sum_{\ell=0}^{L_{\max}} \sum_{m=-\ell}^{\ell} \sum_{s=1}^{S(\ell)} x_{\ell,m,s} Y_{\ell}^m(\theta, \varphi) j_{\ell,s}(k), \quad k \leq 1, \quad (25)$$

where c is the bandlimit, $\{S(\ell)\}$ are determined using the Nyquist criterion as described in [7], $j_{\ell,s}$ 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. Sampling at the Nyquist rate dictates $c = 1/2$ [25]. Because ϕ is real-valued, $\hat{\phi}$ 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 $m \geq 0$.

Let I_{ω} denote the tomographic projection obtained from viewing direction $\omega \in SO(3)$. By the Fourier projection-slice theorem, its 2-D Fourier transform is given by [31]:

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

where $D_{m', m}^{\ell}(\omega)$ is a Wigner-D matrix. This implies that the projections are also c -bandlimited.

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:

$$\mathcal{I} = \sum_{t=1}^M I_{\omega_t} * \delta_{\mathbf{s}_t} + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I), \quad (26)$$

where the viewing directions ω_t are assumed to be drawn from the uniform distribution over $SO(3)$ and $\mathbf{s}_t$ denotes the location of the center of the t th projection in the micrograph. We assume the projection is discretized on a Cartesian grid of size $P \times P$ and 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, similarly to (2). Note that (1) can be also written as a sum of δ functions as in (26).

Define the p th autocorrelation of $\mathcal{I}$ as

$$a_{\mathcal{I}}^p[\ell_1, \dots, \ell_{p-1}] := \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\ell] \mathcal{I}[\mathbf{i} + \ell_1] \cdots \mathcal{I}[\mathbf{i} + \ell_{p-1}],$$

where the summation is for $\mathbf{i}$ ranging over the N^2 pixels of the micrograph. 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

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{i=1}^K a_{\mathcal{I}_i}^p[\ell_1, \dots, \ell_{p-1}] = \gamma \langle a_{I_{\omega}}^p[\ell_1, \dots, \ell_{p-1}] \rangle_{\omega} + b_p[\ell_1, \dots, \ell_{p-1}], \quad (27)$$

$$p = 1, 2, 3, \quad \ell_1, \dots, \ell_{p-1} \in [-(P-1), P-1]^2,$$

where $\langle \cdot \rangle_{\omega}$ denotes averaging over all possible viewing directions ω 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 $\gamma\langle a_{I_\omega}^p \rangle_\omega$ directly from the micrographs. These quantities are functions of the unknown coefficients $x_{l,m,s}$ and we could proceed to invert their relation, as we did in the toy examples.

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 (27) to average over all in-plane rotations as well. This can be done efficiently using Prolate Spheroidal Wave Functions (PSWFs). We use autocorrelations up to and including the third order. Indeed, second-order autocorrelations are not enough, as was observed already in [21] for a simpler problem where the input is not micrographs but rather picked, perfectly centered particles.

F.2 Autocorrelation derivation

In this section we prove relation (27). We note that mathematically taking infinitely many micrographs is equivalent to take 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 (26). Denote by $\mathcal{S}_t$ the support of the t th particle. Then, we have

$$\begin{aligned} a_{\tilde{\mathcal{I}}}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] &= \frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\ &= \frac{1}{N^2} \sum_{t=1}^M \sum_{\mathbf{i} \in \mathcal{S}_t} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots \tilde{\mathcal{I}}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\ &= \frac{MP^2}{N^2} \cdot \frac{1}{M} \sum_{t=1}^M \frac{1}{P^2} \sum_{i,j=0}^{P-1} I_{\omega_t}[\mathbf{i}] I_{\omega_t}[\mathbf{i} + \boldsymbol{\ell}_1] \cdots I_{\omega_t}[\mathbf{i} + \boldsymbol{\ell}_{p-1}] \\ &= \frac{MP^2}{N^2} \frac{1}{M} \sum_{t=1}^M a_{I_{\omega_t}}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] \\ &\rightarrow \gamma \langle a_{I_\omega}^p[\boldsymbol{\ell}_1, \dots, \boldsymbol{\ell}_{p-1}] \rangle_\omega, \end{aligned} \tag{28}$$

