Structure from Noise: Confirmation Bias in Particle Picking in Structural Biology

Amnon Balanov*[1], Alon Zabatani[1], and Tamir Bendory[1]

[1]School of Electrical and Computer Engineering, Tel Aviv University, Tel Aviv 69978, Israel

July 5, 2025

Abstract

Confirmation bias is a fundamental challenge in cryo-electron microscopy (cryo-EM) and cryo-electron tomography (cryo-ET), where prior expectations can lead to systematic errors in data interpretation. This bias may emerge at multiple stages of the reconstruction pipeline, and in particular in the critical particle-picking stage, where 2D particles (in cryo-EM) or 3D subtomograms (in cryo-ET) are extracted from highly noisy micrographs or tomograms. Focusing on two widely used methodologies, template matching and deep neural networks, we combine theoretical analysis with controlled experiments to demonstrate that both methods, when applied to pure noise, can produce persistent molecular structures, a phenomenon we term *structure from noise*. This artifact highlights a critical vulnerability in current workflows: the potential for particle-picking algorithms to inject strong prior-driven bias into downstream analyses. We then propose practical mitigation strategies to reduce the impact of such biases. Together, our findings deepen the theoretical understanding of confirmation bias in cryo-EM and cryo-ET and call for cautious interpretation of reconstructions, primarily when relying on template-driven particle picking.

1 Introduction

Cryo-electron microscopy (cryo-EM) and cryo-electron tomography (cryo-ET) have revolutionized structural biology by enabling visualization of macromolecular complexes in near-native states and at increasingly high resolutions. These complementary techniques are now at the heart of the field. Cryo-EM (Figure 1(a)) reconstructs the 3D structures of isolated macromolecules from 2D projection images acquired at unknown orientations, and has become essential for resolving atomic-level structures of proteins that are difficult to crystallize [32, 1, 37, 39, 58]. In contrast, cryo-ET (Figure 2(a)) captures 3D volumes of macromolecular assemblies in situ by collecting a tilt series of 2D projections at known angles [13, 41, 59, 51, 56]. While cryo-ET typically yields lower resolution, it preserves the

*Corresponding author: amnonba15@gmail.com

native spatial context and is uniquely suited for probing the complexity of cellular environments.

As these techniques mature, a critical concern has emerged: the presence of biases, both experimental and algorithmic, that can affect the accuracy of the final 3D reconstructions. This concern is well recognized in the structural biology community, and practical strategies have been developed to mitigate known sources of bias [49]. However, a deeper theoretical understanding of how such biases arise and propagate through the downstream reconstruction pipeline remains lacking.

Among the algorithmic sources of bias in single-particle reconstruction are incorrect alignment parameters, inaccurate contrast transfer function (CTF) correction, improper image normalization, erroneous masking in real or Fourier space, and the use of incorrect initial templates [49]. This work focuses on the latter and aims to develop a theoretical framework for quantifying its impact. In particular, we seek to establish both theoretical and empirical foundations for analyzing *confirmation bias during the particle-picking stage* in cryo-EM and cryo-ET.

Originally studied in the social sciences, confirmation bias refers to the tendency to interpret information in ways that support pre-existing beliefs or expectations, rather than allowing the data to speak for itself [26, 24, 34]. In the context of structural biology, confirmation bias, viewed as a specific instance of the broader category of model bias, can manifest when researchers, often unconsciously, assume that an unknown structure resembles a previously determined model or a prediction generated by computational tools. These assumptions, introduced early in the pipeline, can bias downstream processing and produce reconstructions that resemble prior models, even when unsupported by the data [44, 49, 2]. This raises a fundamental question that guides this work:

> *To what extent does confirmation bias in particle-picking algorithms shape and potentially distort the downstream pipeline, and in particular the final 3D structure in cryo-EM and cryo-ET?*

The Einstein from Noise phenomenon. To illustrate how methodological assumptions can generate structure from noise, we revisit a striking example widely discussed in the cryo-EM community: the Einstein from Noise phenomenon. In this thought experiment, researchers analyze a dataset believed to contain noisy, shifted instances of a known image, typically an image of Einstein, but in reality, the data consists entirely of pure noise. Despite the absence of any underlying signal, aligning and averaging the images relative to the template yields a reconstruction that closely resembles the original image [45, 46].

This phenomenon played a key role in a significant scientific controversy over the structure of an HIV molecule [30, 22, 53, 50, 29], and has since served as a cautionary tale for methodological overfitting. For a detailed statistical treatment, see [3, 4]. While often cited in discussions of particle picking using template-matching, it is crucial to note that the Einstein from Noise effect arises from a distinct mechanism: in template-matching particle picking methods, confirmation bias arises from selecting micrograph patches that most closely match the templates (essentially a form of *selection bias*). In contrast, the Einstein from Noise effect emerges from aligning already extracted particles to a reference template. We further elaborate on this distinction in Section 4.

The risk of confirmation bias. The risk of confirmation bias is especially pronounced in scenarios with low signal-to-noise ratios (SNRs), a defining characteristic of cryo-EM and cryo-ET datasets. These datasets also exhibit significant uncertainty in particle orientations, rendering reconstructions highly sensitive to initialization and heuristic choices [47, 8]. Figures 1 and 2 illustrate the cryo-EM and cryo-ET setups and their respective particle reconstruction workflows, highlighting critical stages where confirmation bias may emerge. The computational workflows for both modalities involve several common stages: particle detection and extraction (i.e., particle picking), alignment and classification, and 3D reconstruction. In cryo-EM, individual 2D particles are picked from micrographs and grouped into 2D classes (Figure 1(a)); in cryo-ET, smaller subvolumes, termed subtomograms, are extracted from a reconstructed tomogram and then aligned and averaged to enhance the SNR (Figure 2(a)) [59, 56]. In both pipelines, downstream reconstruction accuracy depends critically on the fidelity of the initial *particle-picking stage—the main focus of this work.*

Empirical examples: Confirmation bias in particle picking. This study presents a theoretical and empirical framework to investigate confirmation bias during the crucial *particle-picking* stage of cryo-EM and cryo-ET. Particle picking is the process of extracting 2D particle projections from noisy micrographs (in cryo-EM, Figure 1(b–c)) or 3D subtomograms from noisy tomograms (in cryo-ET, Figure 2(b)). While there are many different particle-picking algorithms, most fall into two main categories: template-matching algorithms (e.g., [14, 23, 31, 12]), and deep learning algorithms (e.g., [11, 54]). Template-matching methods work by scanning the micrograph or tomogram with a predefined reference structure (or template) and selecting regions that correlate highly with it. This technique has been a foundational tool in cryo-EM and cryo-ET since its inception, and it remains a widely used approach today [36]. In contrast, deep learning methods train neural networks, often convolutional neural networks (CNNs), to detect particles based on features learned directly from annotated training data. While both approaches perform well under moderate-to-high SNR, we demonstrate that in low-SNR conditions, common to both cryo-EM and cryo-ET, they can introduce systematic structural biases.

