

Part I

Main text

Estimating the tails of the spectrum of the Hessian of the log-likelihood for *ab-initio* single-particle reconstruction in electron cryomicroscopy

Aaditya V. Rangan^{1,2*}, Wai-Shing Tang^{2,3}, Pilar Cossio^{2,3}, Kexin Zhang^{4,5}, Nikolaus Grigorieff^{4,5},

1 Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

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

3 Center for Computational Biology, Flatiron Institute, New York, USA

4 RNA Therapeutics Institute, University of Massachusetts Chan Medical School, Worcester, USA

5 Howard Hughes Medical Institute, Janelia Research Campus, Ashburn, USA * avr209@nyu.edu

Abstract

Electron cryomicroscopy (cryo-EM) is a technique in structural biology used to reconstruct accurate volumetric maps of molecules. One step of the cryo-EM pipeline involves solving an inverse-problem. This inverse-problem, referred to as *ab-initio* single-particle reconstruction, takes as input a collection of 2d-images – each a projection of a molecule from an unknown viewing-angle – and attempts to reconstruct the 3d-volume representing the underlying molecular density.

Most methods for solving this inverse-problem search for a solution which optimizes a posterior likelihood of generating the observed image-data, given the reconstructed volume. Within this framework, it is natural to study the Hessian of the log-likelihood: the eigenvectors and eigenvalues of the Hessian determine how the likelihood changes with respect to perturbations in the solution, and can give insight into the sensitivity of the solution to aspects of the input.

In this paper we describe a simple strategy for estimating the smallest eigenvalues and eigenvectors (i.e., the ‘softest modes’) of the Hessian of the log-likelihood for the *ab-initio* single-particle reconstruction problem. This strategy involves rewriting the log-likelihood as a 3d-integral. This interpretation holds in the low-noise limit, as well as in many practical scenarios which allow for noise-marginalization.

Once we have estimated the softest modes, we can use them to perform many kinds of sensitivity analysis. For example, we can determine which parts of the reconstructed volume are trustworthy, and which are unreliable, and how this unreliability might depend on the data-set and the imaging parameters. We believe that this kind of analysis can be used alongside more traditional strategies for sensitivity analysis, as well as in other applications, such as free-energy estimation.

1 Introduction

Electron cryomicroscopy (cryo-EM) is a popular technique in structural biology which is used to produce accurate volumetric maps of the electron density of macromolecular structures and biomolecules, such as proteins [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. Certain stages of a cryo-EM analysis pipeline involve solving an inverse-problem: e.g., given a collection of picked-particle images, the goal is to find the density map (i.e., volume) which produced them [11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22]. In this paper we will focus on one such inverse-problem, so-called ‘*ab-initio* single-particle reconstruction’. We will describe a simple strategy for analysing the robustness of solutions to this problem, and provide numerical methods for quantifying the uncertainty in the solution.

To provide some background, a standard single-particle reconstruction pipeline typically involves several stages:

1. Freezing a suspension containing multiple copies of the molecule: As described in [23, 24], there are several techniques for preparing a suspension of the molecule of interest, with details that depend on the properties of the molecule and its surrounding buffer. The goal for this stage is to produce a thin frozen film which can be penetrated by the transmission electron microscope, while also containing multiple well separated copies of the molecule.
2. Imaging the suspension to produce micrographs: The electron-beam interacts with the electron-density in the frozen film, and elastically scattered electrons in the beam are subject to a density-dependent phase-shift [9]. This phase-shift can be measured at each point in the image-plane, producing a ‘micrograph’. The values of the phase-shift can be related to the electron density within the sample. The details of this transformation depends on imaging-parameters such as the defocus level, which determines the contrast transfer function (CTF).

3. Scanning the micrographs for ‘picked-particles’: Once the micrographs are recorded, individual molecules within the micrograph must be ‘picked out’. There are several strategies for this ‘particle picking’ task [25, 26, 27, 28], and it is common practice to cluster the picked-particles, filter and denoise them [29, 30, 26]. The goal for this stage is to produce a collection of picked-particle images, each corresponding to a well-centered projection of a single molecule.
4. Analysis of the picked-particle images: At this stage, the viewing-direction (i.e., viewing-angle) of each picked-particle projection is unknown. Determining these viewing angles is typically done alongside reconstructing the molecular volume itself (see [16] but also [15, 17, 19]). Prior information regarding the molecule can often be used to aid in the reconstruction of the viewing-angles and volume. However, when such prior information is absent or deliberately ignored to avoid bias, the reconstruction process is referred to as *ab-initio* single-particle reconstruction [31, 11, 32, 16, 33, 34]).

Generally speaking, if we were to know the density $F(\mathbf{x})$ of a molecular volume, as well as a collection of image-specific viewing-angles $\{\tau_j\}$, we could use a forward model for the imaging process (involving the imaging-parameters and noise-level) to estimate what each picked-particle image $\{A_j\}$ might look like [31, 35]. The *ab-initio* single-particle reconstruction mentioned in the final step 4 above involves the reverse: taking the picked-particle images $\{A_j\}$ and attempting to infer the volume $F(\mathbf{x})$ and viewing-angles $\{\tau_j\}$. Thus, *ab-initio* reconstruction is a type of inverse-problem.

In practice, most methods for solving this inverse-problem search for a solution which optimizes a posterior likelihood of generating the observed picked-particle images, given the reconstructed volume (see [19, 14, 17, 1, 3, 2, 4, 13] and software from [21, 16, 20, 18, 11, 12]). More specifically, the solution is typically a ‘model’ that optimizes the log-likelihood $\log P(\text{data} | \text{model})$. Because the solution lies at a critical-point of $\log P(\text{data} | \text{model})$, any perturbations to the solution will result in a shift to the overall likelihood which depends on the Hessian of $\log P$.

While it is natural to consider the Hessian of $\log P$, it is not immediately clear how to calculate the spectrum of this object. For example, the posterior likelihood $P(\text{data} | \text{model})$ is usually defined using a standard noise-model (e.g., independent gaussian noise applied to each pixel in each picked-particle image) [31, 35]. In this scenario the likelihood $P(\text{data} | \text{model})$ is typically written as a product of the individual image-likelihoods $P(A_j | \text{model})$, each of which involves an accumulation of many exponential terms. Taking a second-derivative of the overall log-likelihood with respect to the model-parameters can be difficult, especially because the model can involve many degrees of freedom (e.g., the number of voxels in the volume $F(\mathbf{x})$ and the number of viewing-angles in $\{\tau_j\}$ can both be large).

In this paper we make a simple observation. Namely, when the noise-level is sufficiently small, the log-likelihood $\log(P(\text{data} | \text{model}))$ can be written as a single three-dimensional integral. This simplification allows for efficient calculation of the log-likelihood, as well as its gradient and Hessian. A similar simplification can often be made away from the small-noise limit by marginalizing over the noise-level.

Using this simple observation, we can approximate the spectrum of the Hessian of the log-likelihood. In particular, we can estimate the extremal eigenvalues and eigenvectors of the Hessian. The eigenvectors corresponding to the smallest eigenvalues are sometimes called the ‘softest modes’ in the molecular-dynamics literature [36]. The softest-mode indicates the volumetric- and viewing-angle-perturbations which are most easily confused for the current solution. The softest eigenvalue is a measure of the fisher-information restricted to the direction of this particular perturbation.

Calculation of the softest-modes allows for various kinds of sensitivity analysis. For example, we can use the soft subspace to determine which parts of the reconstructed volume are trustworthy, and which are unreliable. We can also use this approach to determine the benefit of adding more images to the image-pool; often the information gained will depend on how those new images are aligned to the volume, as well as other image properties such as the defocus-level. Additionally, when using cryo-EM data to estimate free-energy landscapes, we can use the spectrum of the hessian to help guide our choice of collective-variables. We believe that this kind of analysis can be used alongside more traditional strategies for sensitivity analysis, such as bootstrapping and crossvalidation. For example, we believe the structure of the soft subspace is a useful complement to the fourier-shell-correlation, which is often used to characterize the robustness of an *ab-initio* reconstruction [37, 38, 39, 40].

This paper is structured as follows. We first briefly review the mathematical details underlying our approach. We then provide several examples illustrating different applications of this kind of analysis. We conclude by discussing some extensions of this work. Many of the technical details are explained in the Appendix.

2 Methods Summary

In this section we briefly summarize our methodology, with the details deferred to the Appendix.

As described above, our methodology predominantly applies to the single-volume reconstruction problem. To set the stage, let’s assume that we are given a collection of J picked-particle images $\{A_j\}$. To ease the presentation for now, we’ll ignore the image-specific displacements $\{\delta_j\}$ and contrast-transfer-functions $\{CTF_j\}$.

For any particular image A and volume F , we can define the likelihood: $P(A | F, P(\tau); \sigma)$, which is the probability of observing the image A , given that the true signal is produced by the volume F with image-noise σ . In this expression the

term $P(\tau)$ is the prior-distribution of viewing-angles $\tau \in SO3$ associated with the data.

The setting we consider assumes that a single-particle reconstruction algorithm has already been run, and has produced a volume F which optimizes the posterior likelihood of observing the collection of images $\{A_j\}$:

$$P(\{A_j\} | F, P(\tau); \sigma) = \prod_{j=1}^J P(A_j | F, P(\tau); \sigma). \quad (1)$$

For our purposes we'll perform most of our calculations in fourier-space, referring to \hat{A}_j and \hat{F} as the 2d- and 3d-fourier-transforms of A_j and F , respectively, and using $\hat{\sigma}$ to denote the image-noise in fourier-space.

When the image-noise is relatively small, we can approximate the likelihood in Eq. 1 with the so-called 'low-temperature' limit:

$$P(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\}, \hat{\sigma}) \propto \exp\left(-\frac{1}{\hat{\sigma}^2} \cdot \mathcal{L}(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\})\right),$$

where the negative-log-likelihood \mathcal{L} takes the form of a sum of 2-dimensional integrals:

$$\mathcal{L}(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\}) \sim \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbb{R}^2} |\hat{A}_j(\mathbf{k}) - \hat{S}(\mathbf{k}; \tau_j^{\text{opt}}; \hat{F})|^2 d\mathbf{k}. \quad (2)$$

In this last expression we use $\hat{S}(\mathbf{k}; \tau; \hat{F})$ to denote the fourier-transform of a 2-dimensional projection of any given volume \hat{F} along viewing-angle τ , and denote by τ_j^{opt} the 'optimum' viewing-angle for image- j . As we'll discuss below, the viewing-angle τ_j^{opt} implicitly depends on \hat{F} , and is often referred to as the 'optimal-alignment' of \hat{A}_j to \hat{F} .

The expression in Eq. 2 can be rewritten as a single 3-dimensional integral:

$$\mathcal{L}(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\}) \sim \frac{1}{2} \iiint_{\mathbb{R}^3} \left(M_{20}(\boldsymbol{\kappa}) |\hat{F}(\boldsymbol{\kappa})|^2 - 2\Re(M_{11}^\dagger(\boldsymbol{\kappa}) \hat{F}(\boldsymbol{\kappa})) + M_{02}(\boldsymbol{\kappa}) \right) \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|}, \quad (3)$$

where M_{20} , M_{11} , M_{02} are each 3-dimensional functions corresponding to the various moments of the data, and the integral is performed over frequency-space $\boldsymbol{\kappa} \in \mathbb{R}^3$.

Considering the image-pool $\{\hat{A}_j\}$ to be fixed, we can consider the negative-log-likelihood \mathcal{L} to be a function of the volume \hat{F} and optimal-alignments $\{\tau_j^{\text{opt}}\}$. By using both Eq. 2 and 3, we can take various derivatives of this log-likelihood. For our purposes we will assume that, with the image-pool fixed, the volume and the optimal-alignments $[\hat{F}, \{\tau_j\}]^\top$ are situated at a local minimum of $\mathcal{L}(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\})$. Consequently, we expect that perturbations to these quantities will produce second-order perturbations to \mathcal{L} of the form:

$$\mathcal{L}(\{\hat{A}_j\} | \hat{F} + \Delta\hat{F}, \{\tau_j^{\text{opt}} + \Delta\tau_j\}) = \mathcal{L}(\{\hat{A}_j\} | \hat{F}, \{\tau_j^{\text{opt}}\}) + \frac{1}{2} \cdot \begin{bmatrix} \Delta\hat{F} \\ \{\Delta\tau_j\} \end{bmatrix}^\dagger \cdot \begin{bmatrix} H_{FF} & H_{F\tau} \\ H_{\tau F} & H_{\tau\tau} \end{bmatrix} \cdot \begin{bmatrix} \Delta\hat{F} \\ \{\Delta\tau_j\} \end{bmatrix}. \quad (4)$$

In this formula we use ' \dagger ' to denote the adjoint (i.e., conjugate-transpose). We also use the terms H_{FF} , $H_{F\tau}$, $H_{\tau F}$ and $H_{\tau\tau}$ to refer to blocks of the second-derivative (i.e., Hessian) of \mathcal{L} . The diagonal-blocks H_{FF} and $H_{\tau\tau}$ are trivial to calculate: the former is diagonal in fourier-space while the latter decomposes into small block-diagonal elements for each image. The off-diagonal-blocks $H_{F\tau}$ and $H_{\tau F}$ require slightly more care. In our calculation we use the volumetric-representation in Eq. 3 to calculate $H_{F\tau}$ for any particular choice of $\{\Delta\tau_j\}$, and we use the 2-dimensional representation in Eq. 2 to calculate $H_{\tau F}$ for any particular choice of $\Delta\hat{F}$.

In most practical scenarios of interest the volume \hat{F} and optimal-alignments $\{\tau_j^{\text{opt}}\}$ are implicitly related: the optimal-alignment τ_j^{opt} for any particular image \hat{A}_j is the viewing-angle for which that image aligns best to \hat{F} . Thus, we can consider the τ_j^{opt} each to be functions of \hat{F} , and correspondingly interpret \mathcal{L} as a function of \hat{F} alone.

With this perspective the negative-log-likelihood takes the form:

$$\mathcal{L}(\{\hat{A}_j\} | \hat{F} + \Delta\hat{F}) = \mathcal{L}(\{\hat{A}_j\} | \hat{F}) + \frac{1}{2} \cdot \Delta\hat{F}^\dagger \cdot [H] \cdot \Delta\hat{F}, \quad (5)$$

with $[H]$ representing the Hessian of $\mathcal{L}(\{\hat{A}_j\} | \hat{F})$ with respect to volumetric-perturbations of \hat{F} .

With this formalism in place, we can immediately see that the influence of any perturbation $\Delta\hat{F}$ on the likelihood is determined by the rayleigh-quotient. That is to say, the log-likelihood-ratio

$$\Delta\mathcal{L} := \mathcal{L}(\{\hat{A}_j\} | \hat{F} + \Delta\hat{F}) - \mathcal{L}(\{\hat{A}_j\} | \hat{F}) = \frac{1}{2} \cdot \Delta\hat{F}^\dagger \cdot [H] \cdot \Delta\hat{F}, \quad (6)$$

contrasting the original volume \widehat{F} with the perturbed volume $\widehat{F} + \Delta\widehat{F}$ is given by the rayleigh-quotient $\Delta\widehat{F}^\dagger[H]\Delta\widehat{F}$. Thus, given any volumetric-perturbation $\Delta\widehat{F}$, the rayleigh-quotient associated with that $\Delta\widehat{F}$ can be thought of as a measure of ‘information’ in that direction.

More generally, the critical-points of the rayleigh-quotient correspond to the eigenvectors of $[H]$. Eigenvectors corresponding to the largest eigenvalues correspond to volumetric-perturbations $\Delta\widehat{F}$ for which the nearby volumes $\widehat{F} \pm \Delta\widehat{F}$ are most easily distinguishable from the original volume \widehat{F} . I.e., these ‘stiff modes’ correspond to volumetric-perturbations which produce a significant mismatch between the projections of the perturbed volume and the image-pool. On the other hand, the eigenvectors corresponding to the smallest eigenvalues indicate volumetric-perturbations for which the nearby volumes are *least* distinguishable from the original \widehat{F} . I.e., these ‘soft modes’ correspond to volumetric-perturbations which produce only subtle shifts to the relevant projections of the perturbed volume; shifts that might still allow for plausible alignment to the images.

With this in mind, we are motivated to estimate the spectrum of H from Eq. 5, specifically the lower-tail of the spectrum, comprising the smallest eigenvalues as well as their corresponding eigenvectors (i.e., the softest-modes). In terms of computational details, the Hessian H involves an accumulation of terms combining the various blocks H_{FF} , $H_{F\tau}$, $H_{\tau F}$ and $H_{\tau\tau}$ from Eq. 4 with the volumetric-perturbations $\Delta\widehat{F}$ and the (implicitly-determined) alignment-perturbations $\{\Delta\tau_j\}$. As a result, the H from Eq. 5 is not quite as easy to diagonalize as the H_{FF} from Eq. 4. However, we can still compute quantities such as $H \cdot \Delta\widehat{F}$ relatively easily, allowing us to use iterative strategies such as lanczos-iteration to estimate the extreme eigenvalues and eigenvectors of H . During this calculation we make sure to focus only on true deformations, projecting away any combinations of volumetric- and alignment-perturbations $[\Delta\widehat{F}, \{\Delta\tau_j\}]$ which correspond to trivial rigid rotations.

As mentioned above, the details of this calculation are provided in the Appendix. Importantly, for our calculation, the soft modes (as well as the rest of the spectrum of the Hessian) depend on the volume \widehat{F} and the image-pool $\{\widehat{A}_j\}$, as well as on the optimal-alignments for the images (which are determined implicitly). More generally, the soft modes also depend on the image-specific CTF-functions $\{CTF_j\}$ (see applications below) as well as other image-parameters, such as the image-specific displacements $\{\delta_j\}$ (mentioned in the discussion).

3 Applications

As a straightforward application, one can simply inspect the softest modes of a given single-particle reconstruction. The softest modes can then be used to directly estimate how robust that reconstruction might be. More specifically, the magnitude of the ‘softest eigenvalue’ associated with the softest mode determines the size of the largest volumetric-perturbation which can be ‘confused’ with the original reconstruction. That is to say, the largest volumetric-perturbation that can’t easily be rejected in favor of the original reconstruction (under a standard hypothesis test). The softest modes can also be used to identify (i) which regions of the original reconstruction are least trustworthy, and (ii) which images-alignments are least trustworthy.

3.1 case-study: trpv1

As a first case-study we’ll analyze the single-particle reconstruction of the trpv1 molecule described by the volume EMD-5778 [41]. This single-particle reconstruction was originally performed using the image-pool from EMPIAR-10005. This image-pool is of very high quality and, as we’ll show later, the single-particle reconstructions formed using this image-pool are incredibly robust.

Thus, to illustrate our methodology, we’ll first construct a hypothetical scenario where we use an image-pool that is not as good as the full image-pool from EMPIAR-10005. To do this, we’ll consider the scenario shown in Fig 1 where the number of images J is not too large (say, several hundred to a thousand), and the image viewing-angles are not well distributed. This hypothetical scenario might be similar to situations involving *ab-initio* low-resolution reconstruction, where the picked-particle images are limited in number and type (e.g., drawn from only a few similar class averages).

On the left of Fig 1 we illustrate our hypothetical viewing-angle distribution, which is technically nonzero everywhere, but heavily concentrated towards the poles. We’ll imagine drawing a collection of J synthetic-images (i.e., noisy projections) of the true molecule from these viewing-angles. We’ll further assume that these synthetic-images $\{\widehat{A}_j^{\text{syn}}\}$ are each perfectly centered, and that there is no attenuation from a CTF. With these assumptions, we’ll imagine that we have performed an *ab-initio* low-resolution single-particle reconstruction (e.g., using [42]).

In practice, a single-particle reconstruction can easily produce any one of a number of low-quality volumes that lie at various local minima of the likelihood function. Thus, to ensure that our example is reproducible, we’ll assume that the single-particle reconstruction was indeed of high-quality, producing the volume shown on the right of Fig 1. We have chosen this volume to agree with the true volume up to the resolution of the reconstruction. Thus, for this particular case-study, we know that this volume is a ‘global optimum’ of the likelihood in Eq. 2.




Figure 1: On the left we show an idealized viewing-angle-distribution $\mu(\tau)$ which is concentrated around the poles. On the right we show a low-res ($2\pi K = 48$) reconstruction F of the TRPV1 molecule obtained by appropriately aligning many synthetic-images (i.e., noisy projections) taken from these equatorial viewing-angles. This low-res reconstruction F agrees with the reconstruction EMD-5778 referenced in EMPIAR-10005.

Given this scenario, we can use (i) the reconstructed volume \hat{F} and (ii) the synthetic-image-stack $\{\hat{A}_j^{\text{syn}}\}$ to calculate the softest mode of the Hessian $[H]$ from Eq. 5. This softest mode, denoted by $[\Delta\hat{F}, \{\Delta\tau_j\}]$ is illustrated in Fig 2.

On the right side of Fig 2 we show the volumetric-perturbation $\Delta\hat{F}$, and on the left we show the (implicitly determined) alignment-perturbations associated with each original viewing-angle. We illustrate another view of the same alignment-perturbations in Fig 3. Note that, while the calculation of this softest mode did not explicitly include the viewing-angle distribution shown in Fig 1, this distribution determines the distribution of the synthetic-images in the image-pool, and thus influences the Hessian $[H]$, playing a strong role in the calculation.

Indeed, one can readily see that the softest mode ΔF is structured so that polar-projections of ΔF have very little variation. This in turn implies that any polar-projection of a perturbed volume $F + \Delta F$ would look almost exactly like the corresponding polar-projection of the original volume F . Put another way, if we consider a particular synthetic-image \hat{A}^{syn} constructed using a polar viewing-angle τ , then the likelihood $P(\hat{A}^{\text{syn}} | \hat{F}, \tau)$ of observing \hat{A}^{syn} given \hat{F} and τ will be very close to the likelihood $P(\hat{A}^{\text{syn}} | \hat{F} \pm \Delta\hat{F}, \tau \pm \Delta\tau(\tau))$ of observing the same synthetic-image given the perturbed volume $\hat{F} \pm \Delta\hat{F}$ and the perturbed viewing-angle $\tau \pm \Delta\tau(\tau)$. A similar story holds across the pool of synthetic-images (drawn with viewing-angles concentrated near the poles, as shown in Fig 1). Accumulating these calculations across the image-pool produces an overall shift in likelihood characterized by the rayleigh-quotient in Eq. 6, which ΔF is specifically constructed to minimize.




Figure 2: In this figure we show a collection of τ -specific alignment-perturbations $\{\Delta\tau(\tau)\}$ (left) along with a corresponding volumetric-perturbation δF (right). The arrows on the left indicate the direction of the alignment-perturbations for each (β, α) , with $\Delta\gamma(\tau) = 0$. Larger/darker arrows indicate larger values of $\Delta\tau(\tau)$, with the arrow-area proportional to $|\Delta\tau|$ (and dark purple arrows corresponding to the largest values of $|\Delta\tau|$). Locations with individual perturbations smaller than 10^{-4} in relative magnitude are left blank. On the right we show an isosurface of $|\Delta F|$, colored by the sign of ΔF (with blue and pink corresponding to negative and positive, respectively). The projections of ΔF are shown along each axis (background and bottom). Note that projections of this softest mode from the equator (background and sides) have large variations and are easily visible. By contrast, projections of this softest mode from the poles (bottom) have essentially no variation and are difficult to detect. The softest mode has this structure precisely because the synthetic-image pool included images corresponding predominantly to polar viewing-angles.

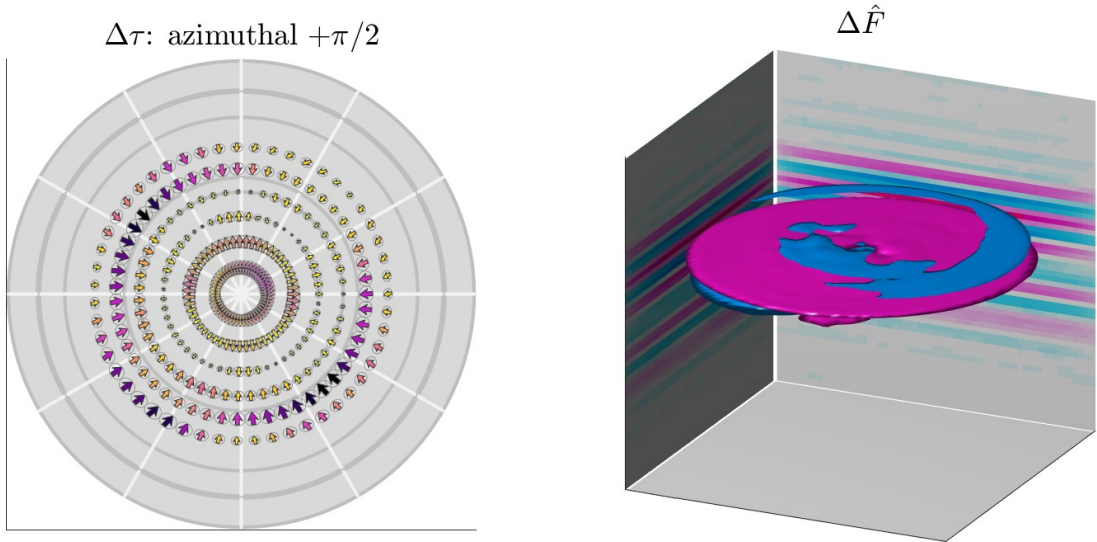


Figure 3: This figure is similar to Fig 2. In this figure we show a collection of τ -specific alignment-perturbations $\{\Delta\tau(\tau)\}$ (left) along with a corresponding volumetric-perturbation δF (right). The arrows on the left indicate the direction of the alignment-perturbations for each (β, α) , with $\Delta\gamma(\tau) = 0$. Note that, unlike Fig 2, we display $\Delta\tau(\tau)$ from the poles, rather than from the equator. On the right we show an isosurface of $|\Delta F|$, colored by the sign of ΔF (with blue and pink corresponding to negative and positive, respectively). As in most of our figures, the values of $\Delta\tau(\tau)$ are implicitly linked to the volumetric-perturbation ΔF . These perturbations have been selected to align with the ‘softest’ observed eigenvector of the Hessian H associated with the single-particle-reconstruction problem associated with Fig 1. For this calculation we use synthetic-images (i.e., noisy-projections of the original volume F).

The panels in Figs 2-3 illustrate the structure of the softest mode, but not its interpretation with regards to the sensitivity of the reconstruction (which depends on the magnitude of the associated softest eigenvalue). Such an illustration

is provided in Fig 4. This shows ‘nearby’ volumes of the form $F \pm \Delta F$, with ΔF oriented along the softest mode. The magnitude of these perturbations is chosen so that the log-likelihood-ratio from Eq. 6 is equal to $\log(20)$, implying that the perturbed-volume is only 20 times less likely than the original (optimal) volume F . Thus, the perturbed-volumes represent the ‘most extreme’ volumetric perturbations that would not be rejected in favor of the true volume under a standard hypothesis test. In this way, the mode’s structure can be used to quantify the magnitude of the deformations associated with the volumetric uncertainty, and also identify which (implicitly linked) alignment-perturbations are largest. In terms of technical details, for this calculation we have assumed $J = 1024$ images and used a signal-to-noise ratio estimated from the image-stack in EMPIAR-10005.

Another illustration of the same perturbations is shown in Fig 5. This time the isosurface of each perturbed-volume is colored to indicate the magnitude of the perturbation with respect to the original isosurface. One can immediately see that volumetric-fluctuations along this particular mode induce variations along certain portions of the isosurface, but not others.




Figure 4: This figure illustrates level surfaces of $F \pm \Delta F$, with ΔF chosen to be the softest-mode taken from Fig 3. These level-surfaces are colored in accordance with the sign and magnitude of ΔF as it intersects the level-surface of $F \pm \Delta F$. The magnitude of the perturbation is chosen so that the likelihood-ratio between these perturbed volumes and the original volume (from Fig 1) is $1/20$. In this case the magnitude of the perturbation has an l_2 -norm which is $\sim 10\%$ of the l_2 -norm of the original volume. For this calculation we have assumed $J = 1024$ images and used a signal-to-noise ratio estimated from the image-stack in EMPIAR-10005.




Figure 5: This figure is similar to Fig 4. This time the level-surfaces of $F \pm \Delta F$ are colored in accordance with the distance between each level-surface shown and the corresponding level-surface of the undeformed volume F . The color ranges from white (no perturbation) to dark-purple ($1/400$ of the viewing-window diameter).

For this particular scenario the second-softest mode also has a small eigenvalue, and is illustrated in Figs 6-7. Note that, by considering the ‘soft eigenspace’ comprising these two softest modes, one can see that the lower region of the reconstruction is less trustworthy than the upper region. We remark that this uncertainty is *not* directly due to the underlying biology (e.g., the flexibility of the ankyrin domains), but rather to the inherent sensitivity of the single-particle reconstruction problem for this particular volume and image-pool. That is to say, the softest modes reveal how well posed the single-particle reconstruction problem is.




Figure 6: This figure is similar to Fig 3, except that we show the second softest observed eigenvector of H . That is to say, the ΔF shown in this figure is orthogonal to the ΔF shown in Fig 3.




Figure 7: This figure is similar to Fig 5, except this time we choose ΔF to be the second-softest mode taken from Fig 6. In this case the magnitude of the perturbation has an l_2 -norm which is $\sim 8\%$ of the l_2 -norm of the original volume. Note that, again, the volumetric-perturbations indicate uncertainty regarding the level-surface towards the bottom of the molecule. Note also that the upper part of the molecule is relatively stable (c.f. Figs 5 and 7).