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 (27). 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 (28) for $p = 2$. The cross terms vanish in the limit. The fourth term is zero unless $\boldsymbol{\ell} = 0$, in which case it converges to σ^2 . Thus, we conclude

$$a_{\mathcal{I}}^2[\boldsymbol{\ell}] \rightarrow \gamma \langle a_{I_\omega}^2[\boldsymbol{\ell}] \rangle_\omega + \sigma^2 \delta[\boldsymbol{\ell}], \tag{29}$$

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

For the third moments, we get 8 terms:

$$\begin{aligned}
 a_{\mathcal{I}}^3[\ell_1, \ell_2] = & \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \ell_1] \tilde{\mathcal{I}}[\mathbf{i} + \ell_2]}_{(1)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \varepsilon[\mathbf{i} + \ell_1] \varepsilon[\mathbf{i} + \ell_2]}_{(2)} \\
 & + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \ell_1] \tilde{\mathcal{I}}[\mathbf{i} + \ell_2]}_{(3)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \ell_1] \varepsilon[\mathbf{i} + \ell_2]}_{(4)} \\
 & + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \ell_1] \tilde{\mathcal{I}}[\mathbf{i} + \ell_2]}_{(5)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \tilde{\mathcal{I}}[\mathbf{i}] \varepsilon[\mathbf{i} + \ell_1] \varepsilon[\mathbf{i} + \ell_2]}_{(6)} \\
 & + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \varepsilon[\mathbf{i} + \ell_1] \tilde{\mathcal{I}}[\mathbf{i} + \ell_2]}_{(7)} + \underbrace{\frac{1}{N^2} \sum_{\mathbf{i}} \varepsilon[\mathbf{i}] \tilde{\mathcal{I}}[\mathbf{i} + \ell_1] \varepsilon[\mathbf{i} + \ell_2]}_{(8)}.
 \end{aligned}$$

We address these terms one by one:

- Term (1) is treated by (28) for $p = 3$;
- Term (2) is the third-order autocorrelation of pure noise which vanishes in the limit;
- Terms (3)-(5) depend linearly on the noise and hence vanish in the limit;
- For term (6), if $\ell_1 \neq \ell_2$ the term vanishes in the limit. If $\ell_1 = \ell_2$ then

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

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

$$a_{\mathcal{I}}^3[\ell_1, \ell_2] \rightarrow \gamma \langle a_{I_{\omega}}^3[\ell_1, \ell_2] \rangle_{\omega} + \gamma \sigma^2 \langle a_{I_{\omega}}^1 \rangle_{\omega} \left(\delta[\ell_1 - \ell_2] + \delta[\ell_1] + \delta[\ell_2] \right), \quad (30)$$

where the second term is the bias $b_3[\ell_1, \ell_2]$. Note that $\gamma \langle a_{I_{\omega}}^1 \rangle_{\omega}$ is approximately the mean of the micrograph since $a_{\mathcal{I}}^1 \approx a_{\mathcal{I}}^1 \approx \gamma \langle a_{I_{\omega}}^1 \rangle_{\omega}$. Therefore, we do not need prior knowledge of γ to effectively debias the third-order autocorrelation.