Figure 1(b–c) demonstrates confirmation bias in cryo-EM particle picking using a template-matching algorithm applied to micrographs composed entirely of pure noise. In these methods, predefined particle templates are correlated with input micrographs to identify potential particle locations. After template matching with different initial templates, the selected particles have been processed using RELION's VDAM 2D classification using a reference-free initialization [25]. Remarkably, despite the absence of any informative signal in the data, the resulting 2D class averages closely resemble projections of the original 3D templates. This behavior is consistent across multiple templates, algorithms, and experimental configurations, and also appears in the cryo-ET pipeline, as shown in Figure 2(b). We refer to this phenomenon, where downstream reconstructions produce biologically plausible structures from pure noise, as *structure from noise.*

Main results and work structure. To systematically investigate confirmation bias during the particle-picking stage, we examine the bias in the two leading approaches: classical template matching and deep learning-based methods. For template matching, we develop a

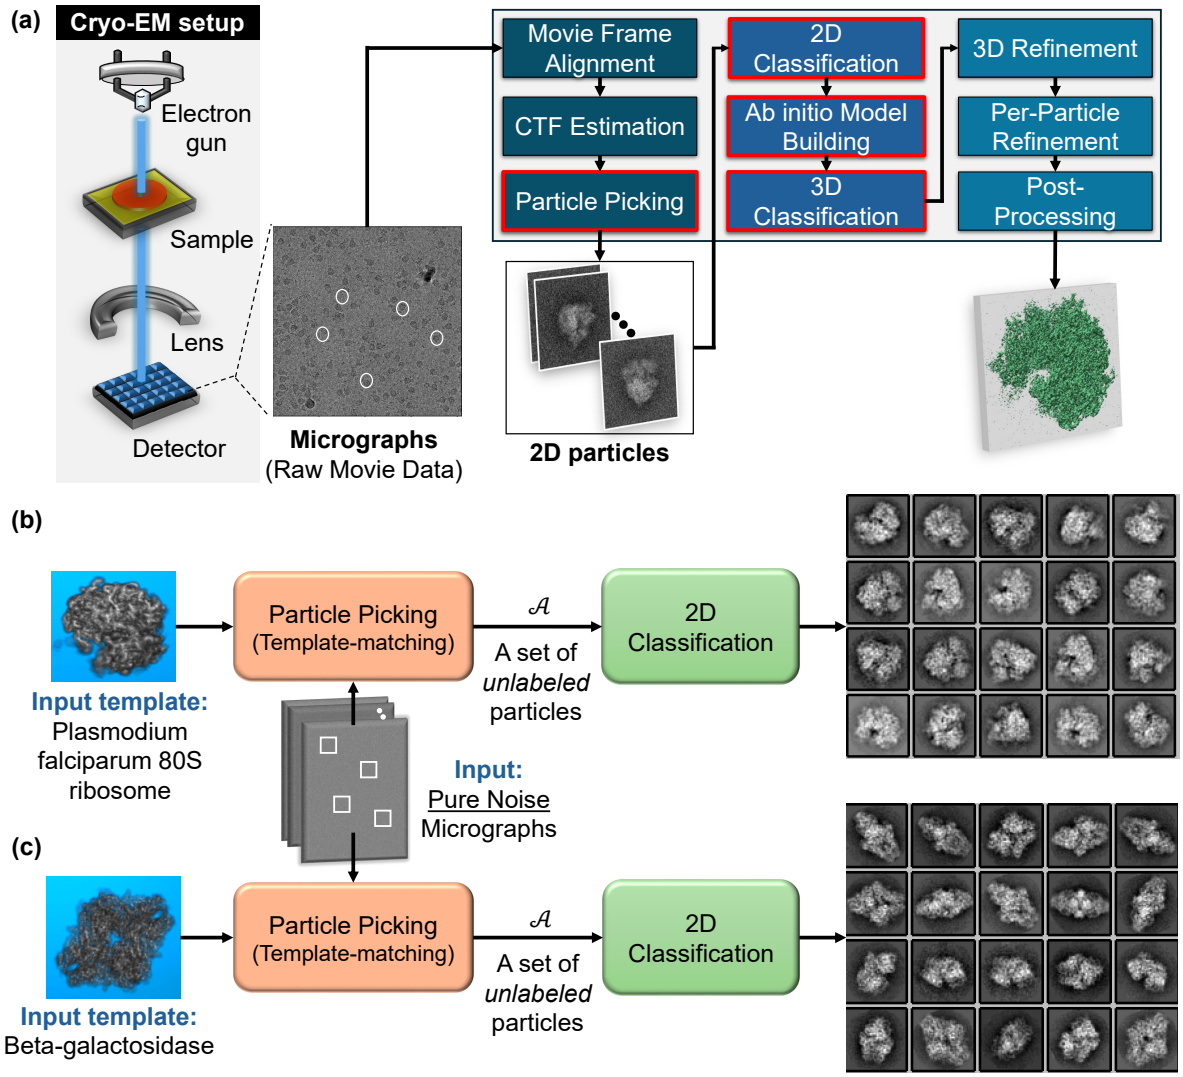

Figure 1: The cryo-electron microscope (cryo-EM) computational pipeline and the structure from noise phenomenon in particle-picking algorithms. (a) A typical cryo-EM workflow, where single-particle electron microscopy is used to reconstruct 3D structures from 2D projections [8]. The process begins with sample preparation, followed by imaging of randomly oriented particles embedded in vitrified ice. The computational pipeline includes preprocessing steps, such as motion correction and contrast transfer function (CTF) estimation and correction, followed by the particle-picking stage to extract individual particles. These particles are subsequently subjected to 2D classification to group similar images, followed by 3D ab initio model reconstruction and refinement to generate a 3D reconstruction. Post-processing steps are applied to refine the final 3D structure. Red squares indicate stages in the cryo-EM pipeline vulnerable to confirmation bias. **(b)+(c) Confirmation bias in particle picking based on template matching.** Our experiment, outlined in Algorithm 1, is applied to micrographs consisting entirely of pure noise. Here, we aim to detect particles using templates that are projections of a 3D structure. In particular, we examine two 3D structures: (b) Plasmodium falciparum 80S ribosome [57], and (c) Beta-galactosidase [7]. It is important to note that template matching, although ubiquitous, is just one approach to particle picking, and bias also exists with other methods, as we demonstrate later. The extracted particles are then classified using a 2D classification algorithm. Remarkably, the resulting 2D class centers closely resemble the templates used during particle picking, even though the input micrographs contain only noise, demonstrating the *structure from noise* bias. The experiments were conducted with $L = 20$ distinct projections of the 3D templates; the 2D classification used the same number of classes. The projection size was $d = 48 \times 48$, and $M = 2 \times 10^5$ particles were extracted from 1000 micrographs. The template-matching threshold is $T = 3.8$. The 2D classification utilized the RELION VDAM algorithm [25], with its parameters detailed in Section 3.