3.2 case-study: trpv1 (empirical viewing-angle distribution)

As an illustration of this fact, we repeat the same calculation, but this time using the empirical distribution of viewing-angles taken from the first $J = 1024$ images of EMPIAR-10005, as shown in Fig 8. We perform this calculation using the empirically determined displacements and CTF-functions for these images. The softest mode is shown in Fig 9.

What is not immediately apparent from Fig 9 is that, even though a handful of images admit large alignment-perturbations, the softest eigenvalue is actually much larger for this example than for the original example shown in Fig 2. This larger eigenvalue implies that the likelihood penalty for a volumetric-deformation is much larger, implying in turn much less uncertainty, as illustrated in Fig 10. While the lower region of the molecule is still the most uncertain, the overall level of uncertainty is very small; the difference between the original level surface and the deformed surfaces in Fig 10 is barely noticeable.

We remark that the increased reliability of this second reconstruction (c.f. Figs 10 and 5) is due almost entirely to the difference in viewing-angles (c.f. the viewing-angle distributions in Figs 8 and 1). Put simply, the broader distribution of viewing-angles in the second reconstruction renders the result much more robust. This observation by itself aligns with intuition, and is not surprising. Nevertheless, we believe that calculations of this kind can be useful for identifying which regions of the molecule are less reliable than others, and quantifying just how unreliable those regions might be.




Figure 8: This figure is similar to Fig 1, except this time we illustrate the empirical viewing-angle distribution for the picked-particles within EMPIAR-10005.

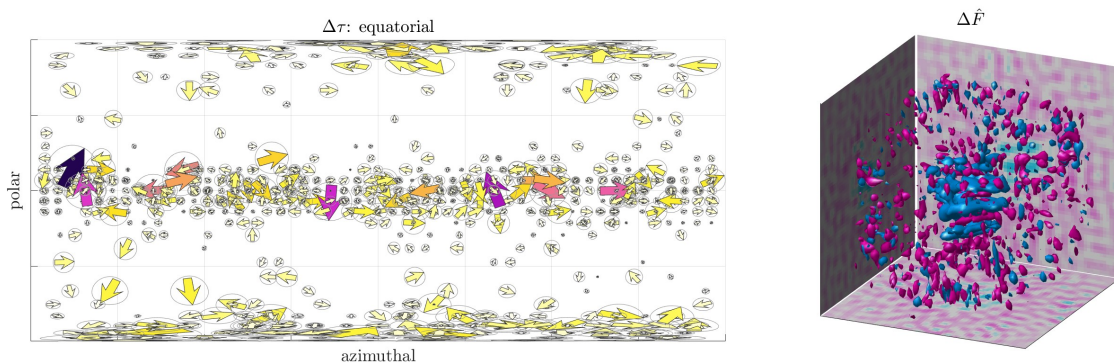


Figure 9: This figure is similar to Fig 2. In this figure we show a collection of τ -specific alignment-perturbations $\{\Delta\tau(\tau)\}$ (left) along with a corresponding volumetric-perturbation δF (right). The arrows on the left indicate the direction of the alignment-perturbations for each (β, α) , with $\Delta\gamma(\tau) = 0$. On the right we show an isosurface of $|\Delta F|$, colored by the sign of ΔF (with blue and pink corresponding to negative and positive, respectively). These perturbations have been selected to align with the ‘softest’ observed eigenvector of the Hessian H associated with the single-particle-reconstruction problem associated with Fig 1. For this calculation we use the first $J = 1024$ images from EMPIAR-10005, along with their CTF-functions and their most-likely displacements. Note that, while most of the alignments are relatively robust, there are a few images for which the viewing-angles can be significantly deformed (see dark-purple arrows on the left).




Figure 10: This figure is similar to Figs 5 and 7, except this time we choose ΔF to be the softest mode taken from Fig 9. In this case the magnitude of the perturbation has an l_2 -norm which is only $\sim 2\%$ of the l_2 -norm of the original volume. Note that, this time, the level-surfaces do not change very much (i.e., the entire level-surface is relatively stable).

3.3 case-study: the effect of viewing-angle distribution

The examples above involved the same reconstruction of trpv1, subject to different distributions of viewing-angles. While the polar-distribution in Fig 1 was less robust than the empirical distribution in Fig 8, both had softest modes which indicated that the lower region of the molecule was the most uncertain.

This is not always the case, and in some scenarios shifts in the viewing-angle distribution can dramatically change the structure of the softest modes. One example of this is shown in Figs 11 and 12. In this example we use $J = 1024$ synthetic-images to explore the spectrum of H for the ISWINCP molecule at `emd_9718.map`. The first scenario in Fig 11 assumes a uniform distribution of viewing-angles, while the second scenario in Fig 12 assumes a collection of viewing-angles that is concentrated towards the equator. Note that the regions of the molecule which are most uncertain depend on the viewing-angle distribution.




Figure 11: This figure is similar to Fig 10. Here we show the softest mode associated with the ISWINCP molecule at `emd_9718.map`. For this calculation we use synthetic-images with a resolution of $2\pi K = 48$, assuming uniform viewing-angles.




Figure 12: This figure is similar to Fig 10. Here we show the softest mode associated with the ISWINCP molecule at `emd_9718.map`. For this calculation we use synthetic-images with a resolution of $2\pi K = 48$, assuming viewing-angles concentrated in a band around the equator (± 15 degrees).

Another example of this is shown in Figs 13-15. In this example we use $J = 1024$ synthetic-images to explore the sensitivity of the MlaFEDB molecule at `emd_9718.map`. The first scenario in Fig 13 assumes a viewing-angle distribution concentrated around the poles. The second scenario in Fig 14 assumes a viewing-angle distribution that is uniform. The third scenario in Fig 15 assumes a viewing-angle distribution concentrated around the equator. Note that, for this molecule, the regions which are most uncertain depend strongly on the viewing-angle distribution.




Figure 13: This figure is similar to Fig 10. Here we show the softest mode associated with the MlaFEDB molecule at emd_9718.map. For this calculation we use synthetic-images with a resolution of $2\pi K = 48$, assuming viewing-angles concentrated in a band around the poles (± 30 degrees).




Figure 14: This figure is similar to Fig 10. Here we show the softest mode associated with the MlaFEDB molecule at emd_9718.map. For this calculation we use synthetic-images with a resolution of $2\pi K = 48$, assuming uniform viewing-angles.




Figure 15: This figure is similar to Fig 10. Here we show the softest mode associated with the MlaFEDB molecule at `emd_9718.map`. For this calculation we use synthetic-images with a resolution of $2\pi K = 48$, assuming viewing-angles concentrated in a band around the equator (± 15 degrees).

3.4 case-study: the effect of defocus

In the previous section we used the softest modes to illustrate how changes in the viewing-angle distribution might impact the sensitivity of single-particle-reconstruction. A similar phenomenon also applies to changes in defocus. To illustrate this, we turn back to the `trpv1` molecule used above.

We consider three scenarios, corresponding to three different distributions of defocus (see Fig 16). Now, for each of these defocus-values, we can estimate the lower-tail of the Hessian for the optimal single-particle-reconstruction. For these calculations we use synthetic-images drawn from the empirical-distribution of viewing-angles. These results are shown in Fig 17. Note that, as the defocus increases, the softest eigenvalues of the Hessian increase, corresponding to the increase in information shown in the CTFs in Fig 16. The softest modes for each scenario are shown in Figs 18-20. The corresponding volumetric-deformation for the largest defocus-value is shown in Fig 21.

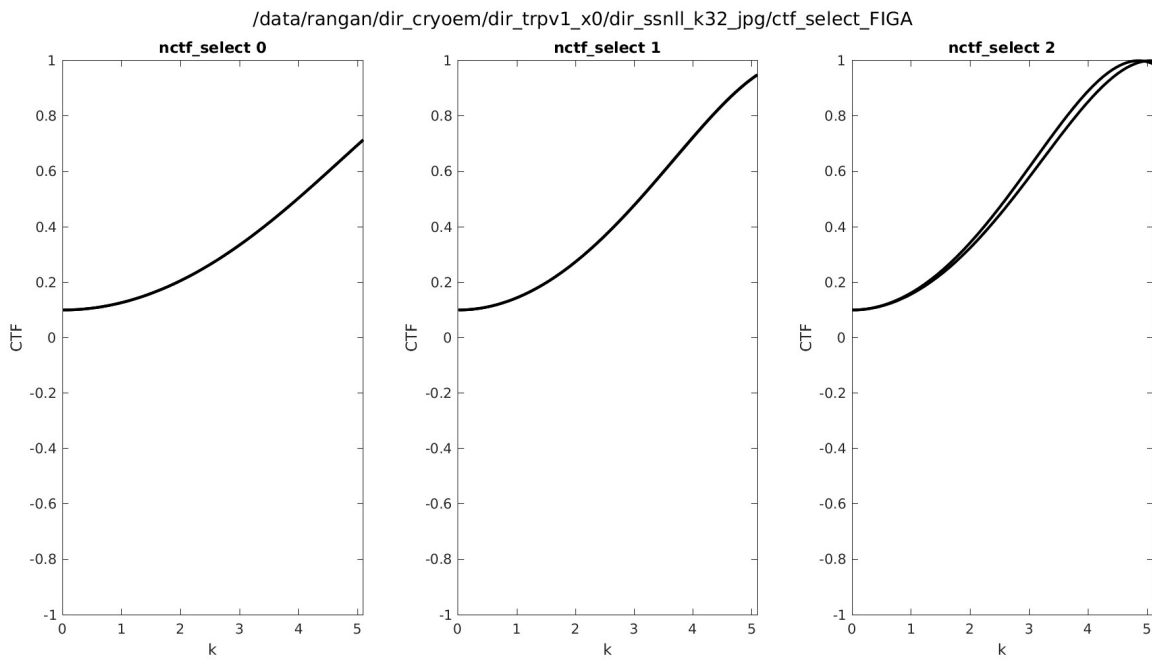


Figure 16: In this figure we illustrate three different collections of CTF-functions, corresponding to defocus-values of $\sim 10k$, $17k$ and $23k$, respectively. Note that, for this calculation, $2\pi K$ was limited to 32. As a result, the larger defocus-values carry more information.




Figure 17: Here we show the rayleigh-quotients associated with the softest 7 modes for each defocus-value. Note that, for this calculation, $2\pi K$ was limited to 32. Note also that the effective dimension of the ‘soft subspace’ shrinks as the defocus-value increases.




Figure 18: This figure is similar to Fig 9. In this case we show the perturbation associated with the softest mode the defocus of $\sim 10k$. Note that, for this calculation, $2\pi K$ was limited to 32. Note also that in this figure (and the following figures) we use the same area for each arrowhead on the left; the magnitude of $\Delta\tau(\tau)$ is represented only by the color.

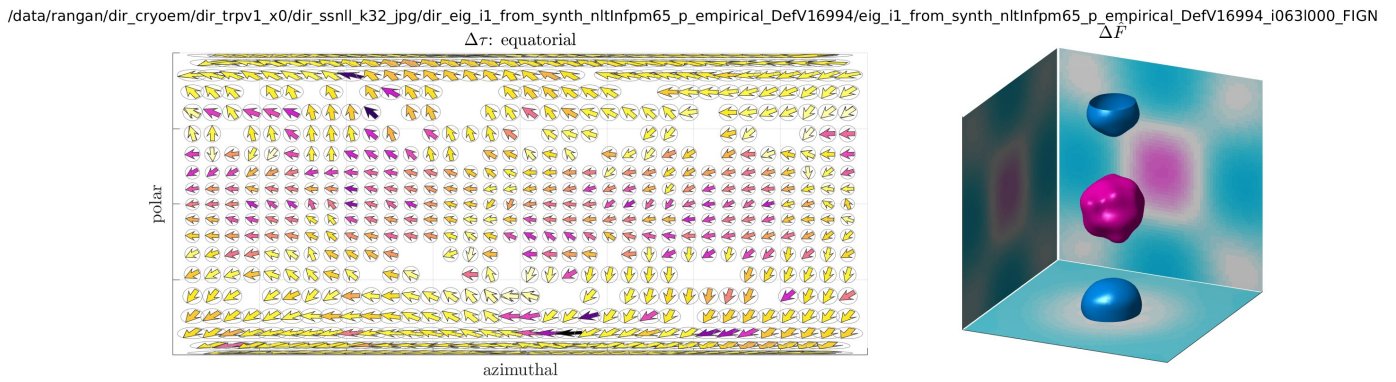


Figure 19: This figure is similar to Fig 18. In this case we show the perturbation associated with the softest mode the defocus of $\sim 17k$. Note that the structure of the perturbation changes (as a result of the change in defocus).




Figure 20: This figure is similar to Fig 18. In this case we show the perturbation associated with the softest mode the defocus of $\sim 23k$. Note that the structure of the perturbation changes (as a result of the change in defocus).

'dir_ssnll_k32_jpg/dir_eig_i1_from_syntn_nltlnfpm65_p_empirical_DefV23120/eig_i1_from_syntn_nltlnfpm65_p
 $\Delta\hat{F}$: $-0.006, J$ $\Delta\hat{F}$: $+0.006, J$




Figure 21: This figure is similar to Fig 10. In this case we show the perturbation associated with the softest mode at a defocus of $\sim 23k$. The level-surfaces change most towards the bottom of the molecule. A very similar structure manifests for the other defocus-values (not shown).

3.5 case-study: paths in volume-space

As demonstrated above, the structure of the softest subspace depends on the viewing-angle distribution, as well as on imaging parameters such as the defocus. The softest subspace also obviously depends on the particular volume which is provided as a solution to the single-particle-reconstruction problem. This dependence can be of interest in certain scenarios, such as free-energy estimation [43].

To illustrate this phenomenon, we consider a one-dimensional family of conformations for an RNA-molecule, shown in Fig 22. This family of conformations corresponds to a particular path of interest in volume-space, and can be parametrized by a ‘collective variable’ s . As described in [43], it is in principle possible to use cryo-EM data to estimate the free-energy density $G(s)$ along this path. This kind of free-energy estimate involves calculating, for each volume $F(\mathbf{x}; s)$ along the path, the likelihood that the image-pool might produce that particular volume. Those values of s for which $F(\mathbf{x}; s)$ is more likely will correspond to lower values of $G(s)$.

When performing this calculation, it is important to consider the interplay between the path-tangent $\partial_s F(\mathbf{x}; s)$ (i.e., the direction in volume-space corresponding to perturbations in the path-variable s) and the softest modes of the Hessian $H(s)$ of the likelihood at each location along the path. Locations s where the path-tangent $\partial_s F(\mathbf{x}; s)$ lies along a stiff-mode of the Hessian $H(s)$ correspond, locally, to locations along the path that are highly discriminable from their neighbors (given the image-pool). In other words, these are locations where projections of $F(\mathbf{x}; s)$ are not easily confused for projections of nearby volumes $F(\mathbf{x}; s \pm \Delta s)$. Conversely, locations where the path-tangent $\partial_s F(\mathbf{x}; s)$ lies parallel to a soft-mode of the Hessian $H(s)$ correspond, locally, to locations along the path that are easily confused with their neighbors. In other words, these are locations where projections of $F(\mathbf{x}; s)$ are quite similar to projections of nearby volumes $F(\mathbf{x}; s \pm \Delta s)$.

Thus, the reliability of any free-energy estimate $G(s)$ will be influenced by the rayleigh-quotient:

$$\mathcal{R}(s) := [\partial_s \hat{F}(\cdot; s)]^\dagger [H(s)] [\partial_s \hat{F}(\cdot; s)].$$

Locations along the path for which $\mathcal{R}(s)$ is small will correspond to locations where the estimate of $G(s)$ is inherently unreliable, simply due to the poor conditioning of the likelihood function. Estimates of $\mathcal{R}(s)$ are shown in Fig 22. This calculation was performed in the low-temperature limit, for which the likelihood can be approximated by Eq. 1. Similar calculations can be performed away from the low-temperature limit by marginalizing over the noise-parameter (see [44] and section 23 in the Appendix).

For the collective-variable shown in Fig 22, the values of $\mathcal{R}(s)$ are smallest when the RNA-molecule is ‘closed’; i.e., there are many projections along which neighboring closed configurations seem quite similar to one another. By contrast, the values of $\mathcal{R}(s)$ are largest when the RNA-molecule is ‘open’; i.e., there are comparatively fewer projections along

which neighboring open configurations seem similar. These differences in $\mathcal{R}(s)$ can be important when interpreting free-energy estimates. Moreover, it may be worth considering $\mathcal{R}(s)$ when designing a collective variable. By ensuring that the path-tangent travels along stiff-modes of the Hessian of the likelihood-function for each point in volume-space, it may be possible to mitigate one source of uncertainty in the final free-energy estimate.




Figure 22: In this figure we illustrate the rayleigh-quotient $\mathcal{R} := [\partial_s \widehat{F}(\cdot; s)]^\dagger [H(s)] [\partial_s \widehat{F}(\cdot; s)]$ (top) along the conformational path corresponding to the sequence of volumes shown on the bottom. The path shown here connects an ‘open’ RNA molecular conformation (left) to a ‘closed’ configuration (right). Portions of the path for which \mathcal{R} is low correspond to zones where locally resolving the free-energy will be difficult, due simply to the fact that other nearby conformations produce similar projections. Conversely, portions of the path for which \mathcal{R} is high correspond to zones where locally resolving the free-energy will be more robust, with less opportunity for confusion. In the calculation of \mathcal{R} we assume that (i) the values of $\Delta\tau(\tau)$ are determined implicitly, with (ii) $\|\partial_s F(\mathbf{x}, s)\| \equiv 1$ at each point along the path. For demonstration, we have also assumed that the viewing-angle-distribution is uniform (for each conformation), and that the CTF is identically 1 (although these assumptions can easily be relaxed). This calculation was performed at relatively low resolution, with $2\pi K = 32$.

4 Discussion

In cryo-EM, single-particle reconstruction algorithms typically search for volumes which maximize a likelihood function. Consequently, the effect of any local perturbation depends on the Hessian of the log-likelihood. In the context of *ab-initio* reconstruction, this log-likelihood can often be approximated with either a low-temperature limit or a noise-marginalization, yielding the simple accumulation of 2d-integrals shown in Eq. 1. This simple structure allows the log-likelihood to be represented as a volumetric integral akin to Eq. 3. Together, these 2d- and 3d-representations of the log-likelihood allow us to recover the extremal eigenvalues and eigenvectors of the Hessian, opening the door to several kinds of sensitivity analyses. We believe that these analyses can complement more traditional strategies for sensitivity analysis, such as bootstrapping and crossvalidation (e.g., measuring the fourier-shell-correlation between reconstructions built using disjoint subsets of the image-pool) [37, 38, 39, 40].

In the applications described above, we have illustrated how the ‘softest modes’ of the hessian can be used to quantify how robust a particular single-particle reconstruction might be. For example, one can examine the softest subspace to identify which regions of the molecule are most uncertain. The same can be said for the image-alignments, and the softest modes can quantify which (implicitly determined) image-alignments are most uncertain. The softest subspace of the Hessian can also be used to probe how the uncertainty depends on properties of the image-pool, such as the viewing-angle distribution and range of defocus-values.

We expect this kind of analysis to be useful in scenarios where the viewing-angle distribution is naturally nonuniform, and where experimentalists have a range of defocus values to choose from. One such example might be the reconstruction

of membrane proteins from picked-particle images of cellular lammellae [45, 46]. There are scenarios where the picked-particles can be biased towards regions where the cell membrane intersects the micrograph transversely. In such scenarios, the membrane-proteins in any picked-particle images would typically be viewed from a ‘side’ (i.e., a direction tangent to the membrane) rather than from the ‘top’ (i.e., the direction orthogonal to the membrane). As illustrated above, nonuniformity in the viewing-angle distribution can result in large uncertainties for particular regions of the molecule. This can be particularly problematic when these large uncertainties occur in regions of interest (e.g., the pores of a membrane channel). By analyzing how the softest modes of the Hessian depend on the viewing-angle distribution and the imaging parameters, one can determine which additional viewing-angles (and choice of defocus-values) would help reduce this uncertainty the most.

While most of our examples involve single-particle reconstruction, the structure of the Hessian can provide important information in other applications as well. For example, when using cryo-EM data to estimate free-energy landscapes one often considers paths in conformation-space. Each location along a conformational-path can be associated with a conformation-specific Hessian, which in turn produces conformation-specific soft modes. In such a scenario, the relative orientation between the conformational-tangent-space and the softest subspace of the conformation-specific Hessian plays a role in determining the local sensitivity of any free-energy estimate.

The analysis of the Hessian can also shed light on some of the difficulties associated with multi-particle reconstruction. As described in section 24 of the Appendix, multi-particle reconstruction can be ill-posed and/or yield spurious heterogeneity even in situations where the analogous single-particle reconstruction problem is well-posed. While we do not have a generic solution to this ill-posedness, we suggest applying the analysis above – namely calculating and inspecting the softest modes – for each of the volumes produced in a multi-particle reconstruction. These softest modes can reveal regions of high sensitivity which could indicate ill-posedness.

Finally, while our presentation above largely ignored image-specific translations $\{\delta_j\}$, one can easily incorporate these translations into the definitions of the likelihood and the corresponding Hessian (see section 12 in the Appendix). In future work we will expand our computational strategy to account for such translations, and plan to use strategies from [47] to accelerate this computation.

Part II

Appendix

In this part we introduce the notation used in the main text and describe the objects we are dealing with in the context of cryo-EM. We first discuss images, and then turn to volumes in section 9. Many of the conventions we establish for the former will carry over to the latter. When possible, we will use notation similar to [47, 48, 42].

5 Image notation

We use $\mathbf{x}, \mathbf{k} \in \mathbb{R}^2$ to represent spatial position and frequency, respectively. In polar coordinates these vectors are represented as:

$$\mathbf{x} = (x \cos \theta, x \sin \theta) \tag{7}$$

$$\mathbf{k} = (k \cos \psi, k \sin \psi). \tag{8}$$

The Fourier transform of a 2-dimensional function $A \in L^2(\mathbb{R}^2)$ is defined as

$$\widehat{A}(\mathbf{k}) := \iint_{\mathbb{R}^2} A(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} d\mathbf{x} . \tag{9}$$

We recover A from \widehat{A} using the inverse Fourier transform:

$$A(\mathbf{x}) = \frac{1}{(2\pi)^2} \iint_{\mathbb{R}^2} \widehat{A}(\mathbf{k}) e^{+i\mathbf{k} \cdot \mathbf{x}} d\mathbf{k} . \tag{10}$$

The inner product between two functions $A, B \in L^2(\mathbb{R}^2)$ is written as

$$\langle A, B \rangle = \iint_{\mathbb{R}^2} A(\mathbf{x})^\dagger B(\mathbf{x}) d\mathbf{x} , \tag{11}$$

where z^\dagger is the complex conjugate of $z \in \mathbb{C}$. We will also use Plancherel’s theorem [49],

$$\langle A, B \rangle = \frac{1}{(2\pi)^2} \langle \widehat{A}, \widehat{B} \rangle , \quad \forall A, B \in L^2(\mathbb{R}^2) , \tag{12}$$

ignoring the prefactors of 2π associated with the Fourier transform when they are not relevant.

We represent any given image as a function $A \in L^2(\mathbb{R}^2)$, with values corresponding to the image intensity at each location. Any inner product between A and B can be calculated equally well in either real or frequency space. Because of recent image alignment tools developed in the context of cryo-EM molecular reconstruction [50, 47], we will typically refer to images in frequency space (i.e., \widehat{A} , rather than A).

Abusing notation, we'll refer to $A(\mathbf{k})$ and $\widehat{A}(\mathbf{k})$ in polar coordinates as:

$$A(x, \theta) := A(\mathbf{x}) = A(x \cos \theta, x \sin \theta) \quad (13)$$

$$\widehat{A}(k, \psi) := \widehat{A}(\mathbf{k}) = \widehat{A}(k \cos \psi, k \sin \psi). \quad (14)$$

With this notation, each $\widehat{A}(k, \psi)$ for fixed k and $\psi \in [0, 2\pi)$ corresponds to a ‘ring’ in frequency space with radius k .

6 Rotating images: the Fourier–Bessel basis

Using the notation above, a rotation \mathcal{R}_γ by angle γ can be represented as:

$$\mathcal{R}_\gamma A(x, \theta) := A(x, \theta - \gamma). \quad (15)$$

Since rotation commutes with the Fourier transform, we have:

$$\widehat{\mathcal{R}_\gamma A}(k, \psi) = \mathcal{R}_\gamma \circ \widehat{A}(k, \psi) = \widehat{A}(k, \psi - \gamma). \quad (16)$$

In this manner, a rotation of any image by $+\gamma$ can be represented as an angular shift of each image ring by $\psi \rightarrow \psi - \gamma$.

The effect of a particular rotation can also be applied in a different basis: the Fourier–Bessel basis. To define the Fourier–Bessel coefficients of an image we recall that, for each fixed k , the image ring $\widehat{A}(k, \psi)$ is a 2π -periodic function of ψ . Thus, we can represent each image ring $\widehat{A}(k, \psi)$ as a Fourier series in ψ , obtaining

$$\widehat{A}(k, \psi) = \sum_{q=-\infty}^{+\infty} \check{A}(k, q) e^{+iq\psi}, \quad (17)$$

for $q \in \mathbb{Z}$. The Fourier–Bessel coefficients $\check{A}(k, q)$ of the image ring $\widehat{A}(k, \psi)$ are given by

$$\check{A}(k, q) = \frac{1}{2\pi} \int_0^{2\pi} \widehat{A}(k, \psi) e^{-iq\psi} d\psi. \quad (18)$$

These coefficients can be represented in a more traditional fashion by recalling that the Bessel function $J_q(kx)$ can be written as:

$$J_q(kx) = \frac{1}{2\pi} \int_0^{2\pi} e^{ikx \sin(\psi) - iq\psi} d\psi \quad (19)$$

$$= \frac{1}{2\pi} \int_0^{2\pi} e^{-ikx \cos(\psi + \pi/2) - iq\psi} d\psi, \quad (20)$$

which, when combined with the definition of the Fourier transform, immediately implies that:

$$\check{A}(k, q) = \iint A(x, \theta) \frac{1}{2\pi} \int_0^{2\pi} e^{-ikx \cos(\psi - \theta) - iq\psi} d\psi x dx d\theta \quad (21)$$

$$= \iint A(x, \theta) e^{-iq(\theta + \pi/2)} \frac{1}{2\pi} \int_0^{2\pi} e^{-ikx \cos(\psi + \pi/2) - iq\psi} d\psi x dx d\theta \quad (22)$$

$$= \iint A(x, \theta) e^{-iq(\theta + \pi/2)} J_q(kx) x dx d\theta, \quad (23)$$

which is the inner product between the original image (in real space) and a ‘Fourier–Bessel’ function.

The rotation \mathcal{R}_γ can now be represented as:

$$\mathcal{R}_\gamma \widehat{A}(k, \psi) = \widehat{A}(k, \psi - \gamma) \quad (24)$$

$$= \sum_{q=-\infty}^{+\infty} \check{A}(k, q) e^{+iq(\psi - \gamma)} \quad (25)$$

$$= \sum_{q=-\infty}^{+\infty} \check{A}(k, q) \cdot e^{-iq\gamma} e^{+iq\psi}, \quad (26)$$

such that the Fourier–Bessel coefficients of the rotated image ring $\mathcal{R}_\gamma \circ \hat{A}(k, \cdot)$ are given by the original Fourier–Bessel coefficients $\check{A}(k, q)$, each multiplied by the phase factor $e^{-iq\gamma}$. Note that (26) naturally allows for rotation by any γ , even values of γ that may not lie on the polar grid used to discretize the images.

In a typical discretization scheme (see below) the images \hat{A} and \hat{B} each require the storage of $\mathcal{O}(N^2)$ values, as do the Fourier–Bessel representations \check{A} and \check{B} . The calculation of the array $\hat{\mathcal{X}}(q)$ involves $\mathcal{O}(N^2)$ operations and $\mathcal{O}(N)$ storage. Once $\hat{\mathcal{X}}$ is calculated, the inner products $\mathcal{X}(\gamma)$ can be recovered on a uniform grid of $\mathcal{O}(N)$ angles $\gamma \in [0, 2\pi)$ with an additional $\mathcal{O}(N \log(N))$ operations using the FFT. More details are given in [50, 47, 42].

7 Image noise model

We assume that the images considered can be modeled as the sum of a ‘signal’ plus a ‘noise’. We assume that the signal corresponds to a 2-dimensional projection of a (smooth) molecular electron density function, while the noise corresponds to detector noise [31, 32, 35, 18]. We do not consider more complicated sources of noise, such as structural noise associated with image preparation (see [51]).

Based on these simple assumptions, we’ll model the noise in real space as independent and identically distributed (iid), with a variance of σ^2 on a unit scale in 2-dimensional real space. Thus, the image A evaluated at pixel (n_1, n_2) has a value modeled by:

$$A_{n_1, n_2} = A(\mathbf{x}_{n_1, n_2}) = A_{n_1, n_2}^{\text{signal}} + A_{n_1, n_2}^{\text{noise}}, \quad (27)$$

where the signal and noise are represented by the arrays $A_{n_1, n_2}^{\text{signal}}$ and $A_{n_1, n_2}^{\text{noise}}$, respectively. Because each entry in $A_{n_1, n_2}^{\text{noise}}$ corresponds to an average over the area element Δx^2 associated with a single pixel, we expect the variance of each $A_{n_1, n_2}^{\text{noise}}$ to scale inversely with Δx^2 . That is to say, we’ll assume that each element of the noise array $A_{n_1, n_2}^{\text{noise}}$ is drawn from the standard normal distribution $\mathcal{N}\left(0, \frac{\sigma^2}{\Delta x^2}\right)$, with a mean of 0 and a variance of $\sigma^2/\Delta x^2$.