F.3 Accounting for all in-plane rotations

We represent our autocorrelations using Prolate Spheroidal Wave Functions (PSWFs) $\{\psi_{k,q}\}$ where $k \geq 0, q \geq 1$ are integers [39]. 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. This is only a concern for the

second and third order autocorrelations. Below, we start with $p = 2$. The PSWFs are eigenfunctions of the truncated Fourier transform and are given in polar coordinates by⁵

$$\psi_{k,q}(r, \varphi) = \begin{cases} \frac{1}{\sqrt{8\pi}} \alpha_{k,q} R_{k,q}(r) e^{ik\varphi}, & r \leq 1, \\ 0, & r > 1, \end{cases} \quad (31)$$

where the range of k, q is determined by Eq. (8) in [24], the $R_{k,q}$ are a family of real, one-dimensional functions and the $\alpha_{k,q}$ are scaling factors which will be defined in the next section. The PSWFs are orthogonal on the unit disk.

For $\ell \in [-(P-1), P-1]^2$, let us define

$$\mathfrak{a}^2[k, q] = \sum_{\ell} a_{\mathcal{I}}^2[\ell] \overline{\psi_{k,q}[\ell]}, \quad (32)$$

where $\psi_{k,q}[\ell] := \psi_{k,q}(\ell/(P-1))$ 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:

$$a_{\mathcal{I}}^2[\ell] \approx \sum_{k,q} \mathfrak{a}^2[k, q] \psi_{k,q}[\ell]. \quad (33)$$

This holds because the continuous PSWFs form an orthogonal basis, and their discretized counterparts are (empirically) almost orthogonal. 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 $\mathfrak{a}^2[k, q]$ can be computed from the micrographs directly. By definition,

$$\begin{aligned} \mathfrak{a}^2[k, q] &= \sum_{\ell} a_{\mathcal{I}}^2[\ell] \overline{\psi_{k,q}[\ell]} \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \left(\sum_{\ell} \mathcal{I}[\mathbf{i} + \ell] \overline{\psi_{k,q}[\ell]} \right) \\ &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] a_{k,q}[\mathbf{i}], \end{aligned} \quad (34)$$

where we defined

$$a_{k,q}[\mathbf{i}] = \sum_{\ell} \mathcal{I}[\mathbf{i} + \ell] \overline{\psi_{k,q}[\ell]}. \quad (35)$$

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 (33), a direct expansion of that patch in the discretized PSWFs yields the sought coefficients:

$$\mathcal{I}[\mathbf{i} + \ell] \approx \sum_{k,q} a_{k,q}[\mathbf{i}] \psi_{k,q}[\ell]. \quad (36)$$

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 (36), and collect the $a_{k,q}$ as per (34) to constitute the second-order autocorrelation of the micrograph.

⁵A different normalization is used in [24].

Crucially, following this formalism, it is now straightforward to account for all in-plane rotations and reflections of the micrograph. Indeed, as can be seen from the definition of the (continuous) PSWFs (31), the effects of rotations and reflections on expansion coefficients of real images are, respectively, phase modulation and conjugation: this is why this the PSWF basis is called *steerable* [24, 45]. By analogy in the discrete case, we have the following approximate expansions for a patch rotated about its center $\mathbf{i}$ by an angle α :

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

and the reflection followed by a rotation by angle α :

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

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

$$\begin{aligned} \mathbf{a}^2[k, q] &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \left(\frac{1}{4\pi} \int_0^{2\pi} (a_{k,q}[\mathbf{i}] + \overline{a_{k,q}[\mathbf{i}]}) e^{-\iota k \alpha} d\alpha \right) \\ &= \delta[k] \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] a_{0,q}[\mathbf{i}], \end{aligned} \quad (37)$$

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 (35). 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

$$\begin{aligned} \mathfrak{a}^3[k_1, q_1; k_2, q_2] &= \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \frac{1}{4\pi} \int_0^{2\pi} \left(a_{k_1, q_1}[\mathbf{i}] \overline{a_{k_2, q_2}[\mathbf{i}]} + \overline{a_{k_1, q_1}[\mathbf{i}]} a_{k_2, q_2}[\mathbf{i}] \right) e^{-\iota(k_1 - k_2)\alpha} d\alpha \\ &= \delta[k_1 - k_2] \frac{1}{N^2} \sum_{\mathbf{i}} \mathcal{I}[\mathbf{i}] \Re\{a_{k_1, q_1}[\mathbf{i}] \overline{a_{k_2, q_2}[\mathbf{i}]} \}. \end{aligned} \quad (38)$$