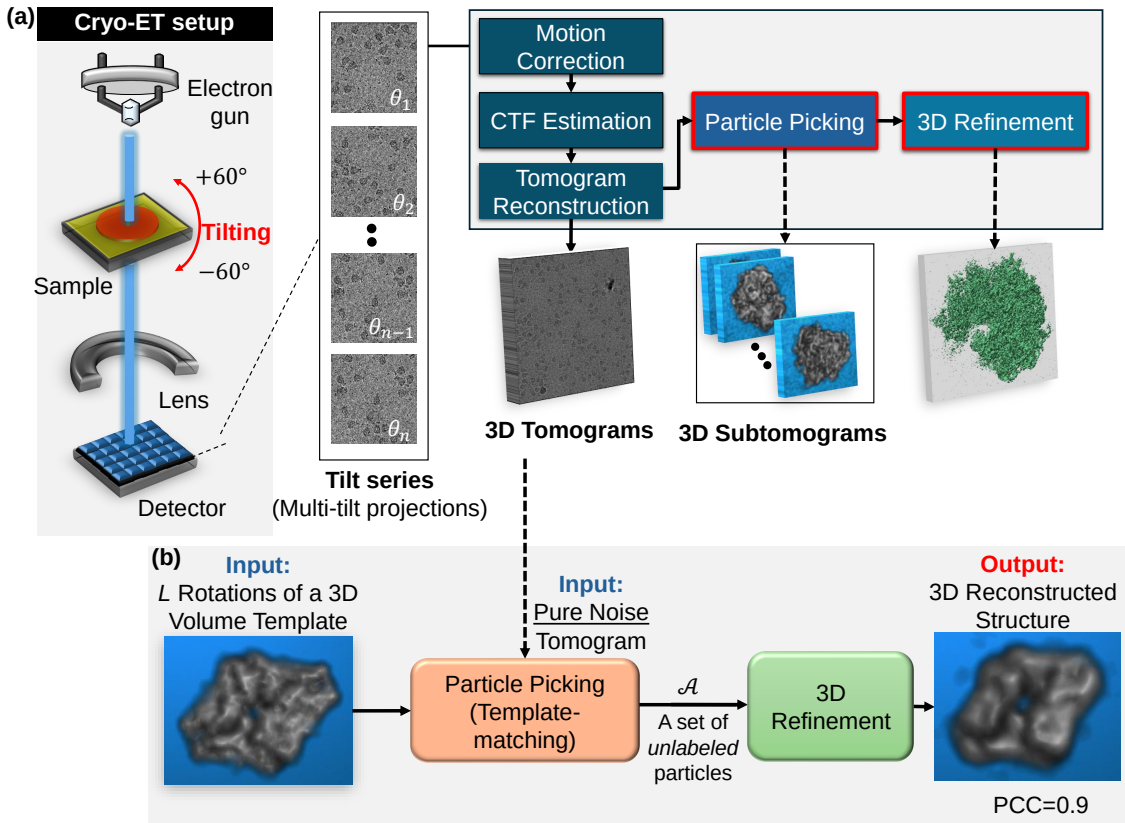

Figure 2: Structure from noise in the particle-picking stage of the cryo-electron tomography (cryo-ET) pipeline. (a) The cryo-ET workflow typically involves imaging a sample from multiple tilt angles (commonly ranging from $-60°$ to $+60°$) to generate 2D projections. These projections are then computationally combined to reconstruct a 3D tomogram. The particle-picking algorithm extracts subtomograms after applying motion correction, tilt-series alignment, and CTF estimation and correction to the tilt series [56]. The extracted particles are then subjected to 3D refinement using subtomogram averaging methods, improving the SNR of the reconstructed 3D volume. **(b)** Confirmation bias during the particle-picking stage: Template matching (as described in Algorithm 1) is applied to a tomogram consisting entirely of noise. The templates are $L = 50$ distinct rotations of a 3D volume. The particle picking produces a series of subtomograms ($M = 2 \times 10^4$), which are used to construct a 3D structure using the expectation-maximization algorithm. Evidently, the resulting 3D structure is highly correlated with the 3D volume used to produce the templates (the Pearson cross-correlation (PCC) is 0.9). The size of the volumes is $d = 24 \times 24 \times 24$, and the template-matching threshold is $T = 3.5$.

theoretical framework and validate it through numerical experiments. While a comparable theory for deep learning methods is currently unavailable, we conduct extensive numerical experiments using the widely adopted particle picker Topaz [11]. Our analysis demonstrates that both approaches introduce systematic biases, often resulting in reconstructions that reflect prior assumptions rather than faithfully representing the underlying data.

In Section 2, we develop a theoretical framework to analyze how confirmation bias arises in template-matching particle picking methods. To quantify this bias, we introduce a metric that measures the similarity between the extracted particles and the original templates, based on the outputs of two key downstream tasks: 2D classification (relevant to cryo-EM) and 3D reconstruction (relevant to cryo-ET). For the theoretical analysis, we focus on the scenario where the input to the particle-picking algorithm consists entirely of pure noise. Remarkably, under standard conditions, both the resulting 2D class averages (in cryo-EM) and the 3D reconstructions (in cryo-ET) converge toward the particle-picking templates, despite the absence of any real underlying signal in the data.

In Section 3, we complement our theoretical analysis with empirical evidence using a standard cryo-EM software pipeline. We evaluate the output of RELION [25], and Topaz [11] when applied to synthetic datasets, including pure-noise and low-SNR micrographs. We find that even in the presence of real particles, low-SNR conditions can result in hybrid reconstructions that combine true structural features with artifacts resembling the templates. Furthermore, we demonstrate that neural network pickers, such as Topaz, while more challenging to analyze theoretically, also introduce reproducible structural artifacts when applied to noise, indicating that learned priors can bias the computational pipeline.

In Section 4, we discuss the broader implications of these findings for the structural biology community. We analyze the trade-offs involved in key parameter choices, particularly selection thresholds, and propose strategies for mitigating bias in the particle-picking stage. We compare the statistical characteristics of confirmation bias introduced during particle picking with those arising at later stages of the reconstruction pipeline, such as 2D classification and 3D reconstruction. Additionally, we outline potential directions for future research to extend this work.

2 Theory

This section introduces a simplified model to investigate confirmation bias in the particle-picking process, where 2D noisy particles (in cryo-EM) or 3D sub-tomograms (in cryo-ET) are identified and extracted. To gain analytical insights, we analyze a particle picker based on template matching (Algorithm 1). While this model does not fully capture the intricate mechanism of real-world particle-pickers, it offers key theoretical insights into the mechanisms by which confirmation bias can emerge at this early stage of the pipeline. In Section 3, to enrich this analysis, we provide empirically grounded evidence that reveals further intricacies of the phenomenon.

For clarity, formal statements and proofs of the main theoretical results are deferred to the appendix. Informal formulations are provided throughout the main text to aid readability. Throughout the text, $\langle a, b \rangle$ represents the inner product between the vectors a and b.

2.1 The particle-picking process

The simplified particle-picking model. We consider a set of N candidate particles, denoted by $\{y_i\}_{i=0}^{N-1} \in \mathbb{R}^d$, where d is the particle dimension (e.g., $d = 36 \times 36$ for cryo-EM, or $d = 36 \times 36 \times 36$ in cryo-ET). These candidates are assumed to be statistically independent and extracted from disjoint, non-overlapping regions of the micrograph. The particle picker is modeled as a binary function $f : \mathbb{R}^d \to \{0, 1\}$, and outputs 1 for particles, and 0 for non-particles. We denote by $M \leq N$ the number of particles extracted by the particle picker.

Particle-picking based on template-matching. The *template-matching particle-picker*, detailed in Algorithm 1, operates by cross-correlating each candidate particle $\{y_i\}_{i=0}^{N-1}$ with L normalized templates $\{x_\ell\}_{\ell=0}^{L-1}$ [1]. The templates $\{x_\ell\}_{\ell=0}^{L-1}$ are either 2D projections from different orientations of a common underlying 3D template volume V_{template} (in the case of cryo-EM), or different rotations of the 3D template model, without projections (in the case of cryo-ET). An observation is classified as a particle if its correlation with at least one template exceeds a fixed threshold T. The extracted particles are stored in the set $\mathcal{A}$, where the number of extracted particles is $|\mathcal{A}| = M$. It is important to note that the extracted particles in the set $\mathcal{A}$ are *unlabeled*, meaning that for any given particle in $\mathcal{A}$, the specific template with which it has the highest correlation is unknown.

Algorithm 1 `Template-Matching Particle Picker`

Input: L templates $\{x_\ell\}_{\ell=0}^{L-1}$, N candidate particles $\{y_i\}_{i=0}^{N-1}$, and a threshold T.
Output: Set $\mathcal{A}$ of selected particles.