Given the assumptions above, we expect that \hat{A} will be modeled by:

$$\hat{A}(\mathbf{k}) = \hat{A}(\mathbf{k}) = \hat{A}^{\text{signal}}(\mathbf{k}) + \hat{A}^{\text{noise}}(\mathbf{k}), \quad (28)$$

where \hat{A}^{noise} is now complex and iid. Because the noise term A^{noise} is real, the noise term \hat{A}^{noise} will be complex, with the conjugacy constraint that $\hat{A}^{\text{noise}}(+\mathbf{k}) = \hat{A}^{\text{noise}}(-\mathbf{k})^\dagger$.

If we were to sample \mathbf{k} on the uniform $N \times N$ cartesian grid associated with the standard 2-dimensional fast Fourier transform, then \mathbf{k}_{n_1, n_2} will correspond to the frequency $\Delta k \cdot (n_1, n_2)$, with frequency spacing $\Delta k = \pi$ and indices (n_1, n_2) considered periodically in the range $[-N/2, \dots, +N/2 - 1]$. In this case the transformation between A_{n_1, n_2} and \hat{A}_{n_1, n_2} will be unitary. When (n_1, n_2) corresponds to an index pair that is periodically reflected onto itself (i.e., the four index pairs $(0, 0)$, $(N/2, 0)$, $(0, N/2)$ and $(N/2, N/2)$) then $\hat{A}_{n_1, n_2}^{\text{noise}}$ will be real and drawn from $\mathcal{N}\left(0, \frac{\sigma^2}{\Delta x^2}\right)$. When (n_1, n_2) corresponds to an index pair that is not periodically reflected onto itself, the real and imaginary components of $\hat{A}^{\text{noise}}(\mathbf{k}_{n_1, n_2})$ will each be drawn iid from $\mathcal{N}\left(0, \frac{\sigma^2}{2\Delta x^2}\right)$ subject to the conjugacy constraint above.

Motivated by this observation, we define $\hat{\sigma}^2 = \frac{\pi^2 \sigma^2}{\Delta x^2}$ as the variance on a unit scale in 2-dimensional frequency space. We expect that the noise term $\hat{A}^{\text{noise}}(\mathbf{k})$ integrated over any area element Δk^2 in frequency space will have a variance of $\frac{\hat{\sigma}^2}{\Delta k^2}$.

8 Image-image similarity: likelihood

Given one ‘noisy’ image $\hat{A} = \hat{A}^{\text{signal}} + \hat{A}^{\text{noise}}$ and another ‘noiseless’ image $\hat{B} = \hat{B}^{\text{signal}}$, we can derive the standard likelihood of observing any value of \hat{A} , given that the signals \hat{A}^{signal} and \hat{B}^{signal} are the same:

$$P(\hat{A}^{\text{signal}} = \hat{B} \mid \hat{A}) \cdot P(\hat{A}) = P(\hat{A} \mid \hat{A}^{\text{signal}} = \hat{B}) \cdot P(\hat{A}^{\text{signal}} = \hat{B}), \quad (29)$$

which can be rearranged into:

$$P(\hat{A}^{\text{signal}} = \hat{B} \mid \hat{A}) = \frac{P(\hat{A} \mid \hat{A}^{\text{signal}} = \hat{B}) \cdot P(\hat{A}^{\text{signal}} = \hat{B})}{P(\hat{A})}. \quad (30)$$

Assuming a uniform prior $P(\hat{A}^{\text{signal}} = \hat{B})$, the posterior probability $P(\hat{A}^{\text{signal}} = \hat{B} \mid \hat{A})$ is proportional to the likelihood $P(\hat{A} \mid \hat{A}^{\text{signal}} = \hat{B})$.

$$P(\hat{A}(\mathbf{k}) \mid \hat{B}(\mathbf{k})) := P(\hat{A}(\mathbf{k}) \mid \hat{A}^{\text{signal}}(\mathbf{k}) = \hat{B}(\mathbf{k})) = \frac{1}{(\sqrt{2\pi} \hat{\sigma})^{\text{dof}}} \exp\left(-\frac{1}{2\hat{\sigma}^2} \iint_{\mathbf{k} \in \Omega_K} |\hat{A}(\mathbf{k}) - \hat{B}(\mathbf{k})|^2 d\mathbf{k}\right), \quad (31)$$

where dof refers to the number of independent degrees of freedom in the image (e.g., dof could be the number of pixels in the image $A(\mathbf{x})$, or the number of quadrature-nodes used to discretize $\widehat{A}(k, \psi)$).

When dealing with two noisy images $\widehat{A} = \widehat{A}^{\text{signal}} + \widehat{A}^{\text{noise}}$ and $\widehat{B} = \widehat{B}^{\text{signal}} + \widehat{B}^{\text{noise}}$, as long as $\widehat{A}^{\text{noise}}$ is independent from $\widehat{B}^{\text{noise}}$ the same general argument applies. The only difference in this scenario is that the effective variance involves both the variance of $\widehat{A}^{\text{noise}}$ and $\widehat{B}^{\text{noise}}$.

9 Volume notation

We use $\boldsymbol{\chi}, \boldsymbol{\kappa} \in \mathbb{R}^3$ to represent spatial position and frequency, respectively. in spherical coordinates the vector $\boldsymbol{\kappa}$ is represented as:

$$\boldsymbol{\kappa} = k \cdot \widehat{\boldsymbol{\kappa}}, \quad \text{with } \widehat{\boldsymbol{\kappa}} = (\cos \phi \sin \theta, \sin \phi \sin \theta, \cos \theta), \quad (32)$$

with polar angle θ and azimuthal angle ϕ representing the unit vector $\widehat{\boldsymbol{\kappa}}$ on the surface of the sphere S^2 .

Using a right handed basis, a rotation about the third axis by angle α is represented as:

$$\mathcal{R}_\alpha^z = \begin{pmatrix} +\cos \alpha & -\sin \alpha & 0 \\ +\sin \alpha & +\cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (33)$$

and a rotation about the second axis by angle β is represented as:

$$\mathcal{R}_\beta^y = \begin{pmatrix} +\cos \beta & 0 & +\sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & +\cos \beta \end{pmatrix}. \quad (34)$$

A rotation \mathcal{R}_τ of a vector $\boldsymbol{\kappa} \in \mathbb{R}^3$ can be represented by the vector of Euler angles $\tau = (\gamma, \beta, \alpha)$:

$$\mathcal{R}_\tau \cdot \boldsymbol{\kappa} = \mathcal{R}_\alpha^z \circ \mathcal{R}_\beta^y \circ \mathcal{R}_\gamma^z \cdot \boldsymbol{\kappa}. \quad (35)$$

We represent any given volume as a function $F \in L^2(\mathbb{R}^3)$, with values corresponding to the volume intensity at each location $\boldsymbol{\chi} \in \Omega_1$. We'll refer to $\widehat{F}(\boldsymbol{\kappa})$ in spherical coordinates as $\widehat{F}(k, \widehat{\boldsymbol{\kappa}})$. With this notation, each $\widehat{F}(k, \cdot)$ corresponds to a 'shell' in frequency space with radius k . The rotation of any volume $\mathcal{R}_\tau \widehat{F}(k, \widehat{\boldsymbol{\kappa}})$ corresponds to the function $\widehat{F}(k, \mathcal{R}_\tau^{-1} \widehat{\boldsymbol{\kappa}})$.

10 Rotation in the spherical-harmonic basis

Using the notation above, we can represent a volume $\widehat{F}(k, \widehat{\boldsymbol{\kappa}})$ as:

$$\widehat{F}(k, \widehat{\boldsymbol{\kappa}}) = \sum_{l=0}^{+\infty} \sum_{m=-l}^{m=+l} \widehat{F}_l^m(k) Y_l^m(\widehat{\boldsymbol{\kappa}}), \quad (36)$$

where $Y_l^m(\widehat{\boldsymbol{\kappa}})$ represents the spherical-harmonic of degree l and order m :

$$Y_l^m(\theta, \phi) = Z_l^m e^{+im\phi} P_l^m(\cos \theta), \quad (37)$$

with P_l^m representing the (unnormalized) associated Legendre polynomial, and Z_l^m the normalization constant:

$$Z_l^m = \sqrt{\frac{2l+1}{4\pi} \times \frac{(l-|m|)!}{(l+|m|)!}}. \quad (38)$$

The coefficients $\widehat{F}_l^m(k)$ define the spherical-harmonic expansion of the k -shell $\widehat{F}(k, \cdot)$.

Using this spherical-harmonic basis allows us to efficiently apply rotations. For example, given a rotation by $\tau = (\gamma, \beta, \alpha)$, we can represent the rotated volume $\widehat{G} := \mathcal{R}_\tau \widehat{F}$ as:

$$\widehat{G}_l^{m_1} = \sum_{m_2=-l}^{m_2=+l} e^{-im_1\alpha} d_{m_1, m_2}^l(\beta) e^{-im_2\gamma} \widehat{F}_l^{m_2}, \quad (39)$$

where $d_{m_1, m_2}^l(\beta)$ represents the degree l Wigner d-matrix associated with the interior Euler angle β .

11 Volume projection in the spherical-harmonic basis

We denote the projection of a volume $F(\boldsymbol{\chi})$ (with $\boldsymbol{\chi} \in \mathbb{R}^3$) along a particular viewing-angle τ by the function $S(\boldsymbol{x}; \tau; F)$ (with $\boldsymbol{x} \in \mathbb{R}^2$). We'll refer to this projection $S(\boldsymbol{x}; \tau; F)$ as the 'template' associated with $F(\boldsymbol{\chi})$ and viewing-angle τ .

By the Fourier-slice theorem we can describe the Fourier-transformed template $\widehat{S}(\mathbf{k}; \tau, \widehat{F})$ via:

$$\widehat{S}(\mathbf{k}; \tau; \widehat{F}) = \left[\widehat{\text{Sl}} \circ \mathcal{R}(\tau) \circ \widehat{F} \right] (\mathbf{k}), \quad (40)$$

where $\widehat{\text{Sl}}$ denotes taking the equatorial slice in the Fourier domain.

In terms of computation, the $\widehat{\text{Sl}}$ operator can be described by restricting the unit-vector $\widehat{\boldsymbol{\kappa}}$ in Eq. 32 to the equatorial-plane, corresponding to setting $\beta \equiv 0$ and associating $\alpha =: \psi$.

$$\left[\widehat{\text{Sl}} \circ \widehat{F} \right] (k, \psi) := \sum_{l=0}^{+\infty} \sum_{m=-l}^{m=+l} \widehat{F}_l^m(k) Z_l^m e^{+im\psi} P_l^m(0). \quad (41)$$

If we consider the effect of a contrast-transfer-function $CTF(\mathbf{k})$, then the template associated with $F(\boldsymbol{\chi})$ and a viewing-angle τ can be written as:

$$\widehat{S}(\mathbf{k}; \tau; CTF; \widehat{F}) = CTF(\mathbf{k}) \odot \left[\widehat{\text{Sl}} \circ \mathcal{R}(\tau) \circ \widehat{F} \right] (\mathbf{k}). \quad (42)$$

If we were also to consider the effect of an image-dependent displacement $\boldsymbol{\delta}$, then the template associated with $F(\boldsymbol{\chi})$ and a viewing-angle τ can be written as:

$$\widehat{S}(\mathbf{k}; \tau; \boldsymbol{\delta}; CTF; \widehat{F}) = CTF(\mathbf{k}) \odot \widehat{\mathcal{T}}(\boldsymbol{\delta}; \mathbf{k}) \odot \left[\widehat{\text{Sl}} \circ \mathcal{R}(\tau) \circ \widehat{F} \right] (\mathbf{k}), \quad (43)$$

where $\widehat{\mathcal{T}}(\boldsymbol{\delta}; \mathbf{k})$ is the fourier-signature corresponding to a 2-dimensional translation by $\boldsymbol{\delta}$:

$$\widehat{\mathcal{T}}(\boldsymbol{\delta}; \mathbf{k}) := \exp(-i\boldsymbol{\delta} \cdot \mathbf{k}). \quad (44)$$

12 Image-volume similarity: likelihood

12.1 single-particle:

Referring back to section 8 and Eq. 31, we can express the likelihood that a particular volume $\widehat{F}(\mathbf{k})$ might produce a particular image $\widehat{A}(\mathbf{k})$ with CTF-function $CTF(\mathbf{k})$ via projection along viewing-angle τ and translation by a displacement-vector $\boldsymbol{\delta}$:

$$P(\widehat{A} | \tau; \boldsymbol{\delta}; \widehat{F}) = P(\widehat{A} | \widehat{S}(\mathbf{k}; \tau; \boldsymbol{\delta}; CTF; \widehat{F})). \quad (45)$$

The likelihood $P(\widehat{A} | \tau; \boldsymbol{\delta}; \widehat{F})$ can be marginalized across τ and $\boldsymbol{\delta}$ to calculate the likelihood that \widehat{F} might produce the image $\widehat{A}(\mathbf{k})$:

$$P(\widehat{A} | \widehat{F}) = \int_{\boldsymbol{\delta} \in \mathbb{R}^2} \int_{\tau \in SO3} P(\widehat{A} | \tau; \boldsymbol{\delta}; \widehat{F}) \cdot P(\tau; \boldsymbol{\delta} | \widehat{F}) d\tau d\boldsymbol{\delta} \quad (46)$$

$$= \int_{\boldsymbol{\delta} \in \mathbb{R}^2} \int_{\tau \in SO3} P(\widehat{A} | \tau; \boldsymbol{\delta}; \widehat{F}) \cdot P(\boldsymbol{\delta} | \tau; \widehat{F}) \cdot P(\tau | \widehat{F}) d\tau d\boldsymbol{\delta}, \quad (47)$$

where $P(\tau | \widehat{F})$ is prior distribution attributed to the viewing-angle τ , $P(\boldsymbol{\delta} | \tau; \widehat{F})$ is the prior distribution attributed to $\boldsymbol{\delta}$ for a particular value of τ , and $P(\tau; \boldsymbol{\delta} | \widehat{F})$ is the joint-prior over the alignments τ and $\boldsymbol{\delta}$. To emphasize the dependence of $P(\widehat{A} | \widehat{F})$ on the prior distributions and noise-level, we will sometimes write:

$$P(\widehat{A} | \widehat{F}) =: P(\widehat{A} | \widehat{F}; P(\tau; \boldsymbol{\delta} | \widehat{F}); \hat{\sigma}).$$

In this manuscript we will typically assume that $\delta \equiv 0$. This corresponds to taking the limit $\sigma_{\boldsymbol{\delta}} \rightarrow 0$, resulting in $P(\boldsymbol{\delta} | \tau; \widehat{F})$ becoming sharply peaked at the origin. Under this limit we'll suppress the notation for $\boldsymbol{\delta}$. More generally, we can consider $P(\boldsymbol{\delta} | \tau; \widehat{F})$ to be an isotropic gaussian centered at the origin:

$$P(\boldsymbol{\delta} | \tau; \widehat{F}) = P(\boldsymbol{\delta}) = \frac{1}{2\pi\sigma_{\boldsymbol{\delta}}^2} \exp\left(-\frac{|\boldsymbol{\delta}|^2}{2\sigma_{\boldsymbol{\delta}}^2}\right),$$

where σ_{δ} is a fixed small distance (e.g., one tenth the diameter of each picture).

Given a pool of J images $\{\hat{A}_j\}$ and their associated CTF-functions $\{CTF_j\}$ (both indexed by $j \in 1, \dots, J$), we can accumulate terms from Eq. 47 to calculate the likelihood that \hat{F}_i might produce this collection of images:

$$P(\{\hat{A}_j\} | \hat{F}) = P(\{\hat{A}_j\} | \hat{F}; P(\tau; \delta | \hat{F}); \hat{\sigma}) = \prod_{j=1}^J P(\hat{A}_j | \hat{F}; P(\tau; \delta | \hat{F}); \hat{\sigma}). \quad (48)$$

12.2 multi-particle:

In the context of multi-particle reconstruction, we might consider I different volumes $\{\hat{F}_i\}$ (indexed by $i \in 1, \dots, I$), each with their own prior distributions. In this context we can calculate the likelihood that the collection $\{\hat{F}_i\}$ might produce the image $\hat{A}(\mathbf{k})$:

$$P(\hat{A} | \{\hat{F}_i\}) = \sum_{i=1}^I P(\hat{A} | \hat{F}_i) \cdot P(\hat{F}_i) \quad (49)$$

$$= \int_{\delta \in \mathbb{R}^2} \int_{\tau \in SO3} \sum_{i=1}^I P(\hat{A} | \tau; \delta; \hat{F}_i) \cdot P(\tau; \delta | \hat{F}_i) \cdot P(\hat{F}_i) d\tau d\delta \quad (50)$$

$$= \int_{\delta \in \mathbb{R}^2} \int_{\tau \in SO3} \sum_{i=1}^I P(\hat{A} | \tau; \delta; \hat{F}_i) \cdot P(\tau, \delta, i | \{\hat{F}_i\}) d\tau d\delta, \quad (51)$$

where $P(\hat{F}_i)$ is the prior probability that the image \hat{A} was drawn from the volume \hat{F}_i , and $P(\tau, \delta, i | \{\hat{F}_i\})$ is the prior probability that the image \hat{A} was drawn from the volume \hat{F}_i with viewing-angle τ and displacement δ . Once again, to emphasize the dependence of $P(\hat{A} | \{\hat{F}_i\})$ on the prior distributions and noise-level, we will sometimes write:

$$P(\hat{A} | \{\hat{F}_i\}) = P(\hat{A} | \{\hat{F}_i\}; P(\tau, \delta, i | \{\hat{F}_i\}); \hat{\sigma}).$$

Just as in Eq. 48, we can accumulate terms from Eq. 51 to calculate the likelihood that $\{\hat{F}_i\}$ might produce a collection of images:

$$P(\{\hat{A}_j\} | \{\hat{F}_i\}) = P(\{\hat{A}_j\} | \{\hat{F}_i\}; P(\tau, \delta, i | \{\hat{F}_i\}); \hat{\sigma}) = \prod_{j=1}^J P(\hat{A}_j | \{\hat{F}_i\}; P(\tau, \delta, i | \{\hat{F}_i\}); \hat{\sigma}). \quad (52)$$

13 Image-volume similarity in the low-temperature limit

13.1 single-particle:

The likelihood in Eq. 48 depends via Eqs. 47, 45 and 31 on the estimated noise-level $\hat{\sigma}^2$, which is related to the signal-to-noise ratio, or ‘snr’, and is sometimes referred to as the estimated ‘temperature’ used in the bayesian framework. In the low-temperature limit that $\hat{\sigma}^2 \rightarrow 0$, the integral in Eq. 47 becomes dominated by the contribution associated with the ‘maximum-likelihood’ alignment-angle for the image \hat{A} . This maximum-likelihood alignment occurs for the viewing-angle, displacement and volume which maximize the summand $P(\hat{A} | \tau; \delta; \hat{F}) \cdot P(\tau, \delta | \hat{F})$ in the integrand of Eq. 47. More specifically, given a particular volume \hat{F} , the maximum-likelihood alignment $\{\tau^{\text{opt}}, \delta^{\text{opt}}\}$ for an image \hat{A} occurs at:

$$\{\tau^{\text{opt}}, \delta^{\text{opt}}\} := \arg \max_{\{\tau, \delta\}} P(\hat{A} | \tau; \delta; \hat{F}), \quad (53)$$

where τ^{opt} and δ^{opt} are restricted to lie in the support of $P(\tau; \delta | \hat{F})$. With this notation, each of the $\{\tau^{\text{opt}}, \delta^{\text{opt}}, i^{\text{opt}}\}$ can be thought of as functions of the image \hat{A} , the volume \hat{F} and the prior distribution $P(\tau, \delta | \hat{F})$.

In the context of single-particle reconstruction, the negative-log-likelihood in the low-temperature limit is dominated by the following leading-order expansion:

$$-\log(P(\{\hat{A}_j\} | \hat{F})) \sim \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbf{k} \in \Omega_K} \left| \hat{A}_j(\mathbf{k}) - \hat{S}(\mathbf{k}; \tau_j^{\text{opt}}; \delta_j^{\text{opt}}; CTF_j; \hat{F}) \right|^2 d\mathbf{k}. \quad (54)$$

This ‘template-wise’ expression represents the likelihood in terms of a collection of 2-dimensional integrals. Later on below we’ll contrast this with an equivalent ‘volumetric’ expression which represents the likelihood in terms of a collection of 3-dimensional integrals.

13.2 multi-particle:

In the context of multi-particle reconstruction, the likelihood in Eq. 52 depends on Eq. 51. In the low-temperature limit that $\hat{\sigma}^2 \rightarrow 0$, the maximum-likelihood alignment occurs for the viewing-angle, displacement and volume-label which maximize the summand $P(\hat{A} | \tau; \boldsymbol{\delta}; \hat{F}_i) \cdot P(\tau, \boldsymbol{\delta}, i | \{\hat{F}_i\})$ in the integrand of Eq. 51.

Given a collection of volumes $\{\hat{F}_i\}$, the maximum-likelihood alignment $\{\tau^{\text{opt}}, \boldsymbol{\delta}^{\text{opt}}\}$ and volume-label i^{opt} for an image \hat{A} occurs at:

$$\{\tau^{\text{opt}}, \boldsymbol{\delta}^{\text{opt}}, i^{\text{opt}}\} := \arg \max_{\{\tau, \boldsymbol{\delta}, i\}} P(\hat{A} | \tau; \boldsymbol{\delta}; \hat{F}_i), \quad (55)$$

where τ^{opt} , $\boldsymbol{\delta}^{\text{opt}}$ and i^{opt} are restricted to lie in the support of $P(\tau, \boldsymbol{\delta}, i | \{\hat{F}_i\})$. With this notation, each of the $\{\tau^{\text{opt}}, \boldsymbol{\delta}^{\text{opt}}, i^{\text{opt}}\}$ can be thought of as functions of the image \hat{A} , the volume-set $\{\hat{F}_i\}$ and the prior distribution $P(\tau, \boldsymbol{\delta}, i | \{\hat{F}_i\})$.

Thus, within the context of multi-particle reconstruction, the negative-log-likelihood in the low-temperature limit takes a form very similar to Eq. 54:

$$-\log(P(\{\hat{A}_j\} | \{\hat{F}_i\})) \sim \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbf{k} \in \Omega_K} \left| \hat{A}_j(\mathbf{k}) - \hat{S}(\mathbf{k}; \tau_j^{\text{opt}}; \boldsymbol{\delta}_j^{\text{opt}}; CTF_j; \hat{F}_{i_j^{\text{opt}}}) \right|^2 d\mathbf{k}. \quad (56)$$

14 Volume reconstruction in the low-temperature limit via least-squares

In the low-temperature limit each image \mathbf{A}_j is assigned to a particular viewing-angle τ_j , displacement $\boldsymbol{\delta}_j$ and volume-label i via Eq. 55. Given these maximum-likelihood assignments, each volume can be reconstructed independently. For example, in the context of single-particle-reconstruction, the single volume \hat{F} can be reconstructed by solving the following linear-system in a least-squares sense:

$$\begin{bmatrix} CTF_1 \circ \hat{\mathcal{T}}(+\boldsymbol{\delta}_1) \circ \{\hat{\mathcal{S}} \circ \mathcal{R}(\tau_1) \circ\} \\ \vdots \\ CTF_J \circ \hat{\mathcal{T}}(+\boldsymbol{\delta}_J) \circ \{\hat{\mathcal{S}} \circ \mathcal{R}(\tau_J) \circ\} \end{bmatrix} \cdot \hat{F}^{\text{lsq}} \approx \begin{bmatrix} \hat{A}_1 \\ \vdots \\ \hat{A}_J \end{bmatrix}. \quad (57)$$

In this expression, the J viewing angles τ_j and displacements $\boldsymbol{\delta}_j$ are assumed to be given as input, as are the image- or micrograph-specific CTF functions CTF_j and images \hat{A}_j . The unknown (to be determined) is the volume \hat{F}^{lsq} . By a least-squares solution to Eq. 57 we specifically mean the volume \hat{F}^{lsq} for which the residual is proportional to the negative-log-likelihood described in Eq. 54.

If it is more convenient, we can center the images on the right-hand-side and rewrite Eq. 57 as:

$$\begin{bmatrix} CTF_1 \circ \{\hat{\mathcal{S}} \circ \mathcal{R}(\tau_1) \circ\} \\ \vdots \\ CTF_J \circ \{\hat{\mathcal{S}} \circ \mathcal{R}(\tau_J) \circ\} \end{bmatrix} \cdot \hat{F}^{\text{lsq}} \approx \begin{bmatrix} \hat{\mathcal{T}}(-\boldsymbol{\delta}_1) \circ \hat{A}_1 \\ \vdots \\ \hat{\mathcal{T}}(-\boldsymbol{\delta}_J) \circ \hat{A}_J \end{bmatrix}. \quad (58)$$

In the context of multi-particle reconstruction the pool of images can be sorted by volume-label, and then the same strategy can be applied for each volume independently.

15 Volume reconstruction in the low-temperature limit via back-propagation

When the number of images J becomes large, the solution to Eq. 58 can be well approximated via a form of back-propagation [42]. To motivate this strategy one can rewrite the likelihood as a volumetric-integral (see section 17). Alternatively, one may observe that, if there were no noise and no CTF-attenuation, then the volume and images would solve:

$$\begin{bmatrix} \hat{\mathcal{S}} \circ \mathcal{R}(\tau_1) \circ \\ \vdots \\ \hat{\mathcal{S}} \circ \mathcal{R}(\tau_J) \circ \end{bmatrix} \cdot \hat{F} = \begin{bmatrix} \hat{\mathcal{T}}(-\boldsymbol{\delta}_1, \mathbf{k}) \circ \hat{A}_1 \\ \vdots \\ \hat{\mathcal{T}}(-\boldsymbol{\delta}_J, \mathbf{k}) \circ \hat{A}_J \end{bmatrix}$$

with each template $\hat{\mathcal{S}} \circ \mathcal{R}(\tau_j) \circ \hat{F}$ corresponding to a cross-section of the volume $\hat{F}(\boldsymbol{\kappa})$, with $\boldsymbol{\kappa} \in \mathbb{R}^3$. Each of these cross-sections can be viewed as a sample of the overall volume $\hat{F}(\boldsymbol{\kappa})$, with values given by the appropriate entry on the

right-hand-side of the equation above. With sufficiently many cross-sectional samples, a functional representation of the volume $\widehat{F}(\boldsymbol{\kappa})$ for any point $\boldsymbol{\kappa} \in \mathbb{R}^3$ can be reconstructed simply by collecting and averaging the images that correspond to that point. That is, for each $\boldsymbol{\kappa} \in \mathbb{R}^3$ we have the formula:

$$\widehat{F}^{\text{bp}}(\boldsymbol{\kappa}) := \frac{1}{|\mathcal{J}(\boldsymbol{\kappa})|} \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} \widehat{\mathcal{T}}(-\boldsymbol{\delta}_{j'}, \mathbf{k}) \cdot \widehat{A}_{j'}(\mathbf{k}), \text{ where: } \mathbf{k} := \mathbf{k}(\boldsymbol{\kappa}; j') = \mathcal{R}(\tau_{j'}) \circ \boldsymbol{\kappa}. \quad (59)$$

In this formula the set $\mathcal{J}(\boldsymbol{\kappa})$ includes only those image-indices j' for which $\widehat{A}_{j'}$ contributes to the point $\boldsymbol{\kappa}$. That is to say, for each $\widehat{A}_{j'}$ there must exist a point $\mathbf{k} \in \mathbb{R}^2$ corresponding to a location in the equatorial-plane in \mathbb{R}^3 onto which the point $\boldsymbol{\kappa} \in \mathbb{R}^3$ is rotated by $\mathcal{R}(\tau_{j'})$. This point $\mathbf{k} \in \mathbb{R}^2$ will implicitly depend on $\boldsymbol{\kappa}$ as well as the chosen image-index j' . The normalizing factor $1/|\mathcal{J}(\boldsymbol{\kappa})|$ ensures that $\widehat{F}^{\text{bp}}(\boldsymbol{\kappa})$ equals the average of all the image-values associated with $\boldsymbol{\kappa}$.

The above formula can easily be generalized to deal with image-specific CTF-values. Indeed, by interpreting each image-value $\widehat{A}_{j'}(\mathbf{k})$ as a noisy sample of $\widehat{F}^{\text{bp}}(\boldsymbol{\kappa})$, we can think of Eq. 59 as a local maximum-likelihood estimate of $\widehat{F}^{\text{bp}}(\boldsymbol{\kappa})$. More specifically, one can imagine sampling the image-values a_j associated with a particular $\boldsymbol{\kappa} \in \mathbb{R}^3$ from the corresponding volume-value f :

$$a_j = f + \epsilon_j, \quad \text{with each } \epsilon_j \text{ drawn independently from } \mathcal{N}(0, \hat{\sigma}^2).$$

Given a collection of J such observations, the local-maximum-likelihood estimate for f is:

$$f^{\text{opt}} = \frac{1}{J} \sum_{j=1}^J a_j.$$

If each image-value a_j is weighted by an image-specific CTF-value c_j , then we can adjust our model and local maximum-likelihood estimate accordingly:

$$a_j = f \cdot c_j + \epsilon_j \quad \implies \quad f^{\text{opt}} = \frac{\sum_{j=1}^J c_j a_j}{\sum_{j=1}^J c_j^2}.$$

Applying this logic to Eq. 59 yields:

$$\tilde{J}(\boldsymbol{\kappa}) := \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} |\text{CTF}_{j'}(\mathbf{k})|^2, \text{ and} \quad (60)$$

$$\widehat{F}^{\text{bp}}(\boldsymbol{\kappa}) := \frac{1}{\tilde{J}(\boldsymbol{\kappa})} \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} \text{CTF}_{j'}(\mathbf{k}) \cdot \widehat{\mathcal{T}}(-\boldsymbol{\delta}_{j'}, \mathbf{k}) \cdot \widehat{A}_{j'}(\mathbf{k}), \quad (61)$$

where once again the $\mathbf{k} \in \mathbb{R}^2$ is implicitly related to $\boldsymbol{\kappa}$ and j' . This strategy coincides with the CTF-weighting of Bayesian inference described in [16, 33].