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:

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

Thus,

$$\begin{aligned} b_3[\ell_1, \ell_2] &= \gamma \sigma^2 \langle a_{I_\omega}^1 \rangle_\omega \left(\delta[\ell_1 - \ell_2] + \delta[\ell_1] + \delta[\ell_2] \right) \\ &\approx \sum_{\substack{k_1, q_1 \\ k_2, q_2}} \mathfrak{b}_3[k_1, q_1; k_2, q_2] \psi_{k_1, q_1}[\ell_1] \overline{\psi_{k_2, q_2}[\ell_2]}. \end{aligned}$$

F.4 Connection to volume

Until now, we have established simple relations between the autocorrelations of the micrographs and the autocorrelations of the volume. Now, we complete the picture by deriving the connection with the volume itself.

Using the 2-D PSWFs, we can express each projection as

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

where

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

where the coefficients

$$\beta_{\ell, s; N, n} := \begin{cases} \frac{\sqrt{8\pi}}{\alpha_{N, n}} Y_\ell^N(\pi/2, 0) \int_0^1 j_{\ell, s}(k) R_{N, n}(k) k dk, & \ell \geq |N| \\ 0, & \ell < |N| \end{cases}, \quad (39)$$

can be precomputed. Since the PSWFs are eigenfunctions of the truncated Fourier transform [24] and hence satisfy

$$\alpha_{N, n} \psi_{N, n}(\mathbf{k}) = \int_{\|\mathbf{r}\|_2 \leq 1} \psi_{N, n}(\mathbf{r}) e^{\iota c(\mathbf{r} \cdot \mathbf{k})} d\mathbf{r}, \quad (40)$$

we can now express the projection in real space as

$$\begin{aligned} I_\omega(r, \varphi) &= \sum_{N,n} \hat{\alpha}_{N,n} b_{N,n}(\omega) \psi_{N,n}(r, \varphi) \\ &= \sum_{\ell=0}^{L_{\max}} \sum_{N,m=-\ell}^{\ell} \sum_{n=0}^{n_{\max}(N)} \sum_{s=1}^{S(\ell)} x_{\ell,m,s} \hat{\beta}_{\ell,s;N,n} D_{N,m}^\ell(\omega) \psi_{N,n}(r, \varphi), \end{aligned} \quad (41)$$

where $n_{\max}(N)$ is chosen according to Eq. (8) in [24], $\alpha_{N,n}$ is the eigenvalue corresponding to the (N, n) th PSWF, $\hat{\alpha}_{N,n} = (c/2\pi)^2 \alpha_{N,n}$, and $\hat{\beta}_{\ell,s;N,n} = \hat{\alpha}_{N,n} \beta_{\ell,s;N,n}$.

Third-order autocorrelation. In (30) and (38) we have shown how the third-order autocorrelations of the micrographs and the volume are related, and how we can present 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 (41):

$$\begin{aligned} \langle a_{I_\omega}^3[\ell_1, \ell_2] \rangle_\omega &= \frac{1}{L^2} \sum_{\mathbf{i}} \langle I_\omega[\mathbf{i}] I_\omega[\mathbf{i} + \ell_1] \overline{I_\omega([\mathbf{i} + \ell_2])} \rangle_\omega, \\ &\approx \frac{1}{L^2} \sum_{\substack{N_1, n_1 \\ N_2, n_2 \\ N_3, n_3}} \langle b_{N_1, n_1}(\omega) b_{N_2, n_2}(\omega) \overline{b_{N_3, n_3}(\omega)} \rangle_\omega \sum_{\mathbf{i}} \psi_{N_1, n_1}[\mathbf{i}] \psi_{N_2, n_2}[\mathbf{i} + \ell_1] \overline{\psi_{N_3, n_3}([\mathbf{i} + \ell_2])}, \end{aligned} \quad (42)$$

where the approximation is due to discretization. Now,

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

where the latter coefficients are given explicitly in (39). Using standard properties of D-Wigner functions, we obtain