Initialize $\mathcal{A} \leftarrow \emptyset$
for $i = 0$ to $N - 1$ **do**
  if $\max_{0 \leq \ell < L-1} \langle y_i, x_\ell \rangle \geq T$ **then**
    Add y_i to $\mathcal{A}$: $\mathcal{A} \leftarrow \mathcal{A} \cup \{y_i\}$
  end if
end for
return $\mathcal{A}$

Main assumptions. Before presenting the theoretical results, we state the main assumptions required for the analysis. For template-matching-based particle picking, we assume an underlying 3D structure denoted by V_{template}, which serves as the reference volume. In cryo-EM, this template is typically projected into 2D images across a range of orientations, while in cryo-ET, it is rotated directly in 3D. We assume that the input micrographs and tomograms are corrupted by additive white Gaussian noise, consistent with the classical Einstein from noise experiment [45]. While this is a simplification, as experimental cryo-EM and cryo-ET data often contain structured contaminants and exhibit weakly correlated

[1] This normalization corresponds to energy normalization commonly used in cross-correlation filtering in cryo-EM. Templates are normalized to unit norm to ensure correlation scores are directly comparable across templates and observations, preventing scaling bias during template matching.

noise, modeling the noise as white and Gaussian offers a clean theoretical baseline. We note that similar artifacts to the Einstein from Noise effect have also been observed under more general noise models [3].

2.2 Theoretical analysis

To analyze confirmation bias in particle picking, we model the candidate particles $\{y_i\}_{i=0}^{N-1}$ as white Gaussian independent and identically distributed (i.i.d) noise (excerpts of pure noise micrographs), i.e., $\{y_i\}_{i=0}^{N-1} \sim \mathcal{N}(0, \sigma^2 I_{d\times d})$, where $\mathcal{N}(0, \sigma^2 I_{d\times d})$ denotes a multivariate Gaussian distribution with mean zero and a diagonal covariance matrix with variance σ^2.

Confirmation bias analysis through downstream tasks. To quantify the confirmation bias introduced during particle picking, we define a metric that compares the extracted particles to the original templates used in the template-matching process. We assess this bias through two downstream tasks: 2D classification (relevant for cryo-EM), where the extracted particles are grouped into classes based on structural similarity and then averaged to increase the SNR, and 3D reconstruction (relevant for cryo-ET), where the extracted subtomograms are aligned and averaged to reconstruct a 3D structure. In both cases, we evaluate how closely the resulting 2D class averages and 3D reconstructions resemble the initial templates used for picking.

2D classification. The 2D classification step in cryo-EM is often modeled using a Gaussian Mixture Model (GMM), a probabilistic framework in which each particle image is assumed to be drawn from a mixture of Gaussian distributions [40]. In this setting, each component of the mixture corresponds to a distinct viewing direction of the underlying 3D structure.

Formally, the extracted particles $z_0, z_1, \ldots, z_{M-1} \in \mathbb{R}^d$ are modeled as independent samples from the GMM:

$$z_0, z_1, \ldots, z_{M-1} \overset{\text{i.i.d.}}{\sim} \sum_{\ell=0}^{L-1} w_\ell \cdot \mathcal{N}(\mu_\ell, \Sigma), \tag{2.1}$$

where μ_ℓ represents the mean of the ℓ-th Gaussian component (i.e., a 2D class average), w_ℓ its mixing weight (satsifying $w_\ell \geq 0$, $\sum_{\ell=0}^{L-1} w_\ell = 1$), and Σ is a shared covariance matrix across components. In the cryo-EM context, each μ_ℓ corresponds to a projection of the 3D structure from a particular viewing direction. For simplicity, we assume that the number of mixture components L matches the number of templates used in particle picking, although we recognize that this correspondence may not hold in practice.

It is important to emphasize that the GMM assumed in (2.1) does not accurately capture the true distribution of particles obtained via template-based particle picking (see Algorithm 1), which does not follow Gaussian statistics; that is, the statistical model is *misspecified*. This mismatch between the assumed and actual data-generating processes must be carefully considered when analyzing the estimated class averages. A detailed discussion of this model misspecification is provided in Appendix B.2.3.

We note that the 2D classification model described in (2.1) omits a common component of practical workflows: in-plane alignment, which is a standard step in cryo-EM 2D classification

pipelines [25]. This procedure estimates the optimal in-plane rotation and translation of each particle image to align it with a class average. Although we exclude in-plane alignment from our theoretical model to preserve analytical tractability, we empirically investigate its effect in Section 3 and show that it does not affect the core properties identified in the simplified model of (2.1).

The goal of 2D classification is to estimate the component means $\{\mu_\ell\}_{\ell=0}^{L-1}$ based on the observed data. A standard approach is maximum likelihood estimation (MLE) [55, 42], yielding estimates $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$. This leads to a key question:

> *To what extent do the estimated class averages $\{\hat{\mu}_\ell\}$ resemble the original templates $\{x_\ell\}$ used in particle picking?*

Ideally, if the estimation process is unbiased, the reconstructed class averages should show no correlation with the templates. In particular, for input particles consisting purely of noise, the class averages should approach zero. Any systematic correlation observed would indicate bias introduced earlier in the pipeline, likely by the particle picking algorithm.

2D classification theoretical results. In Appendix B.2 and Theorem C.2, we prove that the estimated centroids produced by the 2D classification process, $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$, converge to the original templates $\{x_\ell\}_{\ell=0}^{L-1}$ (up to a scaling factor), when the input candidate particles $\{y_i\}_{i=0}^{N-1}$ consist purely of white Gaussian noise and both the number of candidate particles and the threshold parameter T are sufficiently large. The result is summarized informally below.

Theorem 2.1 (Informal, cryo-EM confirmation bias via 2D classification)**.** *Let $\{x_\ell\}_{\ell=0}^{L-1}$ be a set of templates, and let the candidate particles $\{y_i\}_{i=0}^{N-1}$ be drawn i.i.d. from the Gaussian distribution $\mathcal{N}(0, \sigma^2 I_{d\times d})$. Let $\mathcal{A}$ denote the set of extracted particles selected by Algorithm 1 using a threshold parameter T. Let $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ be the maximum-likelihood estimates of the GMM defined in (2.1), computed from the extracted particles $\mathcal{A}$. Then,*

$$\lim_{N,T\to\infty} \frac{\hat{\mu}_\ell}{T} = x_\ell, \quad \ell = 0, \ldots, L-1. \tag{2.2}$$

This result reveals a pronounced bias introduced by template-matching-based particle picking: the estimated class centroids retain full correlation (Pearson cross-correlation (PCC) of 1) with the templates used in the particle-picking stage, even when the input data is pure noise. This is in stark contrast to the unbiased expectation that averaging pure noise yields zero as the number of samples grows.

Figure 3(a) illustrates the effect of a high threshold on the bias. At low threshold values, the centroids do not align with the original templates because the extracted particles are not correctly associated with their corresponding classes. However, as the threshold increases, the estimated centroids progressively align more closely with the original templates. This result reveals a trade-off in selecting the threshold value during the particle picking stage, which will be further discussed in Section 4.