The term $\tilde{J}(\boldsymbol{\kappa})$ can be interpreted as the accumulated evidence at each point $\boldsymbol{\kappa}$ in frequency-space. Similarly, the term $\tilde{J}(\boldsymbol{\kappa}) \cdot \widehat{F}^{\text{bp}}(\boldsymbol{\kappa})$ can be thought of as the accumulated image-data at $\boldsymbol{\kappa}$. The quotient of these two terms (i.e., $\widehat{F}^{\text{bp}}(\boldsymbol{\kappa})$) can be thought of as a first-moment. Later on below we'll see a similarly constructed second-moment within the likelihood.

16 Numerical details for back-propagation

Our algorithm attempts to solve the back-propagation problem exactly up to a given spherical-harmonic degree L (which is itself related to the maximum frequency K determining the resolution).

To describe our approach in more detail, let's consider the following summation $\Sigma(M)$, originally represented as a 2-dimensional 'template-wise' integral:

$$\Sigma(M) = \sum_j \int_{\mathbf{k} \in \Omega(K) \subset \mathbb{R}^2} M_j^\dagger(\mathbf{k}) \cdot [\widehat{\mathcal{S}} \circ \mathcal{R}(\tau_j) \circ \widehat{F}] (\mathbf{k}) d\mathbf{k}, \quad (62)$$

for some collection of data $\{M_j(\mathbf{k})\}$.

Let's first consider the case where $J = 1$, and we have only a single nonzero data-value of $M = 1$ associated with the viewing-angle τ at a point \mathbf{k} for which $\boldsymbol{\kappa} = \mathcal{R}^{-1}(\tau) \circ \mathbf{k}$ is positioned at the north-pole $\boldsymbol{\kappa}_0 = [0, 0, 1]^\top$. Instead of calculating the sum $\Sigma(M)$ by integrating M^\dagger against the template $\widehat{\mathcal{S}} \circ \mathcal{R}(\tau) \circ \widehat{F}$ in 2-dimensions, we can instead just integrate a delta-function 'monopole impulse' concentrated at the north-pole against the original volume \widehat{F} in 3-dimensions. Put

another way, our single nonzero data-value can be associated with the back-propagated volume $\delta(\boldsymbol{\kappa} - \boldsymbol{\kappa}_0)$, again referring to a delta-function impulse concentrated at the north-pole $\boldsymbol{\kappa}_0$.

When restricted to the spherical-shell of radius 1 and projected onto the spherical-harmonics of degree $l \leq L$, the ideal impulse $\delta(\boldsymbol{\kappa} - \boldsymbol{\kappa}_0)$ (itself fully concentrated at $\boldsymbol{\kappa}_0$) produces a band-limited impulse which is a spherical-sinc that is now centered around (but no longer fully concentrated at) its source at the north-pole:

$$\tilde{\delta}(\hat{\boldsymbol{\kappa}}) = \sum_{l=0}^L \sqrt{4\pi} \cdot \sqrt{1+2l} \cdot P_l(\cos(\beta)) = \sum_{l=0}^L 4\pi \cdot Y_l^0(\hat{\boldsymbol{\kappa}}),$$

where P_l is a legendre-polynomial of degree l and, by construction, the spherical-harmonic coefficients $\tilde{\delta}_l^m$ are given by 4π when $m = 0$.

Similarly, a single data-value of $M = 1$ positioned at the equatorial-point $\boldsymbol{\kappa}_1 = [\cos(\psi), \sin(\psi), 0]^\top$ will (given our conventions) produce the back-propagated impulse $[\mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta}](\hat{\boldsymbol{\kappa}})$ (i.e., a spherical-sinc with a source positioned at the appropriate point on the equator).

Due to linearity, we can associate the more general data-pool $\{M_j(\mathbf{k})\}$ with the back-propagated shell (for any radius k):

$$\tilde{M}(\boldsymbol{\kappa}) = \sum_j \int_{\psi \in [0, 2\pi]} M_j(\mathbf{k}) \cdot [\mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta}](\hat{\boldsymbol{\kappa}}) d\psi, \quad (63)$$

which specifies $\tilde{M}(\boldsymbol{\kappa})$ at magnitude k . In this expression we implicitly relate $\boldsymbol{\kappa} \in \Omega(K) \subset \mathbb{R}^3$ to the amplitude and unit-vector $(k, \hat{\boldsymbol{\kappa}})$. Similarly, we relate $\mathbf{k} \in \Omega(K) \subset \mathbb{R}^2$ to the amplitude and angle (k, ψ) . Finally, we relate $\mathbf{k} = \mathcal{R}(\tau_j) \circ \boldsymbol{\kappa}$. In words, this simply amounts to walking through each piece of data M_j and, for each point $\mathbf{k} \in \mathbb{R}^2$, placing an impulse of the appropriate magnitude on the associated point $\boldsymbol{\kappa} = \mathcal{R}^{-1}(\tau_j) \circ \mathbf{k}$ in 3-dimensional space.

With this reorganization we can now represent the sum $\Sigma(M)$ as a 3-dimensional ‘volumetric’ integral:

$$\Sigma(M) = \int_{\boldsymbol{\kappa} \in \Omega(K) \subset \mathbb{R}^3} \tilde{M}^\dagger(\boldsymbol{\kappa}) \cdot \hat{F}(\boldsymbol{\kappa}) \cdot \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|}. \quad (64)$$

To compute $\tilde{M}(\boldsymbol{\kappa})$ efficiently, we employ the following strategy:

1. First, we apply a linear-transformation \tilde{D} to mollify the impulse $\tilde{\delta}$ and ensure that the support of $\tilde{D} \circ \tilde{\delta}(\hat{\boldsymbol{\kappa}})$ is not too large on the surface of the shell.
2. Alongside \tilde{D} , we construct a quadrature-mesh on the surface of the shell which can both (i) resolve the mollified impulse $\tilde{D} \circ \tilde{\delta}$ and (ii) accurately integrate the spherical-harmonics Y_l^m up to degree L .
3. Now, for each j and q , the image-value $\hat{A}_j(k, \psi_q)$ is associated with the compactly-supported mollified impulse $\hat{A}_j(k, \psi_q) \cdot \mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi_q]) \circ \tilde{\delta}(\hat{\boldsymbol{\kappa}})$, which is simply a spherical-sinc with a source at $\mathcal{R}^{-1}(\tau_j) \circ [\cos(\psi), \sin(\psi), 0]^\top$ which we sample on the quadrature-mesh.
4. After accumulating all such values for all j and q on the quadrature-mesh, we use quadrature to recover the spherical-harmonic coefficients of $\tilde{D}\tilde{M}$.
5. Finally, we ‘undo’ the effect of \tilde{D} and calculate $\tilde{M} := [\tilde{D}]^{-1} \circ [\tilde{D}\tilde{M}]$.

This recipe provides a template to calculate \tilde{J} and $\tilde{J} \cdot \hat{F}^{\text{bp}}$:

$$\tilde{J}(\boldsymbol{\kappa}) = \sum_j \int_{\psi \in [0, 2\pi]} |CTF_j(\mathbf{k})|^2 \cdot [\mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta}](\hat{\boldsymbol{\kappa}}), \quad (65)$$

$$[\tilde{J} \cdot \hat{F}^{\text{bp}}](\boldsymbol{\kappa}) = \sum_j \int_{\psi \in [0, 2\pi]} CTF_j(\mathbf{k}) \cdot \hat{T}(-\delta_j, \mathbf{k}) \cdot \hat{A}_j(\mathbf{k}) \cdot [\mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta}](\hat{\boldsymbol{\kappa}}), \quad (66)$$

where once again $\boldsymbol{\kappa} = (k, \hat{\boldsymbol{\kappa}})$, $\mathbf{k} = (k, \psi)$ and $\mathbf{k} = \mathcal{R}(\tau_j) \circ \boldsymbol{\kappa}$.

16.1 designing \tilde{D} :

In this recipe we have the luxury of designing the mollifier \tilde{D} . We choose \tilde{D} so that:

1. \tilde{D} is diagonalized by the spherical-harmonic basis. That is to say: the spherical-harmonic coefficients $[\tilde{D} \circ \tilde{\delta}]_l$ are given by $\tilde{D}_l \cdot \tilde{\delta}_l$, for a collection of coefficients \tilde{D}_l . This allows for \tilde{D} and \tilde{D}^{-1} to be quickly applied in the spherical-harmonic basis.

2. The majority of the amplitude of $\tilde{D} \circ \tilde{\delta}$ is close to the north- and south-pole. Put another way, we would like:

$$I^2(\beta_{\text{north}}, \beta_{\text{south}}; \tilde{D}) := \int_{\beta \in [\beta_{\text{north}}, \beta_{\text{south}}]} \int_{\alpha \in [0, 2\pi]} \left| \tilde{D} \circ \tilde{\delta}(\beta, \alpha) \right|^2 d\alpha \sin(\beta) d\beta \leq \epsilon_{\text{pole}}^2,$$

for parameters including cutoff polar-cap angles β_{north} and β_{south} and an error-tolerance of ϵ_{pole}^2 , where β_{north} and β_{south} are as close to 0 and π as possible, while ϵ_{pole}^2 is as small as possible. This constraint allows $\tilde{D} \circ \tilde{\delta}$ to be compactly supported on each spherical shell.

3. The coefficients \tilde{D}_l are not too small for $l \leq L$. That is, we would like:

$$\min(\tilde{D}_l) \geq \tilde{D}_{\text{min}} > 0,$$

for some lower-bound \tilde{D}_{min} . This constraint bounds the operator-norm of \tilde{D}^{-1} , allowing for both \tilde{D} and \tilde{D}^{-1} to be well-conditioned.

Note that the latter two criteria for \tilde{D} are in conflict with one another: compressing the support of the mollified impulse $\tilde{D} \circ \tilde{\delta}$ can only be done if some of the \tilde{D}_l are small. Thus, while in principle we could use a standard surface-diffusion (i.e., $\tilde{D}_l \sim \exp(l(l+1)t)$ for some fixed diffusion-time t), we find that more efficient mollifiers can be constructed by solving a quadratic program which balances the constraints placed on \tilde{D} .

For example, given a set of coefficients \tilde{D}_l , we can write $I^2(\beta_{\text{north}}, \beta_{\text{south}}; \tilde{D})$ as the quadratic-form:

$$I^2(\beta_{\text{north}}, \beta_{\text{south}}; \tilde{D}) = \sum_{l=0}^L \sum_{l'=0}^L \tilde{D}_l \cdot \tilde{H}(l, l'; \beta_{\text{north}}, \beta_{\text{south}}) \cdot \tilde{D}_{l'},$$

with a kernel:

$$\tilde{H}(l, l'; \beta_{\text{north}}, \beta_{\text{south}}) = 8\pi^2 \int_{\beta \in [\beta_{\text{north}}, \beta_{\text{south}}]} \sqrt{(1+2l)(1+2l')} P_l(\cos(\beta)) P_{l'}(\cos(\beta)) \sin(\beta) d\beta.$$

By fixing $\beta_{\text{north}}, \beta_{\text{south}}$ and \tilde{D}_{min} and running a quadratic-program with a constraint $\tilde{D}_l \geq \tilde{D}_{\text{min}}$, we can quickly find the optimal values of \tilde{D}_l and the associated error ϵ_{pole}^2 . This procedure can be run as a precomputation, scanning through values of $\beta_{\text{north}}, \beta_{\text{south}}$ and \tilde{D}_{min} to find a mollifier which best satisfies the sparsity and error-tolerance constraints provided by the user.

Note that the quadratic-program described above only controls the amplitude of the mollified sinc away from the poles, but does not control the amplitude of the first- or second-derivatives of the mollified sinc. Controlling these derivatives can be easily accomplished by adding to $\tilde{H}(l, l'; \beta_{\text{north}}, \beta_{\text{south}})$ additional kernels that integrate higher-order derivatives of the $P_l(\cos(\beta))$.

Note also that, in the discussion above, we have restricted ourselves to mollifiers \tilde{D} that were diagonal in the spherical-harmonic basis. This is certainly not necessary, as our main design constraint involved being able to apply \tilde{D} ‘quickly’ in the spherical-harmonic basis. This constraint implies that \tilde{D} could be block-diagonal (e.g., like the wigner-d matrix), or more generally sparse, or even factorizable. With these extra degrees of freedom available to \tilde{D} one can design mollifiers that accurately capture $\tilde{\delta}$ and its derivatives, all while having sparse support.

16.2 accommodating derivatives:

We can extend the notion of back-propagation to deal with certain differential quantities. To illustrate this idea, let’s consider the following template-wise summation $\Sigma_\iota(M)$:

$$\Sigma_\iota(M) = \sum_j \int_{\mathbf{k} \in \Omega(K) \subset \mathbb{R}^2} M_j^\dagger(\mathbf{k}) \cdot \partial_{\tau_{j,\iota}} \left[\widehat{\text{Sl}} \circ \mathcal{R}(\tau_j) \circ \widehat{F} \right] (\mathbf{k}) d\mathbf{k}, \quad (67)$$

where $\partial_{\tau_{j,\iota}}$ refers to the ι -component of τ_j (i.e., one of $\{\gamma_j, \beta_j, \alpha_j\}$).

Again, due to linearity, we can associate the data-pool $\{M_j(\mathbf{k})\}$ with the back-propagated shell:

$$\tilde{M}_\iota(\boldsymbol{\kappa}) = \sum_j \int_{\psi \in [0, 2\pi]} M_j(\mathbf{k}) \cdot \left[\partial_{\tau_{j,\iota}} \mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta} \right] (\hat{\boldsymbol{\kappa}}) d\psi, \quad (68)$$

$$= \sum_j \int_{\psi \in [0, 2\pi]} M_j(\mathbf{k}) \cdot \left[\mathcal{R}^{-1}(\tau_j) \circ \tilde{\delta}_{j,\iota,\psi}^{\text{dipole}} \right] (\hat{\boldsymbol{\kappa}}) d\psi, \quad (69)$$

where the dipole-impulse function $\tilde{\delta}_{j,\iota,\psi}^{\text{dipole}}$ is given by:

$$\tilde{\delta}_{j,\iota,\psi}^{\text{dipole}} = \mathcal{R}(\tau_j) \circ \partial_{\tau_j,\iota} \mathcal{R}^{-1}(\tau_j) \circ \mathcal{R}([0, \pi/2, \psi]) \circ \tilde{\delta}. \quad (70)$$

This dipole-impulse corresponds to placing two monopole impulses of opposite signs infinitesimally close to one another near the position $\boldsymbol{\kappa} = \mathcal{R}^{-1}(\tau_j) \circ \mathbf{k}$, with the dipole oriented in a direction (along the surface of the sphere) determined by τ_j , ψ and ι .

With this reorganization we can once again represent the sum $\Sigma(M)$ as a volumetric integral:

$$\Sigma_\iota(M) = \int_{\boldsymbol{\kappa} \in \Omega(K) \subset \mathbb{R}^3} \tilde{M}_\iota^\dagger(\boldsymbol{\kappa}) \cdot \hat{F}(\boldsymbol{\kappa}) \cdot \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|}. \quad (71)$$

17 Volumetric representation of the likelihood in the low-temperature limit

The single-particle negative-log-likelihood shown in Eq. 54 can be written as a volumetric-integral. To see this more clearly, let's write:

$$-\log \left(P \left(\{\hat{A}_j\} \mid \hat{F} \right) \right) \sim \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbf{k} \in \Omega_K \subset \mathbb{R}^2} \left| \hat{A}_j(\mathbf{k}) - CTF_j(\mathbf{k}) \odot \hat{\mathcal{T}}(\boldsymbol{\delta}; \mathbf{k}) \odot \left[\hat{\text{S}}\iota \circ \mathcal{R}(\tau_j) \circ \hat{F} \right](\mathbf{k}) \right|^2 d\mathbf{k}. \quad (72)$$

By interchanging the integral and the sum, we can write:

$$-\log \left(P \left(\{\hat{A}_j\} \mid \hat{F} \right) \right) \sim \frac{1}{2\hat{\sigma}^2} \iiint_{\boldsymbol{\kappa} \in \Omega_K \subset \mathbb{R}^3} \tilde{J}(\boldsymbol{\kappa}) \left\{ \left| \hat{F}(\boldsymbol{\kappa}) - \hat{F}^{\text{bp}}(\boldsymbol{\kappa}) \right|^2 + \hat{F}^{\text{var}}(\boldsymbol{\kappa}) \right\} \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|}, \quad (73)$$

with \tilde{J} and \hat{F}^{bp} given by Eq. 60 and 61, and the second-moment \hat{F}^{var} given by:

$$\hat{F}^{\text{var}}(\boldsymbol{\kappa}) + \left| \hat{F}^{\text{bp}}(\boldsymbol{\kappa}) \right|^2 = \frac{1}{\tilde{J}(\boldsymbol{\kappa})} \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} \left| \hat{\mathcal{T}}(\boldsymbol{\delta}; \mathbf{k}) \odot \hat{A}_{j'}(\mathbf{k}) \right|^2 = \frac{1}{\tilde{J}(\boldsymbol{\kappa})} \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} \left| \hat{A}_{j'}(\mathbf{k}) \right|^2 \quad \text{with: } \mathbf{k} = \mathcal{R}(\tau_{j'}) \circ \boldsymbol{\kappa}. \quad (74)$$

This ‘volumetric’ expression represents the likelihood as a collection of 3-dimensional integrals; this can be compared with the template-wise expression from Eq. 54. Note also that, as mentioned earlier, in the low-temperature-limit the likelihood is maximized when \hat{F} matches the first-moment \hat{F}^{bp} . This first-moment \hat{F}^{bp} itself depends on the maximum-likelihood-alignments $\{\tau_j^{\text{opt}}\}$, $\{\boldsymbol{\delta}_j^{\text{opt}}\}$, the $\{CTF_j\}$ and the images $\{\hat{A}_j\}$.

Using a very similar strategy, we can rewrite the likelihood in Eq. 48 as:

$$\begin{aligned} P \left(\{\hat{A}_j\} \mid \hat{F} \right) &= \int \dots \int_{\{\tau_j\}, \{\boldsymbol{\delta}_j\}} \\ &\times \frac{1}{(\sqrt{2\pi} \hat{\sigma})^{\text{dof}_J}} \exp \left(-\frac{1}{2\hat{\sigma}^2} \iiint_{\boldsymbol{\kappa} \in \Omega_K \subset \mathbb{R}^3} \tilde{J}(\boldsymbol{\kappa}) \left\{ \left| \hat{F}(\boldsymbol{\kappa}) - \hat{F}^{\text{bp}}(\boldsymbol{\kappa}) \right|^2 + \hat{F}^{\text{var}}(\boldsymbol{\kappa}) \right\} \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|} \right) \\ &\times P \left(\{\tau_j\}, \{\boldsymbol{\delta}_j\} \mid \hat{F} \right) \times d\{\tau_j\} d\{\boldsymbol{\delta}_j\}, \end{aligned} \quad (75)$$

where the outer integral is over all possible sets of image-alignments $\{\tau_j\}, \{\boldsymbol{\delta}_j\}$. The terms \tilde{J} , \hat{F}^{bp} and \hat{F}^{var} within the integrand are each functions of the alignments $\{\tau_j\}, \{\boldsymbol{\delta}_j\}$. The probability in the integrand is:

$$P \left(\{\tau_j\}, \{\boldsymbol{\delta}_j\} \mid \hat{F} \right) = \prod_j P \left(\tau_j, \boldsymbol{\delta}_j \mid \hat{F} \right), \quad (76)$$

which can be interpreted as the prior-probability of observing a collection of alignments, given the particle \hat{F} .

To make further manipulation easier, we'll rewrite Eq. 73 as:

$$\mathcal{L} \left(\{\hat{A}_j\} \mid \hat{F}; \{\tau_j^{\text{opt}}\}; \{\boldsymbol{\delta}_j^{\text{opt}}\} \right) \sim \frac{1}{2} \iiint_{\boldsymbol{\kappa}} \left(M_{20}(\boldsymbol{\kappa}) \left| \hat{F}(\boldsymbol{\kappa}) \right|^2 - 2\Re \left(M_{11}^\dagger(\boldsymbol{\kappa}) \hat{F}(\boldsymbol{\kappa}) \right) + M_{02}(\boldsymbol{\kappa}) \right) \frac{d\boldsymbol{\kappa}}{|\boldsymbol{\kappa}|}, \quad (77)$$

where \mathcal{L} is $\hat{\sigma}^2$ times the negative-log-likelihood, and M_{20} , M_{11} and M_{02} represent accumulations of the CTF- and image-data:

$$\mathcal{L} \left(\{\hat{A}_j\} \mid \hat{F}; \{\tau_j^{\text{opt}}\}; \{\delta_j^{\text{opt}}\} \right) := -\hat{\sigma}^2 \log \left(P \left(\{\hat{A}_j\} \mid \hat{F} \right) \right), \quad (78)$$

$$M_{20} \left(\boldsymbol{\kappa}; \{\tau_j^{\text{opt}}\}; \{\delta_j^{\text{opt}}\}; \{CTF_j\} \right) := \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} |CTF_{j'}(\mathbf{k})|^2, \quad (79)$$

$$M_{11} \left(\boldsymbol{\kappa}; \{\tau_j^{\text{opt}}\}; \{\delta_j^{\text{opt}}\}; \{CTF_j\}; \{\hat{A}_j\} \right) := \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} CTF_{j'}(\mathbf{k}) \cdot \hat{\mathcal{T}}(-\delta_{j'}, \mathbf{k}) \cdot \hat{A}_{j'}(\mathbf{k}), \quad (80)$$

$$M_{02} \left(\boldsymbol{\kappa}; \{\hat{A}_j\} \right) := \sum_{j' \in \mathcal{J}(\boldsymbol{\kappa})} \left| \hat{A}_{j'}(\mathbf{k}) \right|^2. \quad (81)$$

18 Synthetic-images

18.1 single-particle:

Given a volume \hat{F} , we can create a distribution of synthetic-images by adding the appropriate amount of noise to the templates associated with \hat{F} . To construct this distribution we must choose the prior joint-distribution of alignments and CTFs, denoted by $P(\tau, \boldsymbol{\delta}, CTF \mid \hat{F})$, as well as the temperature $\hat{\sigma}$. Given these ingredients, we can create a synthetic-image $\hat{A}^{\text{syn}}(\mathbf{k})$ by randomly drawing $\{\tau, \boldsymbol{\delta}, CTF\}$ from $P(\tau, \boldsymbol{\delta}, CTF \mid \hat{F})$, then producing the template $\hat{S}(\mathbf{k}; \tau; \boldsymbol{\delta}; CTF; \hat{F})$, and then finally adding noise in a manner consistent with Eq. 28. In the low-temperature limit we do the same, except in this case the synthetic-images are simply the templates $\hat{S}(\mathbf{k}; \tau; \boldsymbol{\delta}; CTF; \hat{F})$ with no additional noise added. With this construction, each synthetic-image $\hat{A}_j^{\text{syn}}(\mathbf{k})$ is associated with a particular viewing-angle τ_j , displacement $\boldsymbol{\delta}_j$ and CTF_j , and we can express this dependence by writing $\hat{A}_j^{\text{syn}}(\mathbf{k}; \tau_j; \boldsymbol{\delta}_j; CTF_j; \hat{F})$.

18.2 consistency between the likelihood and the synthetic-image distribution:

Note that the likelihood in Eq. 48 is consistent with the synthetic-images described above. To step through the details, let's start by considering the simpler case where we ignore displacements and CTFs (i.e., we assume each $\boldsymbol{\delta}_j \equiv 0$ and each $CTF_j \equiv 1$). In this simple case the volume \hat{F} is associated with the family of templates $\{\hat{S}(\mathbf{k}; \tau; \hat{F})\}$, indexed by the viewing-angle τ . Given a specific distribution of viewing-angles $P(\tau \mid \hat{F})$, the family of templates $\{\hat{S}(\mathbf{k}; \tau; \hat{F})\}$ produces a distribution $P(\hat{S} \mid \hat{F})$ of templates in image-space (here \hat{S} is a dummy-variable indicating the template's location in image-space).

Assuming isotropic noise with variance $\hat{\sigma}^2$, the distribution of synthetic-images $P(\hat{A}^{\text{syn}} \mid \hat{F})$ associated with $P(\hat{S} \mid \hat{F})$ is simply the convolution of $P(\hat{S} \mid \hat{F})$ with a gaussian:

$$P(\hat{A}^{\text{syn}} \mid \hat{F}) = P(\hat{S} \mid \hat{F}) \star \mathcal{N}(0, \hat{\sigma}^2).$$

In this formula \hat{A}^{syn} is a dummy variable indicating location in image-space, and $\mathcal{N}(0, \hat{\sigma})$ represents the gaussian in image-space associated with the isotropic variance of $\hat{\sigma}^2$ per unit-area in \mathbf{k} -space.

With this perspective, the likelihood $P(\{A_j\} \mid \hat{F})$ in Eq. 48 is equal to the probability of drawing the image-set $\{\hat{A}_j\}$ from the distribution of synthetic-images $P(\hat{A}^{\text{syn}} \mid \hat{F})$, provided that the likelihood is calculated using the same prior distribution $P(\tau \mid \hat{F})$ and temperature $\hat{\sigma}$ used to generate the distribution of synthetic-images.

18.3 consistency between \hat{F} and the synthetic-images:

Following directly from the construction of the synthetic-image distribution $P(\hat{A}^{\text{syn}} \mid \hat{F})$ above, we can see that the function \hat{F} along with the prior distribution $P(\tau \mid \hat{F})$ are a global maximum of the likelihood:

$$\left\{ \hat{F}, P(\tau \mid \hat{F}) \right\} = \arg \max_{\{\hat{G}, P(\tau \mid \hat{G})\}} \lim_{J \rightarrow \infty} P \left(\{A_j^{\text{syn}}\} \mid \hat{G}; P(\tau \mid \hat{G}); \hat{\sigma} \right), \quad (82)$$

where we have emphasized the dependence of the likelihood $P(\{A_j^{\text{syn}}\} \mid \hat{G})$ on the prior distribution and temperature. Note that in Eq. 82 we have assumed (via $J \rightarrow \infty$) that $\{A_j^{\text{syn}}\}$ is a pool of synthetic-images which is sufficiently large to sample $P(\hat{A}^{\text{syn}} \mid \hat{F})$.

To see why Eq. 82 might be true, we can consider any other candidate template-distribution $P'(\widehat{S})$. Given any temperature $\hat{\sigma}'$, the associated distribution of synthetic-images $P'(\widehat{A}^{\text{syn}})$ is:

$$P'(\widehat{A}^{\text{syn}}) = P'(\widehat{S}) \star \mathcal{N}(0, [\hat{\sigma}']^2).$$

Now, given an infinite number of synthetic-images $\{\widehat{A}_j^{\text{syn}}\}$ sampled from $P(\widehat{A}^{\text{syn}} | \widehat{F})$, we can ask the following questions: (i) what is the probability of drawing this image-set from $P(\widehat{A}^{\text{syn}} | \widehat{F})$, and (ii) what is the probability of drawing this same image-set from $P'(\widehat{A}^{\text{syn}})$. The logarithm of the ratio of the probabilities from (i) and (ii) is proportional to the Kullback-Leibler divergence:

$$KL(P||P') := \int_{\widehat{A}^{\text{syn}}} P(\widehat{A}^{\text{syn}} | \widehat{F}) \log \left(\frac{P(\widehat{A}^{\text{syn}} | \widehat{F})}{P'(\widehat{A}^{\text{syn}})} \right) d\widehat{A}^{\text{syn}},$$

where the integral is carried out over image-space.

The KL-divergence is always nonnegative, and is only equal to 0 when $P' \equiv P$. This directly implies that the log-probability-ratio is highest when $P(\widehat{A}^{\text{syn}} | \widehat{F}) \equiv P'(\widehat{A}^{\text{syn}})$. Put more simply: given any distribution P , no other distribution P' can possibly produce a collection of samples that better represents P . Consequently, assuming the approximate temperature $\hat{\sigma}'$ matches the true temperature $\hat{\sigma}$ used to generate the synthetic-images, no template-distribution $P'(\widehat{S})$ can produce a higher likelihood than the template-distribution $P(\widehat{S} | \widehat{F})$. Thus, no other pair $\{\widehat{G}, P(\tau | \widehat{G})\}$ can generate a higher likelihood than the true $\{\widehat{F}, P(\tau | \widehat{F})\}$.

Note that it is still possible for other volumes to produce the same likelihood as the true volume \widehat{F} (i.e., \widehat{F} is not guaranteed to be the only global maximum of Eq. 82). Specifically, another volume \widehat{G} might look different from \widehat{F} , but produce a family of templates $\{\widehat{S}(\mathbf{k}; \tau; \widehat{G})\}$ that overlaps (in image-space) the family of templates $\{\widehat{S}(\mathbf{k}; \tau; \widehat{F})\}$ produced by \widehat{F} . In this scenario there exist viewing-angle distributions $P(\tau | \widehat{G})$ and $P(\tau | \widehat{F})$ such that the induced distribution $P'(\widehat{S}) := P(\widehat{S} | \widehat{G})$ is the same as $P(\widehat{S} | \widehat{F})$. Examples along these lines are discussed below.