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

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

$$\begin{aligned} \langle b_{N_1, n_1}(\omega) b_{N_2, n_2}(\omega) \overline{b_{N_3, n_3}(\omega)} \rangle_\omega &= \delta_{N_3, N_1 + N_2} \sum_{\substack{\ell_1, m_1, s_1 \\ \ell_2, m_2, s_2 \\ s_3}}^{\min(L, \ell_1 + \ell_2)} \sum_{\ell_3 = |\ell_1 - \ell_2|} x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} \overline{x_{\ell_3, m_1 + m_2, s_3}} \\ &\times (-1)^{N_1 + N_2 + m_1 + m_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ N_1 & N_2 & -N_1 - N_2 \end{pmatrix} \\ &\times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & -m_1 - m_2 \end{pmatrix} \hat{\beta}_{\ell_1, s_1; N_1, n_1} \hat{\beta}_{\ell_2, s_2; N_2, n_2} \overline{\hat{\beta}_{\ell_3, s_3; N_1 + N_2, n_3}}. \end{aligned} \quad (43)$$

Combining (43) with (42) 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 (38). Hence, to conclude the derivation we expand

$$\langle a_{I_\omega}^3[\ell_1, \ell_2] \rangle_\omega = \sum_{k, q_1, q_2} \mathfrak{a}_x^3[k, q_1, q_2] \psi_{k, q_1}[\ell_1] \overline{\psi_{k, q_2}[\ell_2]},$$

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

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

Then, the final formula reads

$$\begin{aligned} \mathfrak{a}_x^3[k, q_1, q_2] &= \sum_{\substack{\ell_1, m_1, s_1 \\ \ell_2, m_2, s_2 \\ s_3}}^{\min(L, \ell_1 + \ell_2)} \sum_{\ell_3 = |\ell_1 - \ell_2|} x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} \overline{x_{\ell_3, m_1 + m_2, s_3}} \\ &\times (-1)^{m_1 + m_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & -m_1 - m_2 \end{pmatrix} \\ &\times \sum_{N_1 = -\ell_1}^{\ell_1} \sum_{N_2 = -\ell_2}^{\ell_2} (-1)^{N_1 + N_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ N_1 & N_2 & -N_1 - N_2 \end{pmatrix} \\ &\times \frac{1}{L^2} \sum_{\mathbf{i}} \Psi_{\ell_1, N_1, s_1}[\mathbf{i}] \rho_{\ell_2, N_2, s_2}^{(k, q_1)}[\mathbf{i}] \overline{\rho_{\ell_3, N_1 + N_2, s_3}^{(k, q_2)}[\mathbf{i}]}. \end{aligned} \quad (44)$$

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

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 the volume is given in (29) and (37). The connection with the expansion coefficients of the volume can be derived in Fourier space directly from Kam's original formula [21] 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

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

and conclude that

$$\mathfrak{a}_x^2[q] = \frac{1}{\sqrt{8\pi}} \sum_{\substack{\ell, m \\ s_1, s_2}} x_{\ell, m, s_1} \overline{x_{\ell, m, s_2}} \int_0^1 j_{\ell, s_1}(k) j_{\ell, s_2}(k) R_{0, q}(k) k \, dk. \quad (45)$$

The integral on k is precomputed.

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 (25) that

$$\mathfrak{a}_x^1 = \langle a_{I_\omega}^1 \rangle_\omega = \widehat{\phi}(\mathbf{0}) = \frac{1}{\sqrt{4\pi}} \sum_s x_{0,0,s} j_{0,s}(0). \quad (46)$$