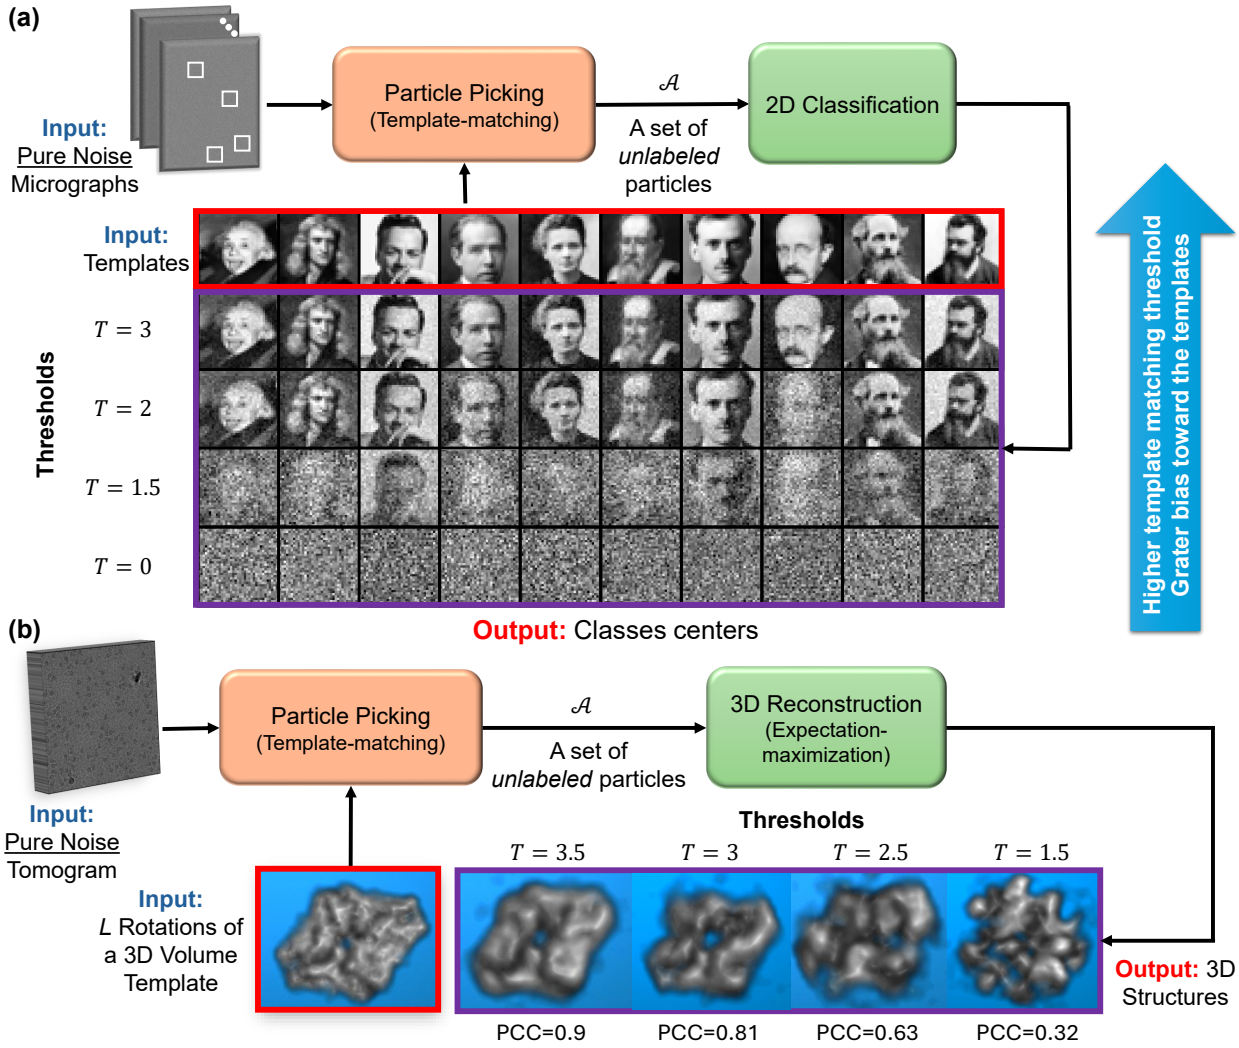

Figure 3: Effect of template-matching thresholds on confirmation bias during the particle-picking stage. (a) 2D Classification Setup. In this example, particle picking is performed on pure noise micrographs using templates of 10 prominent physicists. The resulting particles, which are unlabeled, are then classified using the standard expectation-maximization algorithm. The results demonstrate that increasing the template matching threshold causes the cluster centers to become increasingly biased toward the templates used in the particle-picking stage. This simulation uses $M = 10^5$ particles collected during the particle-picking process and image sizes of $d = 50 \times 50$. The data generation model for this experiment is the same as presented in Section 2.1. **(b) 3D Reconstruction Setup.** In this setup, particle picking is performed on a pure noise tomogram using 50 templates generated by rotating the beta-galactosidase volume [7]. The extracted subtomograms are subjected to 3D refinement, producing a 3D structure. As with the 2D case in panel (a), the results indicate that increasing the template-matching threshold results in the reconstructed 3D structure becoming more biased toward the 3D template used during the particle-picking stage. For each threshold level, the Pearson cross-correlation (PCC) between the 3D reconstructed structure and the 3D volume template is shown. The simulation parameters, based on those outlined in Algorithm 1, are a volume size of $d = 24 \times 24 \times 24$, $L = 50$ distinct template rotations, $M = 2 \times 10^4$ particles collected, noise variance of $\sigma^2 = 1$ and varying template-matching thresholds.

The 3D reconstruction process overview. In cryo-ET, particle picking is performed in 3D, where volumetric subtomograms are extracted from a reconstructed tomogram. This contrasts with cryo-EM, where particle picking operates on 2D micrographs. After 3D particle picking in cryo-ET, a 3D reconstruction step follows, in which the extracted subtomograms are treated as randomly rotated, noise-corrupted observations of an underlying 3D structure V. Formally, the *postulated* statistics of the extracted particle z_i is modeled as:

$$z_i = R_i \cdot V + \epsilon_i, \tag{2.3}$$

where V is the underlying 3D volume, $\{R_i\}_i$ are random 3D rotations, and $\{\epsilon_i\}_{i=0}^{M-1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_V^2 I_{d\times d})$ represents additive Gaussian noise. The goal of the 3D reconstruction process is to estimate the volume V in (2.3), from a set of noisy, random-rotated particles $\{z_i\}_{i=0}^{M-1}$, typically using MLE, yielding an estimate, denoted $\hat{V}$. However, as in the case of 2D classification in cryo-EM, the model in (2.3) does not accurately capture the true distribution of particles obtained through template-based particle picking. In practice, the picking process introduces selection bias and systematic deviations, meaning that the extracted subtomograms do not follow the postulated model. In other words, the statistical model is *misspecified*.

For analytical tractability, we assume that the continuous space of 3D rotations is discretized into a finite set $\{R_\ell\}_\ell$, and that the same set is used both during the template-matching stage of particle picking and in the subsequent 3D reconstruction procedure. This alignment of rotation grids simplifies the analysis and is elaborated further in Appendix B.3. Then, the central question raised is how the reconstructed 3D structure, $\hat{V}$, depends on the underlying template V_{template} used in the template-matching particle-picking process when the input subtomograms consist entirely of white Gaussian noise?

The 3D reconstruction theoretical results. This question is formally addressed in Theorem C.3, which demonstrates that when the candidate subtomograms $\{y_i\}_{i=0}^{N-1}$ provided to the particle picker in Algorithm 1 consist entirely of Gaussian white noise, the resulting reconstructed 3D volume $\hat{V}$ closely resembles the initial template volume V_{template} used during particle picking. This resemblance becomes increasingly pronounced in the asymptotic regime, when the number of selected particles $M \to \infty$ and the selection threshold $T \to \infty$. We provide an informal statement of this result below.

Theorem 2.2 (Informal, cryo-ET confirmation bias via 3D reconstruction)**.** *Let V_{template} be the template volume used in particle picking, and let $\{R_\ell\}_{\ell=0}^{L-1}$ be a discrete set of 3D rotations. Assume the input templates are $\{R_\ell \cdot V_{\mathsf{template}}\}_{\ell=0}^{L-1}$. Let the candidate subtomograms $\{y_i\}_{i=0}^{N-1}$ be drawn i.i.d. from $\mathcal{N}(0, \sigma^2 I_{d\times d})$, and let $\mathcal{A}$ denote the set of particles selected by Algorithm 1 using threshold T. Let $\hat{V}$ be the MLE of the 3D reconstruction applied to the extracted particles $\mathcal{A}$. Then,*

$$\lim_{N,T\to\infty} \frac{\hat{V}}{T} = V_{\mathsf{template}}. \tag{2.4}$$

In words, when the threshold T is high, and the number of particles is large, then the reconstructed structure satisfies $\hat{V} \approx T \cdot V_{\mathsf{template}}$. The important role of the threshold is

demonstrated in Figure 3(b), with behavior analogous to that observed in the 2D classification setting. Appendix C.4 further analyzes the convergence rate in (2.4) as $T \to \infty$, showing its dependence on the noise variance σ_V^2.