18.4 multi-particle:

We can also use a similar strategy to create a pool of synthetic-images associated with multiple volumes $\{\widehat{F}_i\}$. In this case we draw alignments, CTFs and volume-labels from the prior joint-distribution $P(\tau, \delta, CTF, i | \{\widehat{F}_i\})$.

In this multi-volume scenario a statement similar to Eq. 82 holds. Suppressing the notation for δ and the CTF momentarily, we can write:

$$\left\{ \{\widehat{F}_i\}, P(\tau, i | \{\widehat{F}_i\}) \right\} = \arg \max_{\{\{\widehat{G}_i\}, P(\tau, i | \{\widehat{G}_i\})\}} \lim_{J \rightarrow \infty} P(\{A_j^{\text{syn}}\} | \{\widehat{G}_i\}; P(\tau, i | \{\widehat{G}_i\}); \hat{\sigma}), \quad (83)$$

where again we have emphasized the dependence of the likelihood $P(\{A_j^{\text{syn}}\} | \{\widehat{G}_i\})$ on the prior distribution and temperature.

Once again, it is possible for other volume-collections to produce the same likelihood as the true volume-collection $\{\widehat{F}_i\}$. This phenomenon is much more pronounced in the multi-volume setting than in the single-volume setting, and it is rather straightforward to construct alternative volume-sets $\{\widehat{G}_i\}$, each with their own viewing-angle distribution $P(\tau | \widehat{G}_i)$, which collectively produce the same distribution of synthetic-images as $P(\widehat{A}^{\text{syn}} | \{\widehat{F}_i\})$. Examples along these lines are discussed below.

19 Local well-posedness of single-particle reconstruction

For a fixed data-set $\{\widehat{A}_j\}$, the likelihood $P(\{\widehat{A}_j\} | \widehat{F}; P(\tau; \delta | \widehat{F}); \hat{\sigma})$ from Eq. 48 is usually a nonconvex function of \widehat{F} , $P(\tau; \delta | \widehat{F})$ and $\hat{\sigma}$. Even if we assume that the alignment-distribution $P(\tau; \delta | \widehat{F})$ and temperature $\hat{\sigma}$ are known, finding the ‘optimal’ \widehat{F}^{opt} which maximizes the likelihood is often not easy [31, 32, 35, 16]. In general, the iterative strategies used to search for \widehat{F}^{opt} produce an approximate \widehat{F}^{loc} which has a local (but not global) maximum of the likelihood $P(\{\widehat{A}_j\} | \widehat{F}^{\text{loc}})$, which we’ll denote by P^{loc} .

One standard strategy for measuring the quality of any particular configuration $\{\widehat{F}^{\text{loc}}; P(\tau; \delta | \widehat{F}^{\text{loc}})\}$ is to search for nearby configurations that might have a similar likelihood. More specifically, we can consider nontrivial infinitesimal perturbations of the form:

$$\left\{ \widehat{F}^{\text{loc}}; P(\tau; \delta | \widehat{F}^{\text{loc}}) \right\} \rightarrow \left\{ \widehat{F}^{\text{loc}} + \epsilon \Delta \widehat{F}; P(\tau; \delta | \widehat{F}^{\text{loc}}) + \epsilon \Delta P(\tau; \delta) \right\}. \quad (84)$$

In Eq. 84 the function $\Delta\widehat{F}(\boldsymbol{\kappa})$ is the Fourier-transform of a function $\Delta F(\boldsymbol{\chi}) \in L^2(\mathbb{R}^3)$ representing a perturbation of the volume $F^{\text{loc}}(\boldsymbol{\chi})$. Similarly, the function $\Delta P(\tau; \boldsymbol{\delta})$ represents a perturbation of the prior distribution $P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})$. We'll assume that both $\Delta\widehat{F}(\boldsymbol{\kappa})$ and $\Delta P(\tau; \boldsymbol{\delta})$ have unit norm, and that $\epsilon \in \mathbb{R}$ is small. When considering 'nontrivial' perturbations, we'll restrict our attention only to infinitesimal deformations of the solution, and implicitly exclude those perturbations that correspond to infinitesimal rigid rotations or translations of the original volume \widehat{F} and prior-distribution $P(\tau; \boldsymbol{\delta} | \widehat{F})$.

Assuming that the configuration $\{\widehat{F}^{\text{loc}}; P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})\}$ is indeed a local maximum of the likelihood, then the perturbed negative-log-likelihood should be approximately quadratic to leading-order in ϵ for a typical choice of $\Delta\widehat{F}$ and ΔP :

$$-\log\left(P\left(\{\widehat{A}_j\} | \widehat{F}^{\text{loc}} + \epsilon\Delta\widehat{F}(\boldsymbol{\kappa}); P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}}) + \epsilon\Delta P(\tau; \boldsymbol{\delta}); \hat{\sigma}\right)\right) = -\log(P^{\text{loc}}) + \frac{1}{2}\sigma_H(\Delta\widehat{F}, \Delta P)\epsilon^2 + \mathcal{O}(\epsilon^3). \quad (85)$$

In the expression above, $\sigma_H(\Delta\widehat{F}, \Delta P)$ is the concavity of the negative-log-likelihood in the direction of $\{\Delta\widehat{F}, \Delta P\}$. The second-order term in Eq. 85 is equivalent to a rayleigh-quotient of the quadratic kernel H :

$$\sigma_H(\Delta\widehat{F}, \Delta P) = \int \dots \int \Delta P(\tau; \boldsymbol{\delta}) \Delta\widehat{F}(\boldsymbol{\kappa})^\dagger \cdot H(\boldsymbol{\kappa}, \tau, \boldsymbol{\delta}, \boldsymbol{\kappa}', \tau', \boldsymbol{\delta}') \cdot \Delta\widehat{F}(\boldsymbol{\kappa}') \Delta P(\tau', \boldsymbol{\delta}') d\boldsymbol{\kappa}' d\tau' d\boldsymbol{\delta}' d\boldsymbol{\kappa} d\tau d\boldsymbol{\delta}, \quad (86)$$

where H is the hessian of the negative-log-likelihood with respect to perturbations of \widehat{F}^{loc} and $P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})$:

$$H(\boldsymbol{\kappa}, \tau, \boldsymbol{\delta}, \boldsymbol{\kappa}', \tau', \boldsymbol{\delta}') = \partial_{\{\widehat{F}(\boldsymbol{\kappa}); P(\tau, \boldsymbol{\delta} | \widehat{F})\}} \partial_{\{\widehat{F}(\boldsymbol{\kappa}'); P(\tau', \boldsymbol{\delta}' | \widehat{F}')\}} \left\{ -\log\left(P\left(\{\widehat{A}_j\} | \widehat{F}^{\text{loc}}; P(\tau, \boldsymbol{\delta} | \widehat{F}^{\text{loc}}); \hat{\sigma}\right)\right) \right\}, \quad (87)$$

which is directly related to a kind of Fisher information.

If the magnitude $\sigma_H(\Delta\widehat{F}, \Delta P)$ is large for every nontrivial $\{\Delta\widehat{F}, \Delta P\}$, then P^{loc} indicates a sharp minimum of the negative-log-likelihood-function. In this scenario we can say that the single-particle reconstruction problem is well-posed within a local neighborhood of $\{\widehat{F}^{\text{loc}}; P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})\}$. Conversely, if there exist nontrivial directions $\{\Delta\widehat{F}, \Delta P\}$ for which $\sigma_H(\Delta\widehat{F}, \Delta P)$ is small, then these directions correspond to nontrivial perturbations (i.e., deformations) along which the likelihood diminishes only gradually. Put another way, there will be 'nearby' configurations that will have likelihoods similar to the original P^{loc} . Thus, in this latter scenario the single-particle reconstruction problem is poorly-posed near $\{\widehat{F}^{\text{loc}}; P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})\}$.

One could in principle use Eq. 86 to characterize the well-posedness of the reconstruction near any given $\{\widehat{F}^{\text{loc}}; P(\tau; \boldsymbol{\delta} | \widehat{F}^{\text{loc}})\}$ by finding the nontrivial direction $\{\Delta\widehat{F}^{\text{crit}}, \Delta P^{\text{crit}}\}$ which minimizes $\sigma_H(\Delta\widehat{F}, \Delta P)$:

$$\left\{ \Delta\widehat{F}^{\text{crit}}, \Delta P^{\text{crit}} \right\} = \arg \min_{\{\Delta\widehat{F}, \Delta P\}} \sigma_H(\Delta\widehat{F}, \Delta P) \quad \text{with} \quad \sigma_H^{\text{crit}} = \sigma_H(\Delta\widehat{F}^{\text{crit}}, \Delta P^{\text{crit}}). \quad (88)$$

Note that σ_H^{crit} is the smallest eigenvalue of H , and $\{\Delta\widehat{F}^{\text{crit}}, \Delta P^{\text{crit}}\}$ is the associated eigenvector.

We are not aware of an efficient general strategy for calculating σ_H^{crit} and $\{\Delta\widehat{F}^{\text{crit}}, \Delta P^{\text{crit}}\}$. The main impasse is that we do not know how to apply the hessian H (or its inverse) in a computationally efficient manner. However, any given nontrivial direction $\{\Delta\widehat{F}, \Delta P\}$ can be used to give an upper-bound for σ_H^{crit} . Because volume-space is high-dimensional, we use the low-temperature limit (described in the next section) to guide our search.

20 Estimating the hessian in the low-temperature limit

As mentioned in the previous section, we are not aware of any efficient strategies for calculating the smallest eigenvalue of the hessian of the negative-log-likelihood from Eq. 87. However, there are strategies that are useful in the low-temperature limit.

To ease the notation, let's ignore the displacements for now, and let's assume for the moment that the viewing-angles $\{\tau_j\}$ for each image are given. Let's use $\langle \cdot, \cdot \rangle$ to represent the standard inner product over $\mathbf{k} \in \Omega_K \subset \mathbb{R}^2$, and $\int_{\mathbf{k}}$ to represent the integral over $\boldsymbol{\kappa} \in \Omega_K \subset \mathbb{R}^3$ with respect to the weight $d\boldsymbol{\kappa}/|\boldsymbol{\kappa}|$.

With this notation we can use a 2-dimensional representation and define \mathcal{L} to be the negative-log-likelihood scaled by $\hat{\sigma}^2$ from Eq. 54

$$\mathcal{L}\left(\{\widehat{A}_j\} | \widehat{F}; \{\tau_j\}\right) := -\hat{\sigma}^2 \log\left(P\left(\{\widehat{A}_j\} | \widehat{F}; \{\tau_j\}\right)\right), \quad (89)$$

$$:= \frac{1}{2} \sum_{j=1}^J \iint_{\mathbf{k}} \left| \widehat{S}\left(\mathbf{k}; \tau_j; CTF_j; \widehat{F}\right) - \widehat{A}_j(\mathbf{k}) \right|^2 d\mathbf{k}, \quad (90)$$

$$= \frac{1}{2} \sum_{j=1}^J \langle \widehat{S}\left(\tau_j; CTF_j; \widehat{F}\right) - \widehat{A}_j, \widehat{S}\left(\tau_j; CTF_j; \widehat{F}\right) - \widehat{A}_j \rangle, \quad (91)$$

where we have suppressed reference to \mathbf{k} in the integrand, and emphasized the dependence on τ_j and \widehat{F} .

Equivalently, we can use a 3-dimensional representation and write \mathcal{L} from Eq. 77 as:

$$\mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) = \frac{1}{2} \int_{\kappa} \left(M_{20} (\{\tau_j\}; \{CTF_j\}) \cdot |\widehat{F}|^2 - 2\Re \left(M_{11}^\dagger (\{\tau_j\}; \{CTF_j\}; \{\widehat{A}_j\}) \cdot \widehat{F} \right) + M_{02} (\{\widehat{A}_j\}) \right), \quad (92)$$

where again we have suppressed reference to κ in the integrand, and emphasized the dependence on τ_j and \widehat{F} .

The 2-dimensional representation Eq. 91 makes apparent the following relationships:

$$\partial_{\widehat{F}} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) \cdot \Delta \widehat{F} = \sum_j \Re \langle \widehat{S} (\tau_j; CTF_j; \Delta \widehat{F}), \widehat{S} (\tau_j; CTF_j; \widehat{F}) - \widehat{A}_j \rangle, \quad (93)$$

$$\partial_{\tau_j} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) = \Re \langle \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \widehat{F}), \widehat{S} (\tau_j; CTF_j; \widehat{F}) - \widehat{A}_j \rangle, \quad (94)$$

$$\begin{aligned} \partial_{\tau_j} \partial_{\widehat{F}} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) \cdot \Delta \widehat{F} &= \\ \Re \langle \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \Delta \widehat{F}), \widehat{S} (\tau_j; CTF_j; \widehat{F}) - \widehat{A}_j \rangle &+ \Re \langle \widehat{S} (\tau_j; CTF_j; \Delta \widehat{F}), \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \widehat{F}) \rangle, \end{aligned} \quad (95)$$

$$\begin{aligned} \partial_{\tau_j} \partial_{\tau_j} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) &= \\ \Re \langle \partial_{\tau_j} \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \widehat{F}), \widehat{S} (\tau_j; CTF_j; \widehat{F}) - \widehat{A}_j \rangle &+ \Re \langle \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \widehat{F}), \partial_{\tau_j} \widehat{S} (\tau_j; CTF_j; \widehat{F}) \rangle. \end{aligned} \quad (96)$$

In these expressions τ_j involves the three euler-angles $\tau_j = [\gamma_j, \beta_j, \alpha_j]^\top$, and so $\partial_{\tau_j} \mathcal{L}$ and $\partial_{\tau_j} \partial_{\widehat{F}} \mathcal{L}$ each have three components, while $\partial_{\tau_j} \partial_{\tau_j} \mathcal{L}$ is a 3×3 tensor. Note that Eq. 93 provides an expression for the directional-derivative of \mathcal{L} , given a particular volumetric-perturbation $\Delta \widehat{F}$, while Eq. 94 provides a general expression for the gradient of \mathcal{L} with respect to arbitrary perturbations in τ_j . Similarly, Eq. 95 provides an expression for the hessian applied to a particular volumetric-perturbation $\Delta \widehat{F}$, while Eq. 96 gives a general expression for the hessian of \mathcal{L} with respect to arbitrary perturbations in τ_j .

The 3-dimensional representation Eq. 92 demonstrates:

$$\begin{aligned} \partial_{\widehat{F}(\kappa')} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) &= \\ \Re \left(M_{20} (\kappa'; \{\tau_j\}; \{CTF_j\}) \cdot \widehat{F}(\kappa') \right) &- \Re \left(M_{11}^\dagger (\kappa'; \{\tau_j\}; \{CTF_j\}; \{\widehat{A}_j\}) \right), \end{aligned} \quad (97)$$

$$\begin{aligned} \partial_{\tau_j} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) \cdot \Delta \tau_j &= \\ \frac{1}{2} \int_{\kappa} \left(\partial_{\tau_j} M_{20} (\{\tau_j\}; \{CTF_j\}) \cdot |\widehat{F}|^2 \right. &- \left. 2\Re \left(\partial_{\tau_j} M_{11}^\dagger (\{\tau_j\}; \{CTF_j\}; \{\widehat{A}_j\}) \cdot \widehat{F} \right) \right) \cdot \Delta \tau_j, \end{aligned} \quad (98)$$

$$\partial_{\widehat{F}(\kappa')} \partial_{\widehat{F}(\kappa')} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) = M_{20} ((\kappa'); \{\tau_j\}; \{CTF_j\}), \quad (99)$$

$$\begin{aligned} \partial_{\widehat{F}(\kappa')} \partial_{\tau_j} \mathcal{L} \left(\{\widehat{A}_j\} \mid \widehat{F}; \{\tau_j\} \right) \cdot \Delta \tau_j &= \\ \left(\Re \left(\partial_{\tau_j} M_{20} (\kappa'; \{\tau_j\}; \{CTF_j\}) \cdot \widehat{F}(\kappa') \right) \right. &- \left. \Re \left(\partial_{\tau_j} M_{11}^\dagger (\kappa'; \{\tau_j\}; \{CTF_j\}; \{\widehat{A}_j\}) \right) \right) \cdot \Delta \tau_j. \end{aligned} \quad (100)$$

In these expressions $\Delta \tau_j$ also involves three euler-angles $\Delta \tau_j = [\Delta \gamma_j, \Delta \beta_j, \Delta \alpha_j]^\top$, and so $\partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j$ and $\partial_{\widehat{F}} \partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j$ each involve a sum over three components. Note that Eq. 97 provides a general expression for the gradient of \mathcal{L} with respect to any perturbation in \widehat{F} , while Eq. 98 provides an expression for the directional-derivative of \mathcal{L} with respect to a particular perturbation in τ_j . Similarly, Eq. 99 provides a general expression for the hessian with respect to arbitrary perturbations in \widehat{F} , while Eq. 100 gives an expression for the hessian applied to a particular viewing-angle perturbation $\Delta \tau_j$.

Note that we can leverage both the 2- and 3-dimensional representations to write the general gradient:

$$G \cdot \Delta := \begin{bmatrix} G_F & G_\tau \end{bmatrix} \cdot \begin{bmatrix} \Delta \widehat{F} \\ \{\Delta \tau_j\} \end{bmatrix} = \begin{bmatrix} \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j \end{bmatrix}. \quad (101)$$

Specifically, Eq. 97 yields the volumetric-component, while Eq. 94 yields the alignment-component.

Similarly, given a particular perturbation $\{\Delta \widehat{F}, \{\Delta \tau_j\}\}$, we can apply the hessian. To describe this in more detail, let's refer to the perturbation $\Delta := [\Delta \widehat{F}, \{\Delta \tau_j\}]^\top$ as a block-vector. The first block of Δ stores the volumetric-perturbation

$\Delta\hat{F}$, while the second block stores the alignment-perturbations $\{\Delta\tau_j\}$ for each image in sequence. With this organization the hessian is also divided into blocks:

$$H \cdot \Delta := \begin{bmatrix} H_{FF} & H_{F\tau} \\ H_{\tau F} & H_{\tau\tau} \end{bmatrix} \cdot \begin{bmatrix} \Delta\hat{F} \\ \{\Delta\tau_j\} \end{bmatrix} = \begin{bmatrix} \partial_{\hat{F}}\partial_{\hat{F}}\mathcal{L} \cdot \Delta\hat{F} + \sum_{j'} \partial_{\hat{F}}\partial_{\tau_{j'}}\mathcal{L} \cdot \Delta\tau_{j'} \\ \partial_{\tau}\partial_{\hat{F}}\mathcal{L} \cdot \Delta\hat{F} + \sum_{j'} \partial_{\tau}\partial_{\tau_{j'}}\mathcal{L} \cdot \Delta\tau_{j'} \end{bmatrix}. \quad (102)$$

The first block H_{FF} refers to the volumetric-component, which is diagonalized in \mathbf{k} -space as shown in Eq. 99. The second term $H_{F\tau}$ refers to the volumetric-perturbation induced by any given alignment-perturbation, as shown in Eq. 100. The third term $H_{\tau F}$ refers to the alignment-perturbation induced by any given volumetric-perturbation, as shown in Eq. 95. Finally, the block $H_{\tau\tau}$ refers to the alignment-component, which is diagonalized in template-space as shown in Eq. 96.

20.1 Finding eigenvalues and eigenvectors of H :

Given the ability to apply H to any perturbation $\Delta := \{\Delta\hat{F}, \{\Delta\tau_j\}\}$, we can use lanczos-iteration to estimate the extremal eigenvalues and eigenvectors of H . One technical detail that comes into play is that we want to avoid the ‘trivial’ eigenvectors of H that correspond to perturbations Δ which are infinitesimal rigid-rotations of $\{\hat{F}, \{\tau_j\}\}$. Thus, when applying lanczos-iteration we project away the 3-dimensional space associated with these trivial rigid-rotations. The remaining subspace will include those perturbations which correspond to infinitesimal deformations. An example illustrating the convergence of lanczos-iteration is shown in Fig 23.

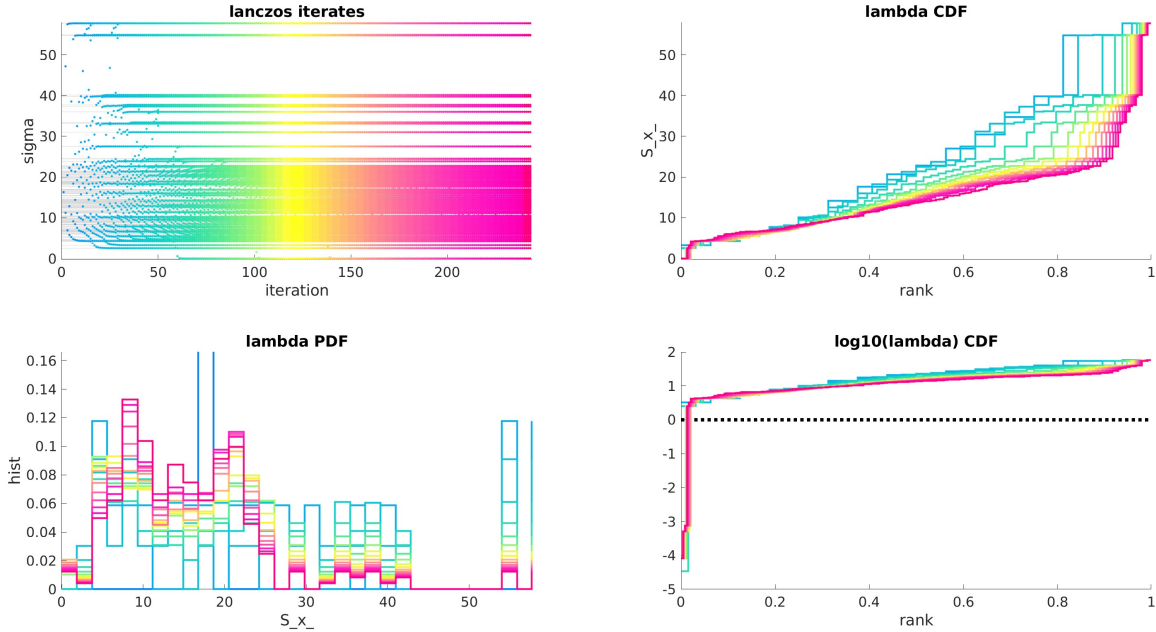


Figure 23: In this figure we illustrate the properties of the hessian H , as well as its lanczos-iterates. For this example we use the TRPV1 molecule at ultra-low resolution: $2\pi K = 8$. This ultra-low resolution allows us to take $L = 2\pi K = 8$, implying only $(1 + L)^2 = 81$ degrees of freedom per frequency-shell. We also project onto the dominant 3 volumetric-principal-modes (capturing over 99.9% of the alignment-information, as described in [48]) and determine the alignment-perturbations implicitly (using the first-order scheme described in section 21). Thus, the total number of degrees-of-freedom for this simple case-study is $3 \times 81 = 243$. The first subplot shows the lanczos-iterates along the vertical, with the iteration-number plotted along the horizontal. The color ranges from cyan (early iterates) to magenta (final iterates). Note that the outliers in the spectrum are captured quite early on. The second subplot shows the cumulative-distribution function for the lanczos-iterates, with the different colors corresponding to different iterations (using the same color-scale shown in the top-left). The third subplot shows the estimated probability-density-function for the lanczos-iterates, while the fourth subplot shows the cumulative-distribution function for the log of the lanczos-iterates. Note that the log of the determinant of the hessian $\log \det(H)$ is the average of the log of the eigenvalues of H , which is proportional to the area under the CDF shown in the bottom-right subplot. Note also that, once the outliers in the spectrum are well approximated, the $\log \det(H)$ is also reasonably well approximated. For this example we have $\sim 25\%$ relative-error for $\log \det(H)$ by iteration 16, and $\sim 10\%$ relative-error by iteration 64.

Note that, while lanczos-iteration does a good job of focusing on the extremal eigenvalues, it does a poor job of

estimating the bulk of the spectrum. Consequently, while we can expect to use the lanczos-iterates to estimate $\log(\det(H))$ in the low-temperature limit, we don't necessarily expect such an estimate to hold for moderate temperature. Nevertheless, the same strategies used above can still be applied when the temperature is not small after marginalization (see section 23).

21 Dealing with implicitly defined τ_j

In the construction above we have assumed that the $\{\tau_j\}$ are given as input. Thus, the gradient in Eq. 101 and the hessian in Eq. 102 both assume that the overall perturbation Δ is given as an (independent) combination of a volumetric-perturbation $\Delta\hat{F}$ and an alignment-perturbation $\{\Delta\tau_j\}$. However, in many scenarios it is more appropriate to assume that the volumetric- and alignment-perturbations are implicitly related to preserve the maximum-likelihood alignment (see Eq. 53). In this section we'll describe how we account for this implicit constraint.

To ease notation, we'll start by considering a single image-index j , setting $\hat{A} := \hat{A}_j$ and $\tau := \tau_j$, and suppressing the notation for the *CTF*. Furthermore, we'll assume that $\tau = \tau^{\text{opt}}$ is determined implicitly by the volume \hat{F} . This implicit constraint minimizes the negative-log-likelihood ℓ for that image:

$$\ell(\hat{A} | \hat{F}, \tau) = \frac{1}{2} \langle \hat{S}(\hat{F}, \tau) - \hat{A}, \hat{S}(\hat{F}, \tau) - \hat{A} \rangle, \quad (103)$$

where $\hat{S}(\hat{F}, \tau)$ is shorthand for the template $\hat{S}(\mathbf{k}; \tau; \text{CTF}; \hat{F})$. Thus, τ should be a critical-point of ℓ , implying that:

$$\partial_{\tau_i} \ell = \frac{1}{2} \langle \partial_{\tau_i} \hat{S}(\hat{F}, \tau), \hat{S}(\hat{F}, \tau) - \hat{A} \rangle + \frac{1}{2} \langle \hat{S}(\hat{F}, \tau) - \hat{A}, \partial_{\tau_i} \hat{S}(\hat{F}, \tau) \rangle = \Re \langle \hat{S}(\hat{F}, \tau) - \hat{A}, \partial_{\tau_i} \hat{S}(\hat{F}, \tau) \rangle = \mathbf{0}, \quad (104)$$

where τ_i runs over each of the three components $[\gamma, \beta, \alpha]^T$ of τ , and each of the 3 components of the gradient $\partial_{\tau} \ell$ are set 0. We expand $\hat{S}(\hat{F}, \tau)$ as follows:

$$\hat{S}(\hat{F} + \Delta\hat{F}, \tau + \Delta\tau) = \hat{S} + \partial_{\hat{F}} \hat{S} \cdot \Delta\hat{F} + \partial_{\tau} \hat{S} \cdot \Delta\tau \quad (105)$$

$$+ \frac{1}{2} [\Delta\hat{F}]^\dagger \cdot [\partial_{\hat{F}\hat{F}} \hat{S}] \cdot [\Delta\hat{F}] + \frac{1}{2} [\Delta\hat{F}]^\dagger \cdot [\partial_{\hat{F}\tau} \hat{S}] \cdot [\Delta\tau] + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\hat{F}} \hat{S}] \cdot [\Delta\hat{F}] + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau], \quad (106)$$

where each term is evaluated at (\hat{F}, τ) unless otherwise specified. Noting that \hat{S} is linear in \hat{F} , we can ignore second-order contributions of $\Delta\hat{F}^2$ to write:

$$\hat{S}(\hat{F} + \Delta\hat{F}, \tau + \Delta\tau) = \hat{S} + \partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] \quad (107)$$

$$+ \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau], \quad (108)$$

where $\Delta\hat{S} := \hat{S}(\mathbf{k}; \tau; \text{CTF}; \Delta\hat{F})$ is simply the original template-operator applied to $\Delta\hat{F}$ rather than \hat{F} .

Now we can write $\langle \hat{S} - \hat{A}, \partial_{\tau_i} \hat{S} \rangle$ as:

$$\begin{aligned} \langle \hat{S} - \hat{A}, \partial_{\tau_i} \hat{S} \rangle = & \left\langle \hat{S} - \hat{A} + \partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] + \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau] \right. \\ & \left. , \partial_{\tau_i} \left[\hat{S} + \partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] + \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau] \right] \right\rangle \end{aligned} \quad (109)$$

Removing the lowest-order term and keeping terms up to second-order we have:

$$\begin{aligned} \langle \hat{S} - \hat{A}, \partial_{\tau_i} \hat{S} \rangle = & \left\langle +\partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] + \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau] \right. \\ & \left. , \partial_{\tau_i} \left[\hat{S} + \partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] + \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau] \right] \right\rangle \\ & + \left\langle \hat{S} - \hat{A}, \partial_{\tau_i} \left[+\partial_{\tau} \hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \hat{S}] \cdot [\Delta\tau] + \Delta\hat{S} + \partial_{\tau} \Delta\hat{S} \cdot \Delta\tau + \frac{1}{2} [\Delta\tau]^\top \cdot [\partial_{\tau\tau} \Delta\hat{S}] \cdot [\Delta\tau] \right] \right\rangle \end{aligned} \quad (110)$$

or when expanded:

$$\begin{aligned}
\langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \widehat{S} \rangle = & \\
& + \langle \partial_{\tau} \widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \widehat{S} + \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau + \partial_{\tau_i} \Delta\widehat{S} \rangle + \langle \frac{1}{2} [\Delta\tau]^{\top} \cdot [\partial_{\tau\tau} \widehat{S}] \cdot [\Delta\tau], \partial_{\tau_i} \widehat{S} \rangle \\
& + \langle \Delta\widehat{S}, \partial_{\tau_i} \widehat{S} + \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau + \partial_{\tau_i} \Delta\widehat{S} \rangle + \langle \partial_{\tau} \Delta\widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \widehat{S} \rangle \\
& + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau + \partial_{\tau_i} \frac{1}{2} [\Delta\tau]^{\top} \cdot [\partial_{\tau\tau} \widehat{S}] \cdot [\Delta\tau] + \partial_{\tau_i} \Delta\widehat{S} + \partial_{\tau_i} \partial_{\tau} \Delta\widehat{S} \cdot \Delta\tau \rangle
\end{aligned} \tag{111}$$

which contains the following four first-order terms:

$$+ \langle \partial_{\tau} \widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \widehat{S} \rangle + \langle \Delta\widehat{S}, \partial_{\tau_i} \widehat{S} \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \Delta\widehat{S} \rangle. \tag{112}$$

as well as the following eight second-order terms:

$$+ \langle \partial_{\tau} \widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau \rangle + \langle \partial_{\tau} \widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \Delta\widehat{S} \rangle + \langle \frac{1}{2} \Delta\tau^{\top} \cdot \partial_{\tau\tau} \widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \widehat{S} \rangle + \langle \Delta\widehat{S}, \partial_{\tau_i} \partial_{\tau} \widehat{S} \cdot \Delta\tau \rangle \tag{113}$$

$$+ \langle \Delta\widehat{S}, \partial_{\tau_i} \Delta\widehat{S} \rangle + \langle \partial_{\tau} \Delta\widehat{S} \cdot \Delta\tau, \partial_{\tau_i} \widehat{S} \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \frac{1}{2} \Delta\tau^{\top} \cdot \partial_{\tau\tau} \widehat{S} \cdot \Delta\tau \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \partial_{\tau} \Delta\widehat{S} \cdot \Delta\tau \rangle. \tag{114}$$

By retaining only first-order terms we see:

$$\sum_{\nu'} [\mathbf{x0}]_{\nu, \nu'} \cdot \Delta\tau_{\nu'} + [\mathbf{x1}]_{\nu} = \mathcal{O}(\Delta\widehat{F}^2), \text{ where:} \tag{115}$$

$$[\mathbf{x0}]_{\nu, \nu'} = \Re \left\{ \langle \partial_{\tau_i} \widehat{S}, \partial_{\tau_{\nu'}} \widehat{S} \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i \tau_{\nu'}} \widehat{S} \rangle \right\}, \tag{116}$$

$$[\mathbf{x1}]_{\nu} = \Re \left\{ \langle \partial_{\tau_i} \widehat{S}, \Delta\widehat{S} \rangle + \langle \partial_{\tau_i} \Delta\widehat{S}, \widehat{S} - \widehat{A} \rangle \right\}, \tag{117}$$

where $[\mathbf{x0}]$ is $\mathcal{O}(1)$ and $[\mathbf{x1}]$ is $\mathcal{O}(\Delta\widehat{F})$, as expected. This first-order expression implies that:

$$\Delta\tau = -[\mathbf{x0}]^{-1} \cdot [\mathbf{x1}] + \mathcal{O}(\Delta\widehat{F}^2). \tag{118}$$

The first-order approximation to $\Delta\tau$ can be improved by using the second-order terms. Let's first group the second-order terms by their tensor-rank:

$$[\mathbf{Y0}]_{\nu, \nu', \nu'} = \Re \left\{ \langle \partial_{\tau_{\nu'}} \widehat{S}, \partial_{\tau_i} \partial_{\tau_{\nu'}} \widehat{S} \rangle + \langle \frac{1}{2} \partial_{\tau_{\nu'}} \partial_{\tau_{\nu'}} \widehat{S}, \partial_{\tau_i} \widehat{S} \rangle + \langle \widehat{S} - \widehat{A}, \frac{1}{2} \partial_{\tau_i} \partial_{\tau_{\nu'}} \partial_{\tau_{\nu'}} \widehat{S} \rangle \right\}, \tag{119}$$

$$[\mathbf{Y1}]_{\nu, \nu'} = \Re \left\{ \langle \partial_{\tau_{\nu'}} \widehat{S}, \partial_{\tau_i} \Delta\widehat{S} \rangle + \langle \partial_{\tau_{\nu'}} \Delta\widehat{S}, \partial_{\tau_i} \widehat{S} \rangle + \langle \Delta\widehat{S}, \partial_{\tau_i} \partial_{\tau_{\nu'}} \widehat{S} \rangle + \langle \widehat{S} - \widehat{A}, \partial_{\tau_i} \partial_{\tau_{\nu'}} \Delta\widehat{S} \rangle \right\}, \tag{120}$$

$$[\mathbf{Y2}]_{\nu} = \Re \left\{ \langle \Delta\widehat{S}, \partial_{\tau_i} \Delta\widehat{S} \rangle \right\}, \tag{121}$$

noting that $[\mathbf{Y0}]$ is $\mathcal{O}(1)$, $[\mathbf{Y1}]$ is $\mathcal{O}(\Delta\widehat{F})$, and $[\mathbf{Y2}]$ is $\mathcal{O}(\Delta\widehat{F}^2)$, as expected.

Now, retaining both first- and second-order terms we see:

$$\sum_{\nu'} [\mathbf{x0}]_{\nu, \nu'} \cdot \Delta\tau_{\nu'} + [\mathbf{x1}]_{\nu} + \sum_{\nu', \nu''} \Delta\tau_{\nu''} [\mathbf{Y0}]_{\nu, \nu'', \nu'} \Delta\tau_{\nu'} + \sum_{\nu'} [\mathbf{Y1}]_{\nu, \nu'} \Delta\tau_{\nu'} + [\mathbf{Y2}]_{\nu} = \mathcal{O}(\Delta\widehat{F}^3) \quad \forall \tau_i \in \{\gamma, \beta, \alpha\}, \tag{122}$$

which, in matrix-vector notation can be written as:

$$[\mathbf{x0}]_{\nu, \cdot} \cdot \Delta\tau + [\mathbf{x1}] + \Delta\tau^{\top} \cdot [\mathbf{Y0}]_{\nu, \cdot, \cdot} \cdot \Delta\tau + [\mathbf{Y1}]_{\nu, \cdot} \cdot \Delta\tau + [\mathbf{Y2}]_{\nu} = \mathcal{O}(\Delta\widehat{F}^3) \quad \forall \tau_i \in \{\gamma, \beta, \alpha\}, \tag{123}$$

where $[\mathbf{x0}]_{\nu, \cdot}$ and $[\mathbf{Y1}]_{\nu, \cdot}$ are column-vectors produced by slicing $[\mathbf{x0}]$ and $[\mathbf{Y1}]$ respectively, while $[\mathbf{Y0}]_{\nu, \cdot, \cdot}$ is a matrix produced by slicing $[\mathbf{Y0}]$. This allows us to solve for $\Delta\tau$ up to second-order:

$$\Delta\tau_{\nu} = -[\mathbf{x0}]_{\nu, \cdot}^{-1} \cdot [\mathbf{x1}] - [\mathbf{x0}]_{\nu, \cdot}^{-1} \left\{ [\mathbf{x1}]^{\top} \cdot [\mathbf{x0}]^{-\dagger} \cdot [\mathbf{Y0}]_{\nu, \cdot, \cdot} \cdot [\mathbf{x0}]^{-1} \cdot [\mathbf{x1}] - [\mathbf{Y1}]_{\nu, \cdot} \cdot [\mathbf{x0}]^{-1} \cdot [\mathbf{x1}] + [\mathbf{Y2}]_{\nu} \right\} + \mathcal{O}(\Delta\widehat{F}^3). \tag{124}$$

Note that, in this expression, we can pick out the first-order solution:

$$[\Delta\mathbf{z1}]_{\nu} := -[\mathbf{x0}]_{\nu, \cdot}^{-1} \cdot [\mathbf{x1}] = \mathcal{O}(\Delta\widehat{F}),$$

as well as the second-order correction:

$$[\Delta^2 \mathbf{Z2}]_\iota := -[\mathbf{X0}]_{\iota'}^{-1} \left\{ [\mathbf{X1}]^\top \cdot [\mathbf{X0}]^{-\dagger} \cdot [\mathbf{Y0}]_{\iota, \iota'} \cdot [\mathbf{X0}]^{-1} \cdot [\mathbf{X1}] - [\mathbf{Y1}]_{\iota, \iota'} \cdot [\mathbf{X0}]^{-1} \cdot [\mathbf{X1}] + [\mathbf{Y2}]_{\iota} \right\} = \mathcal{O}(\Delta \widehat{F}^2).$$

Let's refer to the three components of the first- and second-order terms as the vectors $[\Delta \mathbf{Z1}]$ and $[\Delta^2 \mathbf{Z2}]$, respectively. Furthermore, let's re-introduce the image-index j , so that:

$$\Delta \tau_j = [\Delta \mathbf{Z1}]_j + [\Delta^2 \mathbf{Z2}]_j + \mathcal{O}(\Delta \widehat{F}^3),$$

where now we associate $\Delta \tau_j$ to the alignment of image- j , and note that each term in this expression has three components (corresponding to the index ι , which is no longer written).

With this notation we can now use our expression for the gradient in Eq. 101 and Eq. 102 to write:

$$\begin{aligned} \mathcal{L}(\{\widehat{A}_j\} | \widehat{F} + \Delta \widehat{F}; \{\tau_j + \Delta \tau_j\}) &= \mathcal{L} + G \cdot \Delta + \Delta^\dagger \cdot H \cdot \Delta \\ &= \mathcal{L} + \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j \\ &\quad + \Delta \widehat{F}^\dagger \cdot \partial_{\widehat{F}} \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \Delta \widehat{F}^\dagger \partial_{\widehat{F}} \partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j \\ &\quad + \sum_j \Delta \tau_j^\top \cdot \partial_{\tau_j} \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \partial_{\tau_j} \partial_{\tau_j} \mathcal{L} \cdot \Delta \tau_j \\ &= \mathcal{L} + \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \partial_{\tau_j} \mathcal{L} \cdot [\Delta \mathbf{Z1}]_j \\ &\quad + \sum_j \partial_{\tau_j} \mathcal{L} \cdot [\Delta^2 \mathbf{Z2}]_j + \Delta \widehat{F}^\dagger \cdot \partial_{\widehat{F}} \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \Delta \widehat{F}^\dagger \partial_{\widehat{F}} \partial_{\tau_j} \mathcal{L} \cdot [\Delta \mathbf{Z1}]_j \\ &\quad + \sum_j [\Delta \mathbf{Z1}]_j^\top \cdot \partial_{\tau_j} \partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F} + \sum_j \partial_{\tau_j} \partial_{\tau_j} \mathcal{L} \cdot [\Delta \mathbf{Z1}]_j. \end{aligned} \tag{125}$$

From this expression we can recover the corrections to the gradient and hessian resulting from the implicit relationship between \widehat{F} and the $\{\Delta \tau_j\}$.

As an example of the kinds of manipulations that are required for these corrections, we'll discuss the simplest term: $\sum_j \partial_{\tau_j} \mathcal{L} \cdot [\Delta \mathbf{Z1}]_j$, which serves as an addition to the original volumetric term $\partial_{\widehat{F}} \mathcal{L} \cdot \Delta \widehat{F}$ from the gradient (see Eq. 97).

When unpacked, this additional term can be represented as:

$$\sum_j \partial_{\tau_j} \mathcal{L} \cdot [\Delta \mathbf{Z1}]_j = \sum_j \sum_{\iota, \iota'} \Re \langle \partial_{\tau_j, \iota} \widehat{S}(\widehat{F}, \tau_j), \widehat{S}(\widehat{F}, \tau_j) - \widehat{A}_j \rangle [-\mathbf{X0}^{-1}]_{\iota, \iota'} \Re \langle \partial_{\tau_j, \iota'} \widehat{S}(\widehat{F}, \tau_j), \widehat{S}l \circ \mathcal{R}(\tau_j) \circ \Delta \widehat{F} \rangle \tag{126}$$

$$+ \sum_j \sum_{\iota, \iota'} \Re \langle \partial_{\tau_j, \iota} \widehat{S}(\widehat{F}, \tau_j), \widehat{S}(\widehat{F}, \tau_j) - \widehat{A}_j \rangle [-\mathbf{X0}^{-1}]_{\iota, \iota'} \Re \langle \widehat{S}(\widehat{F}, \tau_j) - \widehat{A}_j, \partial_{\tau_j, \iota'} \widehat{S}l \circ \mathcal{R}(\tau_j) \circ \Delta \widehat{F} \rangle. \tag{127}$$

In this expression the values of $[-\mathbf{X0}^{-1}]_{\iota, \iota'}$ are each just scalars that depend on j , ι and ι' . The same holds for other prefactors such as $\Re \langle \partial_{\tau_j, \iota} \widehat{S}(\widehat{F}, \tau_j), \widehat{S}(\widehat{F}, \tau_j) - \widehat{A}_j \rangle$. Thus, the first term on the right-hand-side (Eq. 126) can be turned into a volumetric integral using monopole-impulses via Eq. 63 and 64. On the other hand, the second term in Eq. 127 involves a derivative, and can be turned into a volumetric integral using dipole-impulses via Eq. 69 and 71. Further correction terms for the hessian involve interactions between dipoles, quadrupoles, etc., producing a multipole expansion.

22 Probing the spectrum of the hessian: case studies

In this section we'll describe one approach that allows us to crudely bound σ_H^{crit} (see Eq. 88) by appealing to the low-temperature limit of the likelihood. Specifically, we'll assume that the noise $\hat{\sigma}$ is sufficiently small that Eq. 54 is a good approximation to Eq. 48. We'll then try and find perturbations $\{\Delta \widehat{F}^{\text{zero}}, \Delta P^{\text{zero}}\}$ for which the leading-order term $\sigma_H^{\text{zero}}(\Delta \widehat{F}^{\text{zero}}, \Delta P^{\text{zero}})$ is small (in this low-temperature limit). These perturbations $\{\Delta \widehat{F}^{\text{zero}}, \Delta P^{\text{zero}}\}$ will then serve as candidates for $\{\Delta \widehat{F}^{\text{crit}}, \Delta P^{\text{crit}}\}$, from which we can bound σ_H^{crit}

By using this approach, we'll be able to investigate certain relationships between the viewing-angle distribution and the stability/robustness of single-particle reconstruction. By examining a few case-studies we'll see that, even if the data is sufficient for a single-particle reconstruction, non-uniformity in the viewing-angle distribution can compromise robustness. Ultimately, we'll recapitulate a well-known phenomenon: images from a variety of viewing-angles are required for a robust reconstruction.

To simplify the discussion, we'll first make several additional assumptions which can be relaxed later on:

Assume a large pool of synthetic-images: We'll assume that the pool of images is well modeled by the distribution of synthetic-images described in section 18. Furthermore, we'll neglect finite-size effects associated with the synthetic-image pool, and assume that J is sufficiently large that the distribution of synthetic-images is well-sampled.

Ignore translations and CTF: We'll assume that each synthetic-image is drawn with a displacement δ identically equal to 0 (i.e., $\sigma_\delta = 0$) and a CTF identically equal to 1.

We'll also suppress the notation for the coordinate \mathbf{k} , noting that the synthetic-images are simply

$$\widehat{A}_j^{\text{syn}} = \widehat{S}(\tau_j; \widehat{F}).$$

Let's now assume that we have a putative (local) minimum $\{\widehat{F}^{\text{opt}}, P(\tau | \widehat{F}^{\text{opt}})\}$ of the negative-log-likelihood from Eq. 54, which we'll denote by $-\log(P^{\text{opt}})$. This means that, for any particular image-index j , the corresponding synthetic-image $\widehat{A}_j^{\text{syn}} = \widehat{S}(\tau_j^{\text{opt}}; \widehat{F}^{\text{opt}})$ is associated with the maximum-likelihood alignment τ_j^{opt} via Eq. 53. Furthermore, we expect that the volume \widehat{F}^{opt} is consistent with these alignments in the sense that \widehat{F}^{opt} can be reconstructed from the synthetic-images using Eq. 73. Without loss of generality we can assume that the viewing-angles are described such that each $\gamma_j \equiv 0$, and hence each $\tau_j = (0, \beta_j, \alpha_j)$. With this notation the density of viewing-angles at a particular (β, α) is given by $P((0, \beta, \alpha) | \widehat{F}^{\text{opt}})$, or just $P(\tau)$ for short.

Note that the assumptions above imply that the maximum-likelihood alignments $\{\tau_j^{\text{opt}}\}$ are actually implicit functions of the $\{\widehat{A}_j^{\text{syn}}\}$ as well as \widehat{F}^{opt} ; if the images or volume were to change, so too would the optimal alignments. The alignments also technically depend on the prior distribution $P(\tau | \widehat{F}^{\text{opt}})$, but only in the sense that each τ_j^{opt} must lie in the support of $P(\tau | \widehat{F}^{\text{opt}})$. Thus, assuming the image-set is fixed and the prior distribution is allowed to be nonzero for all viewing-angles, any particular volumetric-perturbation will give rise to a perturbation in the maximum-likelihood alignments, and vice-versa. We'll take advantage of this implicit relationship in a moment: instead of directly searching for infinitesimal volumetric-perturbations which best preserve the likelihood, we can instead search for infinitesimal perturbations in the maximum-likelihood alignments which best preserve the image-template similarities shown in Eq. 45.

To be more specific, we'll denote by $\{\tau_j^{[0]}\}$ the 'old' maximum-likelihood viewing-angles $\{\tau_j^{\text{opt}}\}$ associated with $\widehat{F}^{[0]} := \widehat{F}^{\text{opt}}$, and define 'temporary' perturbed viewing-angles via:

$$\widehat{F}^{[0]} := \widehat{F}^{\text{opt}}, \quad (128)$$

$$\tau_j^{[0]} := \tau_j^{\text{opt}}, \quad (129)$$

$$\tau_j^{[1]} := \tau_j^{[0]} + \epsilon \Psi(\tau_j^{[0]}), \quad (130)$$

where $\epsilon \in \mathbb{R}$ is small, and $\Psi : SO3 \rightarrow \mathfrak{so3}$ is a flow-field on SO3. Using the $\{\tau_j^{[1]}\}$, we then consider a new back-propagation problem given by Eq. 73 to form a new volume $\widehat{F}^{[1]}$. To describe $\widehat{F}^{[1]}(\boldsymbol{\kappa})$ for any point $\boldsymbol{\kappa} = [\kappa_1, \kappa_2, 0]^T$, we can define $\perp(\boldsymbol{\kappa})$ to be the set of viewing-angles that correspond to rotations for which $\boldsymbol{\kappa}$ is stationary:

$$\perp(\boldsymbol{\kappa}) = \{\tau \mid \mathcal{R}\tau \circ \boldsymbol{\kappa} = \boldsymbol{\kappa}\}, \quad (131)$$

and then, noting that $\tau^{[0]} - \epsilon \Psi(\tau^{[0]})$ is the preimage of $\tau^{[0]}$ under the action of Ψ (to lowest order), we can write:

$$\widehat{F}^{[1]}(\boldsymbol{\kappa}) = \frac{1}{Z(\boldsymbol{\kappa})} \int_{\tau \in \perp(\boldsymbol{\kappa})} \widehat{S}(\mathbf{k}; \widehat{F}^{[0]}, \tau - \epsilon \Psi(\tau)) P(\tau - \epsilon \Psi(\tau)) d\tau \quad \text{where: } \mathbf{k} = \mathcal{R}(\tau) \circ \boldsymbol{\kappa} = [\kappa_1, \kappa_2]^T, \quad (132)$$

$$= \frac{1}{Z(\boldsymbol{\kappa})} \int_{\tau \in \perp(\boldsymbol{\kappa})} \left[\widehat{S} \circ \mathcal{R}(\tau - \epsilon \Psi(\tau)) \circ \widehat{F}^{[0]} \right](\mathbf{k}) P(\tau - \epsilon \Psi(\tau)) d\tau \quad \text{with:} \quad (133)$$

$$Z(\boldsymbol{\kappa}) = \int_{\tau \in \perp(\boldsymbol{\kappa})} P(\tau - \epsilon \Psi(\tau)) d\tau. \quad (134)$$

An analogous equation holds (with a rotated frame of reference) for any general $\boldsymbol{\kappa}$. In words: $\widehat{F}^{[1]}(\boldsymbol{\kappa})$ at any point $\boldsymbol{\kappa}$ can be estimated by simply averaging over all the values of $\widehat{F}^{[0]}$ which are mapped onto $\boldsymbol{\kappa}$ by the viewing-angle perturbations. In the case studies below, we'll choose Ψ to be simple enough that we can estimate the back-propagated $\widehat{F}^{[1]}$ 'by hand'.

Obviously the volumetric-perturbation producing $\widehat{F}^{[1]}$ then induces an alignment-perturbation $\{\tau_j^{[2]}\}$:

$$\tau_j^{[2]} = \arg \max_{\tau} P(\widehat{A}_j^{\text{syn}} \mid \tau; \widehat{F}^{[1]}) \quad (135)$$

$$= \arg \max_{\tau} P(\widehat{S}(\tau_j^{[0]}; \widehat{F}^{[0]}) \mid \tau; \widehat{F}^{[1]}). \quad (136)$$

when then induce another volumetric-perturbation $\widehat{F}^{[2]}$, and so forth. This sequence (or iteration) may eventually converge, but we'll curtail it after the first few steps, typically using $\widehat{F}^{[1]}$ and $\{\tau_j^{[1]}\}$ as an estimate for the perturbed volume and alignments. We'll define $\Delta\widehat{F} := \widehat{F}^{[1]} - \widehat{F}^{[0]}$ and $\Delta\tau_j := \tau_j^{[1]} - \tau_j^{[0]}$. Finally, we'll use these candidate $\Delta\widehat{F}$ and $\{\Delta\tau_j\}$ to measure the rayleigh-product $\sigma_H^{zero}(\Delta\widehat{F}, \{\Delta\tau_j\})$.

The relevant (low-temperature) likelihood is:

$$-\log \left(P(\{\widehat{A}_j\} | \widehat{F}^{[1]}, \{\tau_j^{[1]}\}) \right) \sim \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbf{k} \in \Omega_{\mathcal{K}} \in \mathbb{R}^2} \left| \widehat{A}_j(\mathbf{k}) - \widehat{S}(\mathbf{k}; \tau_j^{[1]}; \widehat{F}^{[1]}) \right|^2 d\mathbf{k} \quad (137)$$

$$= \frac{1}{2\hat{\sigma}^2} \sum_{j=1}^J \iint_{\mathbf{k} \in \Omega_{\mathcal{K}} \in \mathbb{R}^2} \left| \widehat{S}(\mathbf{k}; \tau_j^{[0]}; \widehat{F}^{[0]}) - \widehat{S}(\mathbf{k}; \tau_j^{[1]}; \widehat{F}^{[1]}) \right|^2 d\mathbf{k}. \quad (138)$$

22.1 Rotation Blur:

One of the most straightforward perturbations to analyze is when $\Psi(\tau)$ corresponds to a rotation-blur. To describe this scenario, let's first assume that $P(\tau)$ is uniform. Now let's consider a particular $\tau^{[0]}$, for which the template $\widehat{S}(\mathbf{k}; \tau^0)$ has a pole $\boldsymbol{\kappa}_0 := \mathcal{R}^{-1}\tau^{[0]} \circ [0, 0, 1]^\top$. Now let's define $\mathcal{N}(\tau^{[0]}, \epsilon^2)$ to be an isotropic gaussian-distribution on the surface of the sphere with a mean at $\tau^{[0]}$ and a variance of ϵ^2 . Finally, let's define $\Psi(\tau^{[0]})$ so that $\tau^{[1]} := \tau^{[0]} + \epsilon\Psi(\tau^{[0]})$ is drawn independently from the distribution $\mathcal{N}(\tau^{[0]}, \epsilon^2)$.

In this setting the back-propagated solution $\widehat{F}^{[1]}$ on any k -shell can be constructed by simply convolving the values of $\widehat{F}^{[0]}$ on that shell with $\mathcal{N}(\mathbf{0}, \epsilon^2)$. That is to say, the volume $\widehat{F}^{[1]}$ can be constructed by 'rotationally blurring' the volume $\widehat{F}^{[0]}$: i.e.:

$$\widehat{F}^{[1]}(k, \widehat{\boldsymbol{\kappa}}) = \int_{\widehat{\boldsymbol{\kappa}}'} \widehat{F}^{[0]}(k, \widehat{\boldsymbol{\kappa}}') \mathcal{N}(\widehat{\boldsymbol{\kappa}} - \widehat{\boldsymbol{\kappa}}', \epsilon^2) d\widehat{\boldsymbol{\kappa}}'. \quad (139)$$

This rotational-blur corresponds directly to an isotropic surface-diffusion of the volume $\widehat{F}^{[0]}$:

$$\left[\widehat{F}^{[1]} \right]_l^m = \exp(-l(l+1)\epsilon^2/2) \cdot \left[\widehat{F}^{[0]} \right]_l^m. \quad (140)$$

22.2 Latitudinal perturbation:

Another relatively simple perturbation involves a fully latitudinal perturbation: $\Psi(\tau) = (0, \Theta(\beta), 0) = (0, \sin(2\beta), 0)$. If $\epsilon > 0$, this perturbation pushes each τ in the northern-hemisphere south (towards the equator), while pushing each τ in the southern-hemisphere north (again towards the equator).

If $P(\tau)$ is uniform, then in the limit as $\epsilon \rightarrow 0$ each value of $\widehat{F}^{[1]}(k, \widehat{\boldsymbol{\kappa}})$ is determined by averaging the values of $\widehat{F}^{[0]}(k, \widehat{\boldsymbol{\kappa}}')$ on a small 'circle' (on the surface of the sphere). This circle can be described in terms of the polar-angle a associated with $\widehat{\boldsymbol{\kappa}}$. the circle is centered at a polar-angle of $a_c(a)$ with a radius (extending on the surface of the sphere of radius-1) of $r(a)$.

$$a_c(a) = a + 0.5\epsilon\Theta(a) \quad \text{and: } r(a) = 0.5\epsilon\Theta(a). \quad (141)$$

Note that the radius of this circle vanishes as $\boldsymbol{\kappa}$ approaches either the poles or the equator, and is largest in the temperate zones.

As a result of Eq. 141 the values of $\widehat{F}^{[1]}$ can each be obtained by integrating the nearby values of $\widehat{F}^{[0]}$ against a small local gaussian:

$$\widehat{F}^{[1]}(k, \widehat{\boldsymbol{\kappa}}) = \int_{\widehat{\boldsymbol{\kappa}}'} \widehat{F}^{[0]}(k, \widehat{\boldsymbol{\kappa}}') \mathcal{N}(\widehat{\boldsymbol{\kappa}} - \widehat{\boldsymbol{\kappa}}', \epsilon^2 \sigma^2(\widehat{\boldsymbol{\kappa}})) d\widehat{\boldsymbol{\kappa}}' \quad \text{with: } \sigma^2(\widehat{\boldsymbol{\kappa}}) = \frac{1}{4} \cdot \frac{1}{2} \Theta(a)^2. \quad (142)$$

By looking at this case-study we see that a latitudinal alignment-perturbation $\Theta(\beta)$ results in a volumetric-perturbation $\Delta\widehat{F}$ which is well approximated by a a -dependent local-diffusion about each point $\boldsymbol{\kappa}$. This local diffusion-coefficient drops to 0 at the poles and equator (i.e., $a = 0, \pi/2$ and π). Furthermore, the latitudinal alignment-perturbation also shifts each value of $\widehat{F}^{[0]}$ in the latitudinal-direction, allowing the templates values to be unchanged to first-order if the viewing-angle is also perturbed latitudinally. Thus, the template-perturbation

$$\Delta\widehat{S}(\mathbf{k}; \tau) = \widehat{S}(\mathbf{k}; \widehat{F}^{[1]}; \tau + 0.5\epsilon\Theta(\beta)) - \widehat{S}(\mathbf{k}; \widehat{F}^{[0]}; \tau)$$

will vanish as τ approaches the poles. Consequently, any image with an initial viewing-angle near the poles will align almost as well to $\widehat{F}^{[1]}$ (at the latitudinally-perturbed $\tau^{[1]}$) as it did to $\widehat{F}^{[0]}$ (at its initial $\tau^{[0]}$).

Taking these considerations into account, we see that a latitudinal alignment-perturbation like $\Theta(\beta)$ will introduce a likelihood-penalty to images that are oriented far from the poles, but will not introduce a likelihood-penalty to images that are oriented close to the poles. As a result, any volume constructed from a distribution of viewing-angles concentrated near the poles will be especially susceptible to latitudinal-perturbations of this form. To illustrate this phenomenon we can consider a non-uniform $P(\tau)$ which is concentrated around the polar-caps. As $P(\tau)$ becomes more and more concentrated around the polar-caps, the likelihood of $\{\widehat{F}^{[1]}, \{\tau_j^{[1]}\}\}$ becomes more and more similar to $\{\widehat{F}^{[0]}, \{\tau_j^{[0]}\}\}$, even though $\widehat{F}^{[1]}$ may not be similar to $\widehat{F}^{[0]}$. Equivalently, the average fisher-information per image with-respect to the perturbation $\{\Delta\widehat{F}, \{\Delta\tau_j\}\}$ vanishes as $P(\tau)$ becomes more concentrated near the poles.

Put another way, if the image-pool doesn't have many images that are oriented away from the poles, no amount of near-polar images will help improve the robustness of the reconstructed volume. As mentioned above, this intuitive result aligns with many current methodologies: many images from a variety of viewing-angles are required for a robust reconstruction.