Theorem 2.2 relies on the assumption that the rotation samples used to generate the template set are identical to the samples employed during the 3D reconstruction process. When this alignment condition is not satisfied, the reconstructed volume $\hat{V}$ no longer exactly reproduces the template. Nonetheless, as illustrated empirically in Figure 3(b), the reconstructed volume remains highly correlated with the original template.

The impact of the number of extracted particles and particle size on confirmation bias. The analytical results in this section are derived under the asymptotic regime where the number of particles $M \to \infty$. However, practical applications involve a large but finite number of particles, making it essential to understand the behavior in this setting. In particular, we consider the mean squared error $\mathbb{E}_{\mathcal{A}}\left[\|\hat{\mu}_\ell(\mathcal{A}) - T \cdot x_\ell\|^2\right]$, which measures the difference between the estimated class centroids $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ and the ground truth templates $\{x_\ell\}_{\ell=0}^{L-1}$, scaled by the selection threshold T, where the expectation is taken with respect to the random set of extracted particles $\mathcal{A}$. This error scales as $\mathcal{O}\left(\frac{d}{M}\right)$,[2] where d denotes the ambient dimension of each particle image (i.e., the number of pixels or voxels).

This scaling is well established in classical statistical estimation theory for Gaussian mixture models with shared covariance, reflecting the fact that higher-dimensional data require more samples to achieve a comparable level of estimation accuracy. This implies that smaller particles (i.e., lower d) require fewer samples to converge to the templates. Consequently, smaller particles exhibit faster convergence and are more susceptible to confirmation bias compared to larger particles (see Appendix C.4 for further details).

3 Empirical results

This section presents empirical results that corroborate and extend the theoretical findings from Section 2. We examine how deviations from the simplified model, such as overlapping candidate particles, the presence of true particles (i.e., not only pure noise micrograph), and the use of in-plane alignment, affect the outcomes of 2D classification and 3D reconstruction. Additionally, we evaluate the behavior of modern particle-picking approaches, in particular deep learning-based methods, to assess their susceptibility to confirmation bias in realistic cryo-EM workflows.

Importantly, the 2D classification and 3D reconstruction methods employed in our analysis are reference-free, i.e., they rely on random initialization rather than incorporating prior templates. This design allows us to isolate the confirmation bias introduced during picking. However, we note that the non-convex nature of the downstream reconstruction tasks may also contribute to residual bias.

[2]Here, $\mathcal{O}(d/M)$ means the mean squared error is bounded above by a constant times d/M as $M \to \infty$.

3.1 Template-matching setup

In the previous section, we introduced a simplified model for particle picking based on template matching for the purpose of the analytic analysis, assuming that the input micrograph or tomogram consisted solely of noise. Under the assumptions of a sufficiently large number of particles and a high selection threshold, we demonstrated that the 2D class centers (in cryo-EM) or the reconstructed 3D structure (in cryo-ET) closely resemble the templates. While this model offers important insights into the confirmation bias introduced by template-matching particle-picking, it does not fully reflect the complexity of modern cryo-EM reconstruction pipelines. To address this gap, we now turn to empirical experiments that reflect more realistic workflows. Below, we outline the experimental setup and highlight the key differences from the theoretical model. The corresponding results are presented in Figure 4.

Micrograph generation. In the theoretical analysis presented in Section 2, we examined the case where the input micrograph consists solely of pure noise. This model can be considered an asymptotic case where the SNR of the micrograph approaches zero. However, in practice, true particles are located at unknown positions within the micrograph, with $\mathsf{SNR} > 0$.

We now describe the synthetic process of generating micrographs. The particle placement occurs iteratively. In each iteration, a random location is selected within the micrograph as a potential position for a particle. Prior to placing the particle, the algorithm verifies that the selected position does not overlap with any previously placed particle. This process continues iteratively until all the particles are successfully inserted into the micrograph, ensuring that they do not overlap. We used micrographs of size 2048×2048, each containing 200 projections, randomly sampled from a predefined set of 50 projections of the same structure. Each projection is of size of $d = 36 \times 36$. White Gaussian noise $\mathcal{N}(0, \sigma^2 I_{d\times d})$ was added to each micrograph to maintain SNR levels of approximately 1/25, where $\mathrm{SNR} \triangleq \frac{\mathrm{Var}(x_i)}{\sigma^2}$ and x_i denotes a normalized particle projection. Using a template matching threshold of $T = 4$, we collected a total of $M = 2 \times 10^5$ particles, corresponding to ~ 350 particles picked per micrograph, approximately 75% more than the 200 true planted particles.

Template matching on micrographs. In the theoretical framework presented in Section 2, the candidate particles $\{y_i\}_{i=0}^{N-1}$ were assumed to be drawn i.i.d., corresponding to non-overlapping, disjoint regions of a micrograph composed purely of noise. However, in practical particle picking from micrographs, candidates are identified using a sliding window approach, which involves overlapping regions and thus introduces statistical dependencies between the candidate particles. In particular, to avoid selecting overlapping regions, the detection and extraction process enforces a non-overlapping constraint on the chosen particle patches. A more realistic simulation of this procedure, including the construction of synthetic micrographs and the behavior of template matching under overlapping conditions, is presented in Appendix A and described in detail in Algorithm 2.

2D classification. The theoretical model presented in Section 2 does not account for in-plane alignment and assumes that the number of templates used during particle picking matches the number of class centers in the 2D classification stage, an assumption that rarely

holds in practice. In this section, we incorporate more realistic scenarios that reflect typical cryo-EM workflows, including the application of in-plane alignment and differing numbers of templates and 2D class centers.

To align with standard procedures in the field, we employ RELION's VDAM algorithm for 2D classification [25]. RELION was chosen due to its widespread adoption in the cryo-EM community and its proven track record for robust and reproducible results. The VDAM algorithm, introduced in RELION 4.0, is a variable-metric gradient descent method with adaptive moment estimation. Inspired by techniques from deep learning, particularly the Adam optimizer [16], VDAM maintains exponential moving averages of both the first and second moments of the gradient, enhancing stability and convergence. It also operates on mini-batches, which helps reduce computational costs. The input particles to RELION were extracted using the procedure described in Algorithm 2. For 2D classification, we configured VDAM with the following parameters: $L = 20$ distinct classes, 200 iterations, a regularization parameter set to 3, and no application of CTF correction. We perform in-plane alignment with an angular sampling of 12 degrees.

3D ab-initio modeling. We also performed ab initio 3D reconstruction using RELION's VDAM algorithm. In this context, ab initio reconstruction refers to generating an initial 3D structure without relying on any external reference, thereby reducing potential sources of bias. We employed RELION's VDAM algorithm for this task with the following parameters: $L = 1$ class, 200 iterations, a regularization parameter of 4, no CTF correction, and no symmetry constraints (i.e., $C1$ symmetry).

Results. Figure 4(b) presents results from the empirical setup in which the micrograph contains genuine particle projections, but the template used for particle picking corresponds to a different structure. In this scenario, the 2D class averages reflect a mixture of two competing sources: the true underlying particles and the mismatched template. As a result, while some class averages recover features of the actual particles, others resemble the template used for picking. Similarly, the ab initio 3D reconstruction derived from these particles exhibits structural elements from both the true volume and the template. This highlights the risk of introducing structural bias at the particle-picking stage, where even in the presence of a signal, the use of an incorrect template can propagate artifacts into downstream analyses and compromise the reliability of the final reconstruction.