Figure 24: In this figure we illustrate some properties of a latitudinal-perturbation $\Theta(\beta) = \sin(2\beta)$. On the right we illustrate the local diffusion-coefficient $\mathcal{D}(a)$ (vertical) as a function of the polar-angle a (horizontal) when the viewing-angle distribution is uniform. Note that the local diffusion-coefficient $\mathcal{D}(a)$ drops to 0 when a lies on either the equator or the poles. This means that templates of the original volume F which correspond to slices near the equatorial-zone (i.e., with viewing-angles close to the poles) will align almost as well to the perturbed volume $F + \Delta F$ as they did to the original volume F (after latitudinally adjusting each template's viewing-angle τ by $\Delta\tau(\tau)$). On the left we illustrate the same local-diffusion-coefficient $\mathcal{D}(a; \sigma_\beta)$ as a function of the polar-angle a (vertical) as the viewing-angle distribution becomes more concentrated around the poles (as parametrized by σ_β). The viewing-angle distribution considered for this figure is of the form $\rho(\beta) \propto \exp\left(-\sin(\beta)^2/(2\sigma_\beta^2)\right)$, with σ_β shown along the horizontal. The value of the diffusion-coefficient $\mathcal{D}(a; \sigma_\beta)$ is shown as a color in this array, ranging from cyan (zero) to pink (the maximum shown on the left). Note that, as σ_β shrinks the viewing-angle distribution $\rho(\beta)$ becomes more and more concentrated towards the poles. Such a ‘polar-cap’ distribution introduces essentially no local-diffusion for values of a in a large band around the equator. This means that, as the viewing-angle distribution becomes concentrated near the poles, the range of templates of the original volume F which can be well-aligned to the perturbed-volume $F + \Delta F$ increases to include templates with original-viewing-angles in the tropical- and temperate-zones.

22.3 Longitudinal perturbation:

A slightly more elaborate perturbation involves $\Psi(\tau) := \Phi(\alpha) = \sin(2\alpha)$. If $\epsilon > 0$, this perturbation pushes each τ either east- or west-ward away from the great-circle corresponding to the prime-meridian.

If $P(\tau)$ is uniform, then in the limit as $\epsilon \rightarrow 0$ each value of $\widehat{F}^{[1]}(k, \widehat{\kappa})$ is determined by averaging the values of $\widehat{F}^{[0]}(k, \widehat{\kappa}')$ on a small ‘collapsed circle’ (on the surface of the sphere) projected onto (i.e., extending only in) the azimuthal-direction.

This collapsed circle can be described in terms of the polar-angle and azimuthal-angle (a, b) associated with the point $\hat{\boldsymbol{\kappa}}$. The collapsed circle is centered at an azimuthal-angle of $b_c(a, b)$ and has an azimuthal-radius of $r(a)$.

$$b_c(a, b) = b + \epsilon u(a)\Phi(b) \quad \text{and: } r^2(a) = \epsilon^2 [1 + u(a)], \quad (143)$$

$$u(a) = \frac{1 - |\cos(a)|}{1 + |\cos(a)|} \quad \text{and: } \Phi(b) = \sin(2b). \quad (144)$$

Note that the azimuthal-radius of this collapsed-circle vanishes as $\boldsymbol{\kappa}$ approaches the equator, and is largest near the poles.

As a result of Eq. 144 the values of $\hat{F}^{[1]}$ can each be obtained by integrating the nearby values of $\hat{F}^{[0]}$ against a small local gaussian extending in the azimuthal-direction:

$$\hat{F}^{[1]}(k, \hat{\boldsymbol{\kappa}}) = \int_b \hat{F}^{[0]}(k, \hat{\boldsymbol{\kappa}}') \mathcal{N}(b - b', \epsilon^2 \sigma^2(a)) d\hat{\boldsymbol{\kappa}}' \quad \text{with: } \sigma^2(a) = \frac{1}{2} r^2(a). \quad (145)$$

By looking at this case-study we see that a longitudinal alignment-perturbation $\Phi(\alpha)$ results in a volumetric-perturbation $\Delta \hat{F}$ which is well approximated by a b -dependent local azimuthal-diffusion about each point $\boldsymbol{\kappa}$. This local azimuthal-diffusion-coefficient drops to 0 at the equator (i.e., $a = \pi/2$). However, the longitudinal-perturbation induces an azimuthal-shift in each value of $\hat{F}^{[0]}$ that is both a - and b -dependent. Thus, the original synthetic-images (i.e., templates from $\hat{F}^{[0]}$) won't typically align well to the volume $\hat{F}^{[1]}$. Taking these considerations into account, we see that a longitudinal alignment-perturbation like $\Phi(\alpha)$ will introduce likelihood-penalties for all the images.

However, the results above are strongly dependent on the (near) uniformity of the viewing-angle distribution $P(\tau)$. If $P(\tau)$ were to be concentrated near the equatorial-belt, then the results change substantially. In fact, as $P(\tau)$ becomes more concentrated around the equatorial-belt there are two major changes:

1. The azimuthal-displacement $b_c(a, b)$ becomes dependent mainly on the azimuthal-angle b of $\boldsymbol{\kappa}$ (becoming proportional to $\Phi(b)$) and becomes more and more independent of the polar-angle a .
2. the azimuthal-variance $r^2(a)/2$ drops to zero away from the poles.

In this scenario (i.e., an equatorially-concentrated $P(\tau)$), the situation is quite similar to the poorly-posed scenario involving only equatorial viewing-angles. That is to say, any volume constructed from a distribution of viewing-angles concentrated near the equator will be especially susceptible to longitudinal-perturbations.

To illustrate this phenomenon we can consider a non-uniform $P(\tau)$ which is concentrated around the equator. Similar to the latitudinal-perturbation described above, as $P(\tau)$ becomes more and more concentrated around the equator, the likelihood of $\{\hat{F}^{[1]}, \{\tau_j^{[1]}\}\}$ becomes more and more similar to $\{\hat{F}^{[0]}, \{\tau_j^{[0]}\}\}$, and the average fisher-information per image with-respect to the perturbation $\{\Delta \hat{F}, \{\Delta \tau_j\}\}$ vanishes. Put another way, if the image-pool only has images that have viewing-angles near the equator, their relative orientation cannot be pinned down. Once again, this result echoes conventional-wisdom: many images from a variety of viewing-angles are required for a robust reconstruction.




Figure 25: In this figure we illustrate some properties of a longitudinal-perturbation $\Phi(\alpha) = \sin(2\alpha)$. In this figure we show multiple subplots. Within each subplot (indexed by σ_β) we show a heatmap of the local-azimuthal-diffusion $\mathcal{E}(a, b; \sigma_\beta)$ as a function of the polar- and azimuthal angles (a, b) , corresponding to the vertical- and horizontal-axes respectively. Within each subplot the value of $\mathcal{E}(a, b; \sigma_\beta)$ ranges from cyan (zero) to pink (maximum). The different subplots correspond to different viewing-angle distributions (parametrized by σ_β). The viewing-angle distribution considered for this figure is of the form $\rho(\beta) \propto \exp\left(-\cos(\beta)^2/(2\sigma_\beta^2)\right)$, with σ_β shown in the title of each subplot. In this case, when σ_β is large the viewing-angle distribution is essentially uniform (lower-right subplot). On the other hand, when σ_β becomes small the viewing-angle distribution concentrates around the equator (upper-left subplot). Note that, as σ_β shrinks and the viewing-angle distribution concentrates around the equator, the azimuthal-diffusion $\mathcal{E}(a, b; \sigma_\beta)$ drops close to zero for large regions on the sphere. Thus, for these equatorially-concentrated viewing-angle-distributions, templates of the original volume F will align almost as well to the perturbed volume $F + \Delta F$ as they did to the original volume F after longitudinally adjusting each template's viewing-angle τ by $\Delta\tau(\tau)$.

22.4 Relating the laplacian to a local average:

In the above calculations we've used the following fact. Given a twice-differentiable function $u(\mathbf{x}) : \mathbb{R}^2 \rightarrow \mathbb{C}$, we can use the laplacian Δu to approximate the average of that function along the perimeter ∂B_r of a small disk B_r of radius r .

$$I(r) = \frac{1}{|\partial B_r|} \int_{B_r} u(\mathbf{x}) d\Gamma.$$

$$\partial_r I(r) = \frac{1}{|\partial B_r|} \int_{B_r} \partial_{x_1}^2 u(\mathbf{x}) + \partial_{x_2}^2 u(\mathbf{x}) d\mathbf{x} = \frac{1}{|\partial B_r|} \int_{B_r} \Delta u(\mathbf{x}) d\mathbf{x}.$$

$$\partial_r I(r) \approx \frac{1}{|\partial B_r|} \pi r^2 \Delta u(0) \approx \frac{1}{2} r \Delta u(0).$$

$$I(r) = I(0) + \int_0^r \partial_r I(s) ds \approx u(0) + \frac{1}{4} r^2 \Delta u(0).$$

We can also describe a similar relationship: if we were to define D_r to be the projection of B_r onto the x_1 axis, then:

$$I(r) = \frac{1}{|\partial D_r|} \int_{D_r} u(x_1) d\Gamma \approx u(0) + \frac{1}{4} r^2 \partial_{x_1}^2 u(0),$$

which is what one would expect if u did not vary (or only varied linearly) orthogonal to the x_1 direction (and hence $\partial_{x_2}^2 u \equiv 0$).

22.5 Notes for the longitudinal-perturbation:

For the longitudinal-perturbation we set $\Theta(\alpha) = 2 \sin(2\alpha)$.

Let's pick a particular point on the sphere, and describe it via the polar- and azimuthal-angles (a, b) . The value of the third euler-angle c can be used to parametrize the poles associated with that point. Now all the pole-locations for (a, b) are given by the same formula used to determine the templates:

$$(x, y, z) = [\cos(b) \cos(a) \cos(c) - \sin(b) \sin(c), \sin(b) \cos(a) \cos(c) + \cos(b) \sin(c), -\sin(a) \cos(c)] \quad (146)$$

The norm $(x^2 + y^2) = w^2$ is given by:

$$w^2 = ((\cos(b) \cos(a) \cos(c) - \sin(b) \sin(c))^2 + (\sin(b) \cos(a) \cos(c) + \cos(b) \sin(c))^2) \quad (147)$$

$$= (\cos(b) \cos(a) \cos(c))^2 + (\sin(b) \sin(c))^2 - 2(\cos(b) \cos(a) \cos(c) \sin(b) \sin(c)) \quad (148)$$

$$+ (\sin(b) \cos(a) \cos(c))^2 + (\cos(b) \sin(c))^2 + 2(\sin(b) \cos(a) \cos(c) \cos(b) \sin(c)) \quad (149)$$

$$= ((\cos(b) \cos(a) \cos(c))^2 + (\sin(b) \sin(c))^2 + (\sin(b) \cos(a) \cos(c))^2 + (\cos(b) \sin(c))^2) \quad (150)$$

$$= (\cos(a) \cos(c))^2 + (\sin(c))^2. \quad (151)$$

The polar-angle at (x, y, z) is:

$$\arctan(y/x) = \arctan\left(\frac{\sin(b) \cos(a) \cos(c) + \cos(b) \sin(c)}{\cos(b) \cos(a) \cos(c) - \sin(b) \sin(c)}\right) =: \phi(a, b, c). \quad (152)$$

Note that, obviously, as any particular pole is perturbed azimuthally, the associated point on the sphere (associated with that pole) is perturbed azimuthally by the same amount. Thus, for any pole ϕ , the preimage $\phi^{\text{old}} = \phi^{\text{old}} - \epsilon\Theta(\phi^{\text{old}})$, or, more simply: $\phi^{\text{new}} - \phi^{\text{old}} = -\epsilon \sin(2\phi^{\text{old}})$, which is the azimuthal change in the point induced by the azimuthal change in the pole.

Now we can estimate the average shift Δb for any point on the sphere by averaging over c (parametrizing the poles associated with that point).

$$\int_0^{2\pi} -\epsilon \sin(2\phi) dc = -\epsilon \int_0^{2\pi} 2 \sin(\phi) \cos(\phi) dc = -\epsilon \int_0^{2\pi} 2 \frac{xy}{w^2} dc \quad (153)$$

$$= -\epsilon \int_0^{2\pi} 2 \frac{(\sin(b) \cos(a) \cos(c) + \cos(b) \sin(c))(\cos(b) \cos(a) \cos(c) - \sin(b) \sin(c))}{(\cos(a) \cos(c))^2 + (\sin(c))^2} dc. \quad (154)$$

Now seeking a separable solution we fix the b at, say, $\pi/4$ to get:

$$= -2\epsilon \int_0^{2\pi} \frac{(\sqrt{0.5} \cos(a) \cos(c) + \sqrt{0.5} \sin(c))(\sqrt{0.5} \cos(a) \cos(c) - \sqrt{0.5} \sin(c))}{(\cos(a) \cos(c))^2 + (\sin(c))^2} dc \quad (155)$$

$$= -\epsilon \int_0^{2\pi} \frac{(\cos(a) \cos(c))^2 - (\sin(c))^2}{(\cos(a) \cos(c))^2 + (\sin(c))^2} dc \quad (156)$$

$$= -\epsilon \int_0^{2\pi} \frac{(\cos(a))^2 (\cos(c))^2 - (\sin(c))^2}{(\cos(a))^2 (\cos(c))^2 + (\sin(c))^2} dc. \quad (157)$$

Now using Gradshteyn & Ryzhik: 3.647 (p402) BI (47)(20) we see: ;

$$\begin{aligned} \int_0^{\pi/2} \frac{\cos(x)^p \cos(px)}{a^2 \sin(x)^2 + b^2 \cos(x)^2} dx &= \frac{\pi}{2b} \frac{a^{p-1}}{(a+b)^p} \\ \frac{1}{2\pi} \int_0^{2\pi} \frac{\cos(x)^2}{\sin(x)^2 + b^2 \cos(x)^2} dx &= \frac{1}{b(1+b)} \\ + \int_0^{2\pi} \frac{\alpha^2 \cos(c)^2}{\alpha^2 \cos(c)^2 + \sin(c)^2} dc &= \frac{\alpha^2}{\alpha(1+\alpha)} = -\frac{\alpha}{1+\alpha} \\ + \int_0^{2\pi} \frac{\sin(c)^2}{\alpha^2 \cos(c)^2 + \sin(c)^2} dc &= +\frac{1}{\alpha(1+\alpha)} - \frac{1}{\alpha} \\ + \int_0^{2\pi} \frac{\alpha^2 \cos(c)^2 - \sin(c)^2}{\alpha^2 \cos(c)^2 + \sin(c)^2} dc &= +\frac{\alpha}{1+\alpha} + \frac{1}{\alpha(1+\alpha)} - \frac{1}{\alpha} \quad (158) \\ &= \frac{\alpha^2 - \alpha}{\alpha^2 + \alpha} = \frac{\alpha - 1}{\alpha + 1}. \quad (159) \end{aligned}$$

Note that, the equations above hold for $\alpha > 0$. Thus, in practice we'll need absolute-values in the equations above: i.e., $(|\alpha| - 1)/(|\alpha| + 1)$.

Aside: as we perturb the point-polar-angle a , the pole-azimuthal-angle ϕ changes as $\partial_a \phi$. ; Note that:

$$\partial \arctan(y/x) = \frac{1}{1 + \frac{y^2}{x^2}} \cdot \left(\frac{y'}{x} - \frac{yx'}{x^2} \right) = \frac{x^2}{x^2 + y^2} \cdot \frac{y'x - yx'}{x^2} = \frac{y'x - yx'}{x^2 + y^2}.$$

Thus,

$$\partial_a \phi = \frac{(-\sin(b) \sin(a) \cos(c))(\partial_b \cos(a) \cos(c) - \sin(b) \sin(c)) - (\sin(b) \cos(a) \cos(c) + \cos(b) \sin(c))(-\cos(b) \sin(a) \cos(c))}{(\cos(a) \cos(c))^2 + (\sin(c))^2} \quad (160)$$

$$= \frac{(\cos(c))^2(-\sin(b) \sin(a) \cos(b) \cos(a)) + (\sin(b))^2(\sin(a) \cos(c) \sin(c))}{(\cos(a) \cos(c))^2 + (\sin(c))^2} \quad (161)$$

$$+ \frac{(\cos(c))^2 + (\cos(c))^2 + \sin(b) \cos(a) \cos(b) \sin(a) + (\cos(b))^2(\sin(a) \cos(c) \sin(c))}{(\cos(a) \cos(c))^2 + (\sin(c))^2} \quad (162)$$

$$= \frac{\sin(a) \cos(c) \sin(c)}{(\cos(a) \cos(c))^2 + (\sin(c))^2}. \quad (163)$$

23 The Noise-marginalized limit

In section 13 we discussed the low-temperature limit $\hat{\sigma} \rightarrow 0$, for which the likelihood in Eq. 48 can be greatly simplified, producing Eq. 54. We leverage this simplified expression to produce Eq. 73, which affords certain computational and conceptual advantages.

A similar simplification can occur when the likelihood is marginalized over the noise $\hat{\sigma}$. To demonstrate this, we can return to Eq. 31 and rewrite the likelihood of observing an image \hat{A} , given a particular template \hat{S} and noise-level $\hat{\sigma}$:

$$P(\hat{A} | \hat{S}, \hat{\sigma}) := \frac{1}{(\sqrt{2\pi} \hat{\sigma})^{\text{dof}}} \exp\left(-\frac{1}{\hat{\sigma}^2} \ell(\hat{A} | \hat{S})\right), \quad (164)$$

where $\ell(\hat{A} | \hat{S})$ takes the form:

$$\ell(\hat{A} | \hat{S}) := \frac{1}{2} \iint_{\mathbf{k} \in \Omega_K} |\hat{A}(\mathbf{k}) - \hat{S}(\mathbf{k})|^2 d\mathbf{k}, \quad (165)$$

and dof refers to the number of independent degrees of freedom in the image.

Now we can use Gradshteyn and Ryzhik 3.326.2:

$$\int_0^{+\infty} \frac{1}{\lambda^N} \exp\left(-\frac{\ell}{\lambda^2}\right) d\lambda = \frac{1}{2} \Gamma\left(\frac{N-1}{2}\right) \frac{1}{\ell^{(N-1)/2}} = \frac{1}{2} \Gamma\left(\frac{N-1}{2}\right) \cdot \exp\left(-\frac{N-1}{2} \log \ell\right),$$

and marginalize over the noise $\hat{\sigma}$ to produce:

$$P(\hat{A} | \hat{S}, \hat{\sigma}) = \int_0^{+\infty} P(\hat{A} | \hat{S}, \hat{\sigma}) d\hat{\sigma} \quad (166)$$

$$= \frac{1}{(2\pi)^{\text{dof}/2}} \cdot \frac{1}{2} \Gamma(\text{dof}') \cdot \exp\left(-\text{dof}' \log \ell(\hat{A} | \hat{S})\right), \quad (167)$$

where we have denoted $\hat{\sigma}$ -marginalization via the strikethrough $\hat{\sigma}$, and set $\text{dof}' = (\text{dof} - 1)/2$.

Note that dof' is often quite large in practice, and thus this marginalized likelihood is typically very well approximated using standard asymptotics (e.g., laplace's approximation). In this regime we can marginalize Eq. 47 over $\hat{\sigma}$ and write:

$$P(\hat{A} | \hat{F}, \hat{\sigma}) = \int_{\boldsymbol{\delta} \in \mathbb{R}^2} \int_{\tau \in SO3} P(\hat{A} | \tau; \boldsymbol{\delta}; \hat{F}; \hat{\sigma}) \cdot P(\boldsymbol{\delta} | \tau; \hat{F}) \cdot P(\tau | \hat{F}) d\tau d\boldsymbol{\delta}, \quad (168)$$

where $P(\hat{A} | \tau; \boldsymbol{\delta}; \hat{F}; \hat{\sigma})$ now represents a marginalized likelihood similar to Eq. 167. In the large-dof limit we now have a statement very similar to Eq. 53. Namely, as $\text{dof} \rightarrow \infty$, the integral in Eq. 168 becomes dominated by the contribution associated with the 'maximum-likelihood' alignment-angle for the image \hat{A} . Just as before, this maximum-likelihood alignment occurs for the viewing-angle, displacement and volume which maximize the summand $P(\hat{A} | \tau; \boldsymbol{\delta}; \hat{F}; \hat{\sigma}) \cdot P(\tau, \boldsymbol{\delta} | \hat{F})$ in the integrand of Eq. 168.

In the context of single-particle reconstruction, the negative-log-likelihood for a collection of images in the large-dof limit is now dominated by the following leading-order expansion:

$$-\log \left(P \left(\{\widehat{A}_j\} \mid \widehat{F}, \widehat{\sigma} \right) \right) \sim +\text{dof}' \cdot \sum_{j=1}^J \log \left(\ell \left(\widehat{A}_j \mid \widehat{S}_j^{\text{opt}} \right) \right), \quad (169)$$

with:

$$\widehat{S}_j^{\text{opt}}(\mathbf{k}) := \widehat{S} \left(\mathbf{k}; \tau_j^{\text{opt}}; \delta_j^{\text{opt}}; CTF_j; \widehat{F} \right).$$

23.1 marginalizing with respect to offset μ :

In practice it is often convenient to take into consideration an image-specific real-space offset μ :

$$P(\widehat{A} \mid \widehat{S}, \widehat{\sigma}, \mu) := \frac{1}{(\sqrt{2\pi} \widehat{\sigma})^{\text{dof}}} \exp \left(-\frac{1}{\widehat{\sigma}^2} \ell(\widehat{A} \mid \widehat{S}, \mu) \right), \quad (170)$$

where $\ell(\widehat{A} \mid \widehat{S}, \mu)$ takes the form:

$$\ell(\widehat{A} \mid \widehat{S}, \mu) := \frac{1}{2} \iint_{\mathbf{x} \in \Omega_X} |A(\mathbf{x}) + \mu - S(\mathbf{x})|^2 d\mathbf{x}. \quad (171)$$

To marginalize this likelihood with respect to the image-offset μ we can use Gradshteyn and Ryzhik 2.33:

$$\int_{-\infty}^{+\infty} \exp \left(-\frac{ax^2 \pm 2bx + c}{\lambda} \right) dx = \sqrt{\frac{\lambda\pi}{a}} \exp \left(-\frac{ac - b^2}{\lambda a} \right),$$

which holds for any normalizing factor λ .

Using this relationship, we can write:

$$P(\widehat{A} \mid \widehat{S}, \widehat{\sigma}, \#) = \int_{-\infty}^{+\infty} P(\widehat{A} \mid \widehat{S}, \widehat{\sigma}, \mu) d\mu \quad (172)$$

$$= \frac{1}{(2\pi)^{\text{dof}/2}} \cdot \frac{1}{\widehat{\sigma}^{\text{dof}}} \cdot (2\pi\widehat{\sigma}^2) \cdot \exp \left(-\frac{1}{\widehat{\sigma}^2} \ell(\widehat{A} \mid \widehat{S}, \#) \right), \quad (173)$$

$$\text{with: } \ell(\widehat{A} \mid \widehat{S}, \#) := \frac{1}{2} \iint_{\mathbf{x} \in \Omega_X} |(A(\mathbf{x}) - \bar{A}) - (S(\mathbf{x}) - \bar{S})|^2 d\mathbf{x}, \quad (174)$$

where we have used \bar{A} and \bar{S} to denote the average image- and template-value.

Given the structure of $\ell(\widehat{A} \mid \widehat{S}, \#)$, we can assume without loss of generality that the images are all ‘centered’ (i.e., with mean 0). Similarly, we can assume that the volume F is also mean 0, and hence every template S is centered as well.

After performing this offset-marginalization we can once again marginalize over the noise to obtain the negative-log-likelihood:

$$-\log \left(P \left(\{\widehat{A}_j\} \mid \widehat{F}, \widehat{\sigma}, \# \right) \right) \sim +\text{dof}' \cdot \sum_{j=1}^J \log \left(\ell \left(\widehat{A}_j \mid \widehat{S}_j^{\text{opt}}, \# \right) \right). \quad (175)$$

In many practical scenarios the values of $\ell(\widehat{A}_j \mid \widehat{S}_j^{\text{opt}}, \#)$ are rather tightly distributed across the image-pool j , and it is not unreasonable to approximate the function $\log(\ell)$ as an affine function $\log(\ell) \approx a\ell + b$ for some fixed a and b (i.e., constants across the image-pool). One such example is shown in Fig 26.




Figure 26: In this figure we illustrate values of $\ell(\widehat{A}_j | \widehat{S}_j^{\text{opt}}, \#)$ taken from micrographs of FIB-milled yeast lamellae as described in [46] (see Fig 1 and Fig 15 of [46]). These micrographs have been scanned for picked particle images that resemble mature 60S ribosomes. For each of these picked particle images A_j , we measure the optimal pose-parameters for the 60S ribosome, and calculate the corresponding projection $\widehat{S}_j^{\text{opt}}$, as well as the function $\ell_j := \ell(\widehat{A}_j | \widehat{S}_j^{\text{opt}}, \#)$. On the far left we show the cumulative-distribution function of $\{\ell_j\}$ for the original set of picked-particles (bright red). We also show the corresponding cumulative-distribution-functions for the particles that were detected after adding various amounts of gaussian-noise to the original micrograph. The noise added to the original micrograph is quantified by the additional variance per-pixel (relative to the original noise-level), which ranges across the values $\{0.0, 0.1, 0.2, 0.5, 1.0, 1.5, 2.0, 2.5\}$ (ranging from red to yellow, respectively). In the middle we show the histogram of values for $\{\ell_j\}$ accumulated across all 16014 picked-particles. On the right we show a scatterplot of ℓ_j along the horizontal, and $\log(\ell_j)$ along the vertical. The color indicates the amount of added noise, and uses the same scale as the right subplot. The affine-fit is shown in grey in the background. Note that, for this experimental data, the values of $\log(\ell_j)$ can be very well approximated by an affine function $\log(\ell) \approx a\ell + b$ for values of $a \approx 0.537$ and $b \approx -0.378$ that are fixed and independent of the image-index j . Indeed, for this data-set the average relative error between the the values of $\log(\ell_j)$ and the linear-fit is $\approx 6 \times 10^{-4}$.

When this approximation can be made, the negative-log-likelihood can be approximated by:

$$-\log \left(P \left(\{\widehat{A}_j\} | \widehat{F}, \widehat{\sigma}, \# \right) \right) \approx +\text{dof}' \cdot b \cdot J + \text{dof}' \cdot a \cdot \sum_{j=1}^J \ell \left(\widehat{A}_j | \widehat{S}_j^{\text{opt}}, \# \right), \quad (176)$$

Which is once again amenable to a volumetric-representation very similar to Eq. 73.

24 Ill-posedness: Case studies

24.1 Equatorially-distributed viewing-angles

As described in section 22.3, the single-particle reconstruction problem is generically poorly posed when the viewing angles are concentrated around the equator. More specifically, assume we are given a volume \widehat{F} and a collection of templates taken from a viewing-angle-distribution $\mu(\tau)$ supported on the equator. Now any equatorial-perturbation $\Delta\tau(\tau)$ of these viewing-angles can be applied to the original viewing-angles. These perturbed viewing-angles can then be used to reconstruct a volume \widehat{G} .

As long as the map $\tau \leftarrow \tau + \Delta\tau(\tau)$ is one-to-one, the reconstructed \widehat{G} will then produce a collection of templates that perfectly match those templates used to reconstruct it. In other words, the collection of templates of \widehat{G} will include the original collection of templates from \widehat{F} .

This issue is a result of the fact that, when $\mu(\tau)$ is restricted to the equator, different templates only overlap at the poles. Any invertible equatorial-perturbation simply reorganizes the azimuthal-angle associated with each template, without disturbing this polar overlap. Examples of this phenomenon are shown in Figs 27-34.




Figure 27: On the left we show an idealized viewing-angle-distribution $\mu(\tau)$ which is uniformly distributed along the equator. On the right we show a low-res ($2\pi K = 48$) reconstruction F of the TRPV1 molecule obtained by appropriately aligning many synthetic-images (i.e., noisy projections) taken from these equatorial viewing-angles. This low-res reconstruction F agrees with the reconstruction referenced in EMPIAR-10005.




Figure 28: In this figure we show a collection of τ -specific alignment-perturbations $\{\Delta\tau(\tau)\}$ (left) along with a corresponding volumetric-perturbation δF (right). The arrows on the left indicate the direction of the alignment-perturbations for each (β, α) , with $\Delta\gamma(\tau) = 0$. Larger/darker arrows indicate larger values of $\Delta\tau(\tau)$, with the arrow-area proportional to $|\Delta\tau|$ (and dark purple arrows corresponding to the largest values of $|\Delta\tau|$). On the right we show an isosurface of $|\Delta F|$, colored by the sign of ΔF (with blue and pink corresponding to negative and positive, respectively). The projections of ΔF are shown along each axis (background and bottom). These perturbations have been selected so that, when applied to the reconstruction F shown in Fig 27, the likelihood remains the same. Put another way, each equatorial-template of the original F at angle τ corresponds to a template of the perturbed volume $F + \Delta F$ at angle $\tau + \Delta\tau(\tau)$. Indeed, when inspecting the right-hand subplot, one can see that the equatorial projections of ΔF shown along the background are identically 0. Only the polar-projection shown on the bottom is nonzero (however, for this example there are no viewing-angles that have a polar-component).




Figure 29: In this figure we show one of the perturbed volumes $F - \Delta F$ of the TRPV1 molecule (at $2\pi K = 48$). On the left subplot we show an isosurface (center), along with projections along each axis (background and bottom). On the right subplot we show the same isosurface, re-colored to indicating the value of the volumetric-deformation ΔF at each point. The color on the right ranges from blue (negative) to pink (positive).




Figure 30: This figure has the same format as Fig 29, except that we illustrate $F + \Delta F$.




Figure 31: This figure illustrates an even more extreme volumetric perturbation that, again, has the same collection of equatorial-templates as the original F .




Figure 32: This figure illustrates an even more extreme volumetric perturbation that, again, has the same collection of equatorial-templates as the original F .




Figure 33: This figure illustrates an even more extreme volumetric perturbation that, again, has the same collection of equatorial-templates as the original F .