3.2 Particle picker based on Topaz neural network

Deep learning-based methods have become a widely used approach for particle picking in cryo-EM and cryo-ET, with Topaz being one of the most prominent examples [11]. Unlike template-matching approaches, which introduce bias through direct similarity to predefined structures, deep neural networks (DNNs) introduce bias through more subtle mechanisms. These include the composition of the training dataset (which may reflect prior structural assumptions), architectural choices, and the learning dynamics of the training process itself. Although the mechanisms differ, such models can still exhibit confirmation bias by preferentially selecting particles that resemble features emphasized during training. In our

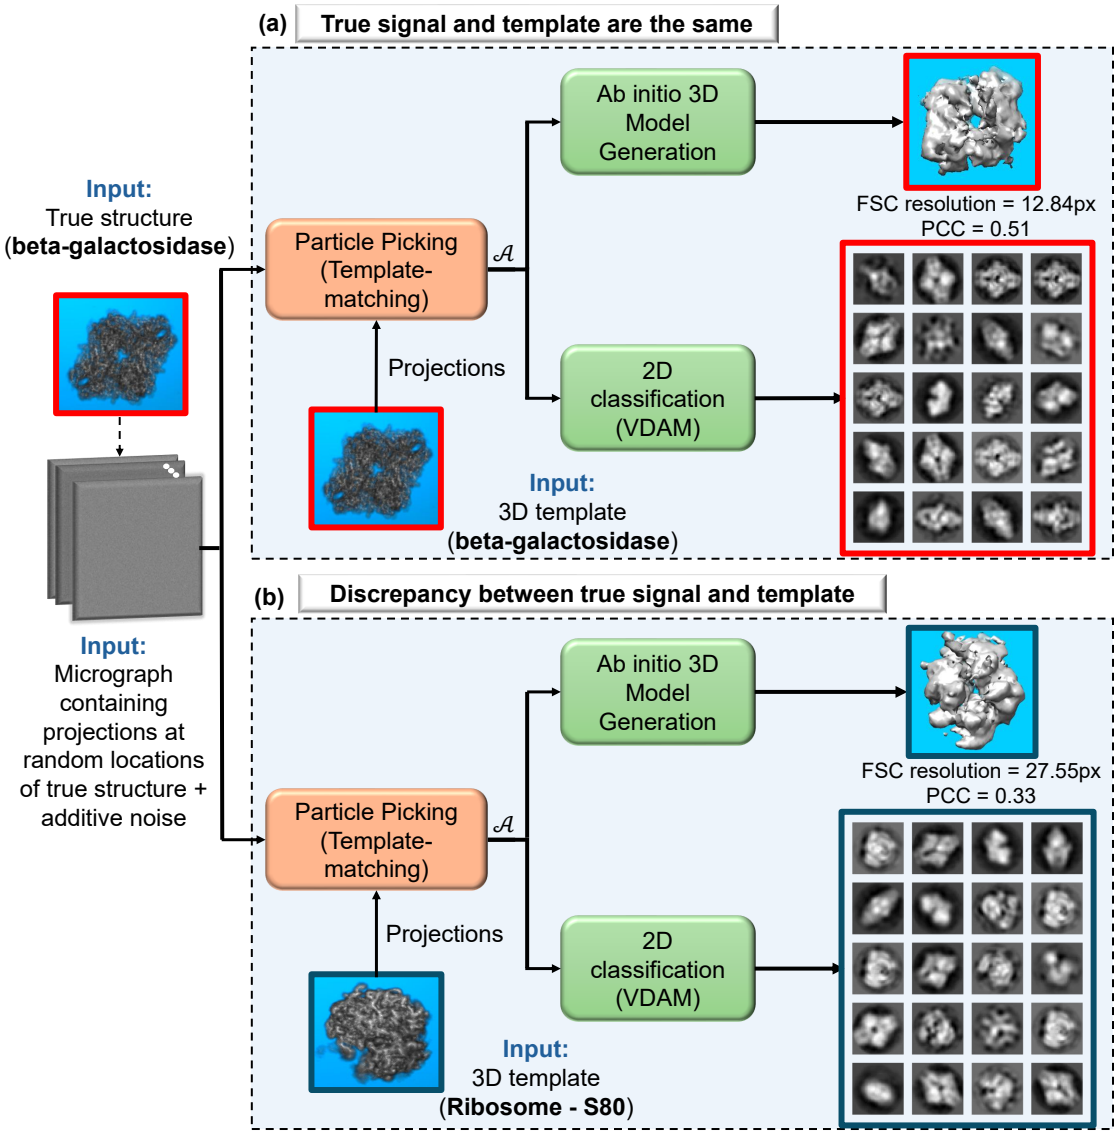

Figure 4: Confirmation bias in particle picking with particles present in the micrograph. We examine a scenario where the micrograph contains true particles, specifically, projections of the Beta-galactosidase 3D structure [57], corrupted by additive white Gaussian noise, with an SNR of approximately 1/25. Template matching is performed using two different sets of templates: **(a)** projections of the true Beta-galactosidase structure (ground truth), and **(b)** projections of a different structure, the Plasmodium falciparum 80S ribosome [7]. For both cases, we apply 2D classification and ab initio 3D reconstruction using RELION's VDAM algorithm [25]. In Panel **(a)**, the 2D classes align closely with the true Beta-galactosidase projections, and the 3D reconstruction shows a strong resemblance to the original structure (FSC resolution of 12.84 pixels, and Pearson cross-correlation (PCC) = 0.51). In contrast, Panel **(b)** reveals a mixture: the 2D classification captures both genuine Beta-galactosidase projections and spurious classes that resemble the ribosome templates. The resulting 3D reconstruction is a hybrid structure, partially resembling both the ground truth and the incorrect template, with a lower PCC of 0.33 and FSC resolution of 27.55 pixels. In both experiments, projection patches were of size $d = 36 \times 36$, and $M = 2 \times 10^5$ particles were extracted using a template-matching threshold of $T = 4$. The VDAM reconstructions were performed without CTF correction over 200 iterations. Additional implementation details are provided in Appendix A.

experiments, we use the widely adopted Topaz network to investigate model bias in deep learning-based particle pickers and find that it also consistently produces structured outputs influenced by its training priors, albeit to a lesser extent than template matching.

To analyze this phenomenon, we used the pre-trained Topaz model provided by Bepler et al. [11] without performing any additional training. We worked with micrographs of size 2048×2048, which were down-sampled in Topaz by a factor of $s = 2$, due to GPU memory constraints. During extraction, we set the radius to $r = 9$ (i.e., original particle projection dimensions were presumed to be $d = 36 \times 36$) and the coordinates scaling factor to $x = 2$ (to account for aforementioned down-sampling).

Figure 5(a–c) illustrates the presence of confirmation bias in Topaz network (as a representative of deep learning-based particle pickers). In Figure 5(a), the average of the particles extracted by Topaz reveals a faint circular structure at the center, which is an unexpected result under a truly unbiased system, where the mean should converge to zero. This artifact closely resembles the central pattern observed in Figure 5(b), which shows the mean of particles selected via template matching using a large and diverse set of templates. The similarity between these results suggests that both approaches introduce structured bias into the picking process. For comparison, Figure 5(c) shows the mean of randomly selected patches from pure noise, which converges to zero, serving as an unbiased reference.

To quantify the impact of confirmation bias on 3D reconstruction, we conducted a controlled experiment using pure-noise micrographs that are split into two independent halves. Each half is processed using one of three particle-picking strategies: template matching, Topaz, or random selection. We then performed independent ab initio 3D reconstructions on each half using the RELION VDAM algorithm. Following standard practice, we assess the consistency between the resulting half-maps using Fourier Shell Correlation (FSC), as shown in Figure 5(d).