Figure 34: This figure illustrates an even more extreme volumetric perturbation that, again, has the same collection of equatorial-templates as the original F . To summarize, because the viewing-angle distribution is supported along the equator, any one-to-one equatorial perturbation of the viewing-angles gives rise to a reconstructed volume with the same collection of equatorial-projections (i.e., templates). In other words, the single-particle reconstruction problem is ill-posed for this example.

24.2 Uniformly-distributed viewing-angles

In the section above we discussed a somewhat degenerate scenario: one where the viewing-angles τ_j for each observed image were all drawn from the equator. In a more general scenario we would expect the distribution $\mu(\tau)$ of viewing-angles to be more uniformly distributed. In this more general situation we typically don't find the same kinds of degeneracies.

Indeed, in the zero-noise limit it is typically the case that there will be a unique volume that satisfies the single-particle reconstruction problem [52, 53].




Figure 35: In this figure we illustrate the frequency-shells $\widehat{F}_A(k, \hat{\mathbf{k}})$, $\widehat{F}_B(k, \hat{\mathbf{k}})$ and $\widehat{F}_C(k, \hat{\mathbf{k}})$. The real- and imaginary-parts of these frequency-shells are shown from both a three-quarters view and a polar view. Each frequency shell adopts the background values of b and b^\dagger in the northern- and southern-hemispheres, respectively. In the regions surrounding the north- and south-poles, these frequency-shells adopt the values of d and d^\dagger , respectively. Note that any great-circle cutting through $\widehat{F}_A(k, \hat{\mathbf{k}})$ will correspond to at least one great-circle cutting through $\widehat{F}_B(k, \hat{\mathbf{k}})$ and at least one great-circle cutting through $\widehat{F}_C(k, \hat{\mathbf{k}})$.




Figure 36: In this figure we illustrate level-sets of the volumes produced by convolving the frequency-shells shown in Fig 35 with a radial-kernel. Note that, by construction, each projection of the volume F_A will correspond to some projection of F_B and also to some projection of F_C . In other words, given the full collection of templates produced by F_A , the single-particle reconstruction problem is poorly posed.

Nevertheless, even though they are not generic, degeneracies can still exist. One of many such examples is shown in Figs 35 and 36. To explain this example, let's first consider a single frequency-shell $\widehat{F}_A(k, \hat{\mathbf{k}})$ (with fixed frequency-amplitude k). On this frequency-shell let's imagine that there is a baseline value of b in the northern-hemisphere and – due to the conjugacy constraint – the corresponding value of b^\dagger in the southern-hemisphere. Let's also assume there exists a localized feature, in this case a circular disc ‘ D ’ around the north-pole where the value of $\widehat{F}_A(k, \hat{\mathbf{k}})$ is elevated from the baseline value of ‘ b ’ to some different value ‘ d ’. Because of the conjugacy-constraint applied to \widehat{F}_A , the presence of D

mandates an antipodal region D' on which \widehat{F}_A takes the value d^\dagger . On this frequency-shell each projection corresponds to a great-circle. Most such great-circles do not intersect D or D' , and thus correspond to projections \widehat{S} which cut through the background-values of b and b^\dagger (i.e., for which $\widehat{S}(\psi) \equiv b$ or b^\dagger for all ψ). However, great-circles that intersect D and D' will produce projections \widehat{S} that have intervals of ψ for which $\widehat{S}(\psi) = d$ and $\widehat{S}(\psi + \pi) = d^\dagger$.

Now consider another frequency-shell $\widehat{F}_B(k, \hat{\mathbf{k}})$ which is similar to $\widehat{F}_A(k, \hat{\mathbf{k}})$, except that the localized feature now lies on an elliptical zone 'E'. The zone E has a major-axis that matches that of D , but a minor-axis which is smaller. Let's assume that \widehat{F}_B takes on the values b on most of the sphere, and the values d and d^\dagger on the elliptical-regions E and E' (where E' is antipodal to E).

Now each great-circle from $\widehat{F}_A(k, \hat{\mathbf{k}})$ corresponds to at least one great-circle from $\widehat{F}_B(k, \hat{\mathbf{k}})$, and vice-versa. A collection of templates corresponding to a uniform distribution of viewing-angles from \widehat{F}_A can now be associated with a collection of templates corresponding to a non-uniform distribution of viewing-angles from \widehat{F}_B . Of course the same story holds if we reduce the minor-axis of the elliptical-region E to produce a shell $\widehat{F}_C(k, \hat{\mathbf{k}})$. Thus, in this case the single-particle reconstruction problem is poorly posed. Give only the collection of templates from F_A , but without knowledge of the viewing-angles nor of the viewing-angle-distribution $\mu(\tau)$, one cannot determine if the true underlying volume is F_A , F_B or any other admissible function (e.g., F_C).

The scenario described above is just one example of the kinds of degeneracies that can occur. Figs 37 and 38 illustrate another similar degenerate scenario. In this case each projection of the volume F_A corresponds to at least one projection of F_B or of F_C . Thus, given only an image-pool comprising the collection of projections from F_A , one cannot determine if the image-pool was produced by the single volume F_A , or a combination of other volumes (e.g., F_B and F_C).




Figure 37: This figure is similar to Fig 35. This time the frequency-shells adopt more than one value in the polar-regions. Note that any great-circle cutting through $\widehat{F}_A(k, \hat{\mathbf{k}})$ will correspond to at least one great-circle cutting through either $\widehat{F}_B(k, \hat{\mathbf{k}})$ and/or $\widehat{F}_C(k, \hat{\mathbf{k}})$.




Figure 38: In this figure we illustrate level-sets of the volumes produced by convolving the frequency-shells shown in Fig 37 with a radial-kernel. Note that, by construction, each projection of the volume F_A will correspond to some projection of F_B and/or to some projection of F_C . In other words, given the full collection of templates produced by F_A , the multi-particle reconstruction problem is poorly posed.

24.3 Noisy images

The previous section discussed degenerate scenarios exhibiting spurious heterogeneity: a collection of projections from one volume can be indistinguishable from a subset of projections collected from a pair of other volumes.

A similar scenario can manifest in real data, particularly when the number of images J is relatively small and the viewing-angle distribution μ_τ^{emp} of the data is non-uniform. As an example, consider the trpv1-molecule and empirical viewing-angle distribution shown in Fig 8.

Now let's consider the multi-particle reconstruction problem, given only the first $J = 1024$ images from the image-stack at EMPIAR_10005. For simplicity, let's ignore translations for now (assuming that each image is appropriately centered). As described in section 13.2, one might approximately solve this problem by constructing a collection of volumes $\{F_i\}$ and, for each image A_j , determining the maximum-likelihood alignments $\{\tau_j^{\text{opt}}\}$ and associated volume-label i_j^{opt} .

There is a trivial (and useless) solution to this problem involving a number of volumes equal to the number of images, where volume j is fit to image A_j . Clearly, this trivial solution is overfit to the data, and one might hope to prevent such overfitting by restricting the total number of volumes. However, as we'll see in a moment, such a strategy is not guaranteed to succeed, and in this scenario we still observe spurious heterogeneity even if the number of volumes is restricted to only 2.

Spurious heterogeneity with synthetic images: To see why this might be the case, recall that the empirical viewing-angle distribution from EMPIAR_10005 is partially concentrated around the equator, with the remainder concentrated near the poles. Let's now consider performing 'surgery' on μ_τ^{emp} , and dividing this into two disjoint distributions, one μ_τ^{equa} supported in a band around the equator, and the second μ_τ^{pole} supported on a pair of antipodal polar caps.

As discussed in section 22.3, the equatorial viewing-angles in μ_τ^{equa} can, by themselves, give rise to a single-particle reconstruction which is not very robust. That is to say, there are volumes that are not close to the true volume, but which can produce a collection of equatorial-templates which are a very close match to the equatorial-templates of the true volume. Generally speaking, the quality of this match, as measured by the likelihood in Eq. 54 will increase as the viewing-angle distribution μ_τ^{equa} becomes concentrated around the equator (see Fig 25).

Similarly, as discussed in section 22.2, the polar viewing-angles in μ_τ^{pole} can, by themselves, also produce a reconstruction which is not very robust. Once again, there are volumes that are not close to the true volume, but which can produce nearly identical polar-templates. Generally speaking, the quality of this match will increase as the viewing-angle distribution μ_τ^{pole} becomes concentrated around the poles (see Fig 24).

For this particular case-study we'll simply chop μ_τ^{emp} into two parts, as indicated in Fig 39. Note that, while the data from μ_τ^{equa} technically covers the surface of the sphere, the data from μ_τ^{pole} does not. As show on the right of Fig 39, there is a 'blind spot' at each pole where the data from μ_τ^{pole} does not constrain the values of \hat{F} near the poles.

We'll return to these blind-spots in a moment, but for now we'll illustrate in Figs 40-43 two volumes F_B and F_C . These two volumes correspond to reconstructions using the distributions μ_τ^{pole} and μ_τ^{equa} , respectively. These volumes are illustrated alongside the original volume F_A , which corresponds to the true volume EMD-5778.

The volume F_B has been crafted so that, for each template $S(\cdot; \tau^A, F_A)$ with a τ^A drawn from μ_τ^{pole} , there exists a template $S(\cdot; \tau^B, F_B)$ which is very close to $S(\cdot; \tau^A, F_A)$. Similarly, the volume F_C has been crafted so that, for each template $S(\cdot; \tau^A, F_A)$ with a τ^A drawn from μ_τ^{equa} , there exists a template $S(\cdot; \tau^C, F_C)$ which is very close to $S(\cdot; \tau^A, F_A)$.

To be more explicit, consider taking a sample of J synthetic-images $\{\widehat{A}_j^{\text{syn}}\}$, each produced using the true volume \widehat{F}_A and a viewing-angle τ_j^A drawn from μ_τ^{emp} . If we were to align $\widehat{A}_j^{\text{syn}}$ to \widehat{F}_A , then the optimal viewing-angle would be $\tau_j^{\text{opt},A} \equiv \tau_j^A$, by construction. This collection of synthetic-images forms our ‘synthetic-image-pool’, and is the data we will use immediately below.

Let’s first align each $\widehat{A}_j^{\text{syn}}$ to \widehat{F}_B , producing an optimal viewing-angle $\tau_j^{\text{opt},B}$. This optimal viewing-angle will only correspond to a close match if $\widehat{A}_j^{\text{syn}}$ corresponds to a polar-template (i.e., if τ_j^A lies in the support of μ_τ^{pole}). To indicate this, let’s define $\tau_j^{\text{opt},B+C} := \tau_j^{\text{opt},B}$ and set the volume-assignment $i_j^{\text{opt},B+C}$ to be ‘B’ for these images.

In a similar fashion, we can align each $\widehat{A}_j^{\text{syn}}$ to \widehat{F}_C , producing an optimal viewing-angle $\tau_j^{\text{opt},C}$. This optimal viewing-angle will only correspond to a close match if $\widehat{A}_j^{\text{syn}}$ corresponds to an equatorial-template (i.e., if τ_j^A lies in the support of μ_τ^{equa}). To indicate this, let’s define $\tau_j^{\text{opt},B+C} := \tau_j^{\text{opt},C}$ and set the volume-assignment $i_j^{\text{opt},B+C}$ to be ‘C’ for these images.

Now we can use the optimal-alignments $\{\tau_j^{\text{opt},A}\}$ to calculate the log-likelihood $\mathcal{L}_A := \log\left(P\left(\{\widehat{A}_j^{\text{syn}}\} \mid \widehat{F}_A\right)\right)$, as described in Eq. 54. Similarly, we can use the optimal-alignments $\{\tau_j^{\text{opt},B+C}\}$ and optimal-assignments $\{i_j^{\text{opt},B+C}\}$ to calculate the log-likelihood $\mathcal{L}_{B+C} := \log\left(P\left(\{\widehat{A}_j\} \mid \{\widehat{F}_B, \widehat{F}_C\}\right)\right)$ as described in Eq. 56.

Now we can compare two different hypothesis. The first hypothesis, denoted by ‘Hyp_A’, is that the synthetic-image-pool was generated by the single volume x_{F_A} and the alignments $\{\tau_j^{\text{opt},A}\}$. The second hypothesis, denoted by ‘Hyp_{B+C}’ is that the synthetic-image-pool was generated by the pair of volumes $\{F_B, F_C\}$ along with the alignments $\{\tau_j^{\text{opt},B+C}\}$ and the volume-assignments $\{i_j^{\text{opt},B+C}\}$. Because we are using synthetic-images, we are guaranteed that the ‘true’ hypothesis Hyp_A is more likely than the ‘spurious’ hypothesis Hyp_{B+C} (see section 18). However, for this case-study we have prepared the volumes F_B and F_C so that, given the empirical noise-level (estimated from EMPIAR.10005), the log-likelihood ratio $\mathcal{L}_A - \mathcal{L}_{B+C}$ is actually slightly less than $\log(20)$. This means that, under standard-practice, the spurious hypothesis Hyp_{B+C} would not be rejected in favor of the true hypothesis Hyp_A.

Spurious heterogeneity with experimental-images: Above we considered a somewhat ideal setting where the image-pool was constructed using only synthetic images drawn from the true volume F_A . If the size J of the image-pool is not too large and the empirical viewing-angle distribution μ_τ^{emp} is nonuniform, then we can encounter spurious heterogeneity. Notably, if we revisit the same setting and use actual experimental-images, we see that spurious heterogeneity can be even more pernicious.

Recall that the distribution μ_τ^{pole} was quite concentrated near the poles, yielding a ‘blind-spot’ near each pole shown in Fig 39. That is to say, the data from the synthetic-images drawn from μ_τ^{pole} do not constrain the values of \widehat{F}_B in these blind-spots. Thus, these blind-spots represents a collection of degrees-of-freedom that are unconstrained by the true projections of the volume.

Above, to produce the polar-aligned volume F_B shown in Figs 40-43, we merely set the values of \widehat{F}_B to be constant in each blind-spot (subject to the conjugacy conditions). However, in practice these blind-spots are not likely to be treated so simply.

One such example is shown on the right of Fig 44. This shows a volume $F_{B'}$ which has the same polar-templates as F_B from Fig 41, but has different data in its blind spot, corresponding to ‘overfitting’ to noise.

To explain this scenario more carefully, first recall the original pool of $J = 1024$ synthetic-images described above. As described above, these J synthetic-images were constructed using projections of the true molecule F_A . When aligning these synthetic-images to the spurious-volumes F_B and F_C , those synthetic-images corresponding to polar-templates fit best to F_B , while those corresponding to equatorial-templates fit best to F_C .

Now let’s add 32 additional noisy experimental-images to the image-pool. These additional noisy images have viewing-angles that are not too far from the poles, and naturally align best to F_B . However, due to the noise, their optimal alignment-angle $\tau_j^{\text{opt},B}$ is somewhat far from the poles, so that the images themselves correspond to values of \hat{k} that pass through the polar ‘blind-spots’ associated with μ_τ^{pole} .

Now, by using these optimally-aligned noisy images (along with the optimally-aligned synthetic-images corresponding to polar-templates), we can reconstruct $F_{B'}$. Unlike the volume F_B shown in the middle of Fig 44, the volume $F_{B'}$ shown on the right of Fig 44 has its polar-data fit to the additional noisy experimental-images. This additional polar-information distorts the volume $F_{B'}$, but does not appreciably change the polar-projections. Thus while the volumes F_B and $F_{B'}$ provides equally good matches to the polar-templates, the volume $F_{B'}$ provides a better match than the ‘cleaner’ volume F_B to the noisy experimental-images.

This results in a classical instance of ‘overfitting’: When considering the expanded image-pool comprising both the idealized synthetic-images and the additional 32 noisy experimental-images, the hypothesis Hyp_{B'+C} is actually *more* likely than Hyp_A. A similar scenario is shown in Fig 45, using a different collection of 64 noisy experimental-images.

Finally, we remark that spurious heterogeneity of this kind can be difficult to correct for a-posteriori. In many cases

it might not be clear which images are ‘good’ images of particles, and which are too strongly polluted by noise. It also might be unclear which features of the reconstructed volume (or volumes) are a result of overfitting to noise, or if such a distinction can even be made. Moreover, some commonly used techniques for validation, such as aligning the images to an ‘averaged’ volume of the form $0.5 \times (F_B + F_C)$, could erroneously provide support for the spurious hypothesis.

In these situations we would recommend calculating the softest modes of the Hessian associated with each reconstructed volume, as described in the main text. While such a calculation won’t immediately identify a heterogeneous reconstruction as spurious, it can certainly provide a warning. For example, in the case-studies shown above the spurious volume F_B is susceptible to latitudinal perturbations, while the spurious volume F_C is susceptible to longitudinal perturbations. Both of these spurious volumes are less robust to perturbations than the volume F_A produced via single-particle reconstruction. This kind of analysis can help indicate possible errors associated with a heterogeneous reconstruction.




Figure 39: In this figure we illustrate the viewing-angles corresponding to μ_τ^{equa} (cyan dots, left) and μ_τ^{pole} (cyan dots, middle and right). Each viewing-angle is indicated by a cyan dot hovering above the surface of the sphere. The fourier-slice corresponding to each viewing-angle is supported on a great-circle orthogonal to the viewing-angle itself. This data (accumulated across all viewing-angles) is illustrated by pink dots covering the surface of the sphere. Note that, while the data from μ_τ^{equa} technically covers the surface of the sphere, the data from μ_τ^{pole} does not. As show on the right, there is a ‘blind spot’ near each pole where the data does not constrain the values of \hat{F} near the poles.




Figure 40: In this figure we illustrate three level-sets of volumes F_A , F_B and F_C . The volume F_A on the left is the reconstruction EMD-5778 of trpv1 referenced in EMPIAR-10005. The empirical-distribution of viewing-angles μ_τ^{emp} for this volume in the data-set EMPIAR-10005 is concentrated around the equator and the poles (see Fig 8). Templates of F_A with viewing-angles $\tau \sim \mu_\tau^{\text{pole}}$ are very likely to have a close match to one of the templates of F_B . Similarly, templates of F_A with viewing-angles $\tau \sim \mu_\tau^{\text{equa}}$ are very likely to have a close match to one of the templates of F_C . Thus, by construction, each projection of the volume F_A will correspond to some projection of F_B and/or to some projection of F_C . In other words, given a collection of synthetic-images produced by F_A , the multi-particle reconstruction problem can be poorly posed.




Figure 41: This is similar to Fig 40, except for a different level-surface.




Figure 42: This is similar to Fig 40, except for a different level-surface.




Figure 43: This is similar to Fig 40, except for a different level-surface.




Figure 44: This is similar to Fig 41. We show the original F_A on the left, and the polar-matched volume F_B in the middle. On the far right we now show an alternative volume $F_{B'}$. This alternative volume again matches the polar-templates from F_A , but has its ‘blind-spots’ fit to a small subset of 32 noisy images.




Figure 45: This is similar to Fig 44. This time the alternative volume on the right is fit to a different subset of 64 noisy images.

References

- [1] A. B. Goncharov. Methods of integral geometry and finding the relative orientation of identical particles arbitrarily arranged in a plane from their projections onto a straight line. *Dokl. Phys.*, 32:173, 1987.

- [2] M. van Heel. Angular reconstitution: A posteriori assignment of projection directions for 3D reconstruction. *Ultramicroscopy*, 21:111–123, 1987.
- [3] A. B. Goncharov and M. S. Gelfand. Determination of mutual orientation of identical particles from their projections by the moments method. *Ultramicroscopy*, 25:317–328, 1988.
- [4] S. Jonic, C.O. Sorzano, P. Thevenaz, C. El-Bez, S. De Carlo, and M. Unser. Spline-based image-to-volume registration for three-dimensional electron microscopy. *Ultramicroscopy*, 103(4):303–317, 2005.
- [5] Yifan Cheng, Nikolaus Grigorieff, Pawel A. Penczek, and Thomas Walz. A primer to single-particle cryo-electron microscopy. *Cell*, 161:439–449, 2015.
- [6] E. Nogales and S.H. Scheres. Cryo-EM: a unique tool for the visualization of macromolecular complexity. *Mol. Cell.*, 58:677–689, 2015.
- [7] Dominika Elmlund and Hans Elmlund. Cryogenic electron microscopy and single-particle analysis. *Annu. Rev. Biochem.*, 84:499–517, 2015.
- [8] K. Murata and M. Wolf. Cryo-electron microscopy for structural analysis of dynamic biological macromolecules. *Biochim. Biophys. Acta Gen. Subj.*, 1862(2):324–334, 2017.
- [9] F Sigworth. Principles of cryo-EM single-particle image processing. *Microscopy*, 65(1):57–67, 2016.
- [10] D. Bhella. Cryo-electron microscopy: an introduction to the technique, and considerations when working to establish a national facility. *Biophys Rev.*, 11:515–519, 2019.
- [11] Nikolaus Grigorieff. FREALIGN: High-resolution refinement of single particle structures. *J. Struct. Biol.*, 157(1):117–125, 2007.
- [12] G. Tang, L. Peng, P.R. Baldwin, D.S. Mann, W. Jiang, I. Rees, and S.J. Ludtke. EMAN2: An extensible image processing suite for electron microscopy. *J. Struct. Biol.*, 157:38–46, 2007.
- [13] Zh. Yang and P.A. Penczek. Cryo-EM image alignment based on nonuniform fast Fourier transform. *Ultramicroscopy*, 108:959–969, 2008.
- [14] A. Singer, R. R. Coifman, F. J. Sigworth, D. W. Chester, and Y. Shkolnisky. Detecting consistent common lines in cryo-EM by voting. *J. Struct. Biol.*, 169:312–322, 2009.
- [15] A. Singer and Y. Shkolnisky. Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming. *SIAM J. Imaging Sci.*, 4(2):543–572, 2011.
- [16] Sjors H W Scheres. A Bayesian view on cryo-EM structure determination. *J. Mol. Biol.*, 415:406–418, 2012.
- [17] Y. Shkolnisky and A. Singer. Viewing direction estimation in cryo-EM using synchronization. *SIAM J. Imaging Sci.*, 5(3):1088–1110, 2012.
- [18] Dmitry Lyumkis, Axel F. Brilot, Douglas L. Theobald, and Nikolaus Grigorieff. Likelihood-based classification of cryo-EM images using FREALIGN. *J. Struct. Biol.*, 183(3):377–388, 2013.
- [19] L. Wang, A. Singer, and Z. Wen. Orientation determination from cryo-EM images using least unsquared deviations. *SIAM J. Imaging Sci.*, 6(4):2450–83, 2013.
- [20] N. Grigorieff. Frealign: An exploratory tool for single-particle cryo-EM. In R A Crowther, editor, *The Resolution Revolution: Recent Advances In cryoEM*, volume 579 of *Methods Enzymol.*, pages 191–226. Academic Press, 2016.
- [21] A. Punjani, J.L. Rubinstein, D.J. Fleet, and M.A. Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nat. Methods*, 14:290–296, 2017.
- [22] A. Punjani, M.A. Brubaker, and D.J. Fleet. Building proteins in a day: efficient 3D molecular structure estimation with electron cryomicroscopy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):706–718, 2017.
- [23] L.A. Passmore and C.J. Russo. Specimen preparation for high-resolution cryo-em. *Methods Enzymol.*, 579:51–86, 2016.
- [24] R.F. Thompson, M. Walker, C.A. Siebert, S.P. Muench, and N.A. Ranson. An introduction to sample preparation and imaging by cryo-electron microscopy for structural biology. *Methods*, 100:3–15, 2016.

- [25] Ayelet Heimowitz, Joakim Andén, and Amit Singer. Apple picker: Automatic particle picking, a low-effort cryo-em framework. Journal of Structural Biology, 204(2):215–227, 2018.
- [26] Donal M. McSweeney, Sean M. McSweeney, and Qun Liu. A self-supervised workflow for particle picking in cryo-em. IUCrJ, 7:719–727, 2020.
- [27] Robert Langlois, Jesper Pallesen, Jordan T. Ash, Danny Nam Ho, John L. Rubinstein, and Joachim Frank. Automated particle picking for low-contrast macromolecules in cryo-electron microscopy. Journal of Structural Biology, 186(1):1–7, 2014.
- [28] Vibhor Kumar, Jukka Heikkonen, Peter Engelhardt, and Kimmo Kaski. Robust filtering and particle picking in micrograph images towards 3d reconstruction of purified proteins with cryo-electron microscopy. Journal of Structural Biology, 145(1):41–51, 2004. Automated Particle Selection for Cryo-Electron Microscopy.
- [29] W. Park, D. R. Madden, D. N. Rockmore, and G. S. Chirikjian. Deblurring of class-averaged images in single-particle electron microscopy. Inverse Prob., 26(035002), 2010.
- [30] Joachim Frank. Chapter 3 - two-dimensional averaging techniques. In Joachim Frank, editor, Three-Dimensional Electron Microscopy of Macromolecular Assemblies, pages 54–125. Academic Press, Burlington, 1996.
- [31] F J Sigworth. A maximum-likelihood approach to single-particle image refinement. J. Struct. Biol., 122(3):328–39, 1998.
- [32] S H W Scheres, M Valle, P Grob, E Nogales, and J.-M. Carazo. Maximum likelihood refinement of electron microscopy data with normalization errors. J. Struct. Biol., 166(2):234–240, 2009.
- [33] Sjors H W Scheres. RELION: Implementation of a Bayesian approach to cryo-EM structure determination. J. Struct. Biol., 180(3):519–530, 2012.
- [34] James M. Bell, Muyuan Chen, Philip R. Baldwin, and Steven J. Ludtke. High resolution single particle refinement in EMAN2.1. Methods, 100:25–34, 2016.
- [35] F. J. Sigworth, Doerschuk P.C., J.-M. Carazo, and S.H.W. Scheres. An introduction to maximum-likelihood methods in Cryo-EM. In Methods in Enzymology. Cryo-EM, Part B: 3D reconstruction, pages 263–294. Academic Press., 2010.
- [36] Jing Leng, Weiguo Gao, Cheng Shang, and Zhi-Pan Liu. Efficient softest mode finding in transition states calculations. The Journal of Chemical Physics, 138(9):094110, 03 2013.
- [37] Richard Henderson, Andrej Sali, Matthew L. Baker, Bridget Carragher, Batsal Devkota, Kenneth H. Downing, Edward H. Egelman, Zukang Feng, Joachim Frank, Nikolaus Grigorieff, Wen Jiang, Steven J. Ludtke, Ohad Medalia, Pawel A. Penczek, Peter B. Rosenthal, Michael G. Rossmann, Michael F. Schmid, Gunnar F. Schröder, Alasdair C. Steven, David L. Stokes, John D. Westbrook, Willy Wriggers, Huanwang Yang, Jasmine Young, Helen M. Berman, Wah Chiu, Gerard J. Kleywegt, and Catherine L. Lawson. Outcome of the first electron microscopy validation task force meeting. Structure, 20(2):205–214, 2012.
- [38] Pawel A. Penczek. Ab initio cryo-EM structure determination as a validation problem. 2014.
- [39] J Bernard Heymann. Validation of 3D EM Reconstructions: The Phantom in the Noise. AIMS Biophysics, 2:21–35, 2015.
- [40] Peter B. Rosenthal and John L. Rubinstein. Validating maps from single particle electron cryomicroscopy. Current Opinion in Structural Biology, 34:135–144, 2015.
- [41] M. Liao, E. Cao, D. Julius, and Y. Cheng. Structure of the TRPV1 ion channel determined by electron cryo-microscopy. Nature, 504:107–12, 2013.
- [42] Aaditya V. Rangan and Leslie Greengard. Robust ab initio solution of the cryo-em reconstruction problem at low resolution with small data sets. Journal of Structural Biology, 215(3):107994, 2023.
- [43] Julian Giraldo-Barreto, Sebastián Ortiz, Erik Thiede, Karen Palacio-Rodríguez, Bob Carpenter, Alex Barnett, and Pilar Cossio. A bayesian approach to extracting free-energy profiles from cryo-electron microscopy experiments. Scientific Reports, 11, 07 2021.
- [44] Wai Shing Tang, Jeff Soules, Aaditya Rangan, and Pilar Cossio. Cryolike: A python package for cryo-electron microscopy image-to-structure likelihood calculations. bioRxiv, 2024.

- [45] Bronwyn A. Lucas and Nikolaus Grigorieff. Quantification of gallium cryo-fib milling damage in biological lamellae. Proceedings of the National Academy of Sciences, 120(23):e2301852120, 2023.
- [46] Kexin Zhang, Pilar Cossio, Aaditya Rangan, Bronwyn Lucas, and Nikolaus Grigorieff. New statistical metric for robust target detection in cryo-em using 2dtm. bioRxiv, 2024.
- [47] Aaditya Rangan, Marina Spivak, Joakim Andén, and Alex Barnett. Factorization of the translation kernel for fast rigid image alignment. Inverse Problems, 36(2), 2020.
- [48] Aaditya V Rangan. Radial-recombination for rigid rotational alignment of images and volumes. Inverse Problems (submitted), 2022.
- [49] R Bracewell. The Fourier Transform and Its Applications. McGraw-Hill, 3rd edition, 1999.
- [50] A. Barnett, L. Greengard, A. Pataki, and M. Spivak. Rapid solution of the cryo-EM reconstruction problem by frequency marching. SIAM J. Imaging Sci., 10(3):1170–1195, 2017.
- [51] W. Baxter, R. Grassucci, Haixiao Gao, and J. Frank. Determination of signal-to-noise ratios and spectral snrs in cryo-em low-dose imaging of molecules. Journal of structural biology, 166 2:126–32, 2009.
- [52] Z. Zhao and A. Singer. Rotationally invariant image representation for viewing direction classification in cryo-em. J Struct Biol., 186(1):153–66, 2014.
- [53] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer. Bispectrum inversion with application to multireference alignment. IEEE Trans Signal Process., 66(4):1037–1050, 2018.