Both Topaz and template matching produce reconstructions with high FSC values between the halves, indicating that the selected particles give rise to consistent, method-dependent structures, even though the input data contain only noise. In contrast, random selection yields near-zero FSC values, as expected under an unbiased model. Among the three methods, template matching exhibits the strongest bias, followed by Topaz, with random picking serving as a near-unbiased baseline. These results provide additional empirical evidence that the FSC gold standard can be misled by systematic errors, consistent with observations in [49].

4 Discussion and Outlook

To contextualize the confirmation bias analyzed in this work, we compare it to other sources of bias that arise throughout the cryo-EM and cryo-ET reconstruction pipeline. We then discuss strategies for mitigating this bias and conclude with several promising directions for future research.

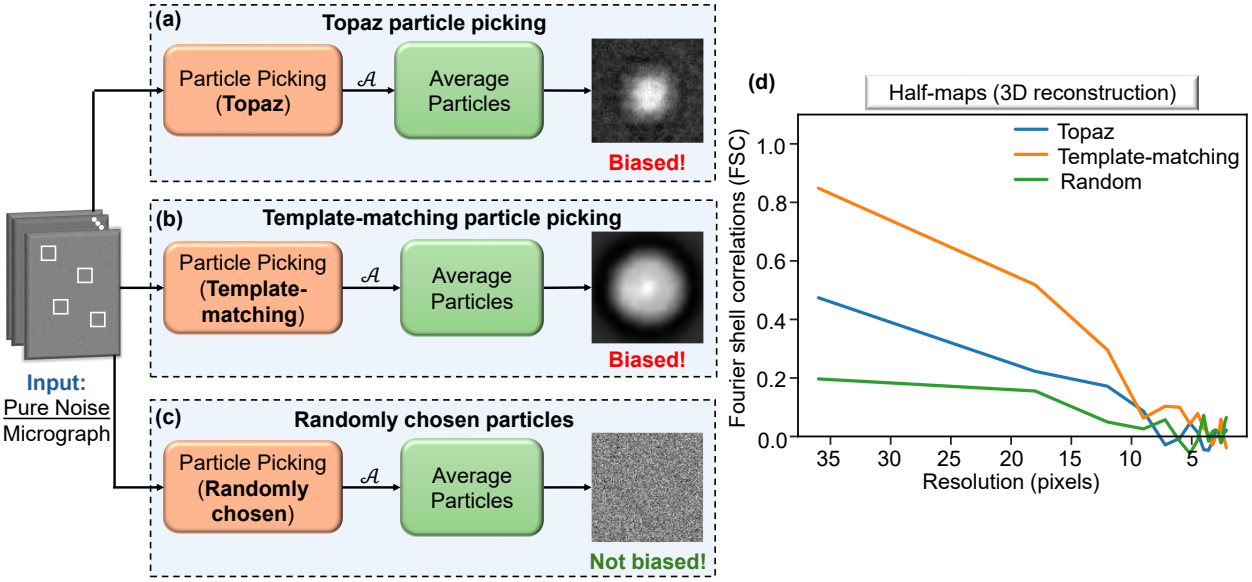

Figure 5: Confirmation bias in Topaz particle picking. (a) The average of particles extracted using the Topaz neural network displays a prominent circular feature at the center. **(b)** Similarly, particle picking via a large set of the Plasmodium falciparum 80S ribosome [7] projection templates yields an average image with a comparable circular pattern. **(c)** In contrast, randomly selecting particles from micrographs results in an average dominated by noise, which converges to zero. Further details on the Topaz network parameters are provided in Appendix A. **(d)** To assess bias, micrographs containing pure noise were split into two independent halves. On each half, particles were picked using three methods: (1) the Topaz network, (2) template matching, and (3) random selection. Each set of picked particles was used to perform independent 3D reconstructions with RELION's VDAM algorithm. The Fourier Shell Correlation (FSC) between the resulting half-maps was then computed. Both the Topaz and template-matching methods yielded significantly higher FSC values than random picking, indicating the presence of consistent biases capable of generating structured output from pure noise. Among the three, template matching exhibited the strongest bias, followed by Topaz, while random picking introduced the least bias. In all experiments, projection size was set to $d = 36 \times 36$, and $M = 10^5$ particles were extracted for each half-map. The VDAM reconstructions were run for 200 iterations without CTF correction; additional technical details are provided in Appendix A.

4.1 Different statistical models for confirmation bias in cryo-EM

Several stages of the cryo-EM and cryo-ET processing pipeline are vulnerable to confirmation bias, stemming from the use of initial templates or strong prior assumptions [49, Section 4.1]. Key examples include: (1) Particle picking, where 2D particles (in cryo-EM) or 3D subtomograms (in cryo-ET) are extracted using template-based or learning-based methods [43, 21, 11, 18]; (2) 2D classification (cryo-EM), where particles are clustered and averaged; (3) 3D reconstruction, including subtomogram averaging in cryo-ET, which is typically carried out using the expectation-maximization algorithm. As demonstrated and analyzed in [4], these algorithms are especially vulnerable to bias in low SNR conditions, where unverified assumptions can significantly shape the final reconstruction.

Each of these stages introduces bias through different statistical mechanisms. Figure 6

compares the forms of confirmation bias that occur during particle picking with those that emerge in downstream tasks, such as 2D classification. Below, we describe and contrast the Einstein from Noise phenomenon, which was presented in Section 1 and the confirmation bias observed in 2D classification with the bias arising from particle picking.

The Einstein from Noise phenomenon. Although frequently mentioned in the context of template-based particle picking, the Einstein from Noise effect arises from a distinct model. In the Einstein from Noise experiment, all observations are aligned and averaged with respect to a fixed template. In contrast, particle picking involves selecting a small subset of patches from a large micrograph, specifically, those with the highest cross-correlation scores relative to the template. This key distinction gives rise to different types of confirmation bias and artifacts. Particle picking tends to produce reconstructions that closely resemble the template, differing only by a global scale factor (Theorems 2.1 and 2.2). In contrast, the Einstein from Noise effect produces an average that is positively correlated with the template but not identical [3].

Confirmation bias in 2D classification and 3D volume reconstruction. In later stages of cryo-EM processing, such as 2D classification and 3D volume estimation, iterative algorithms often use predefined templates for initialization. When applied to data with no true signal (e.g., white Gaussian noise), these methods may still produce output that resembles the initial templates, even though the input data provides no evidence for such structure [2].

Figure 6(b) illustrates this phenomenon in the context of 2D classification. In this example, the input consisted entirely of pure noise, and the initial templates were set to $L = 30$ distinct projections of the Beta-galactosidase structure, each of size $d = 36 \times 36$. A total of $M = 5 \times 10^5$ particles were processed using a 2D classification based on the expectation-maximization algorithm, run for 20 iterations. Even without a signal, the resulting class averages exhibit a positive correlation with the templates. However, unlike the case of template-based particle picking, the final class averages do not match the templates exactly (even up to a scale factor), though they still appear visually similar. It is worth noting that many modern 2D classification and 3D ab initio reconstruction algorithms now adopt reference-free initialization strategies (e.g., random initializations), which help mitigate such bias [25]. Nevertheless, when templates are introduced, whether explicitly or implicitly, similar confirmation effects can still occur.

Hallucinated heterogeneity in conformational analysis. A distinct form of confirmation bias arises in methods for continuous heterogeneity analysis, particularly when state-of-the-art models are applied to structurally homogeneous datasets. In [20], the authors benchmarked RECOVAR [20], 3DFlex [38], and cryoDRGN [60] on a synthetic homogeneous dataset with high levels of noise. Despite the data consisting of a single conformation, all three methods showed signs of hallucinated heterogeneity. This demonstrates that both deep-learning methods (3DFlex and cryoDRGN) and classical statistical methods (RECOVAR) are not immune to this form of bias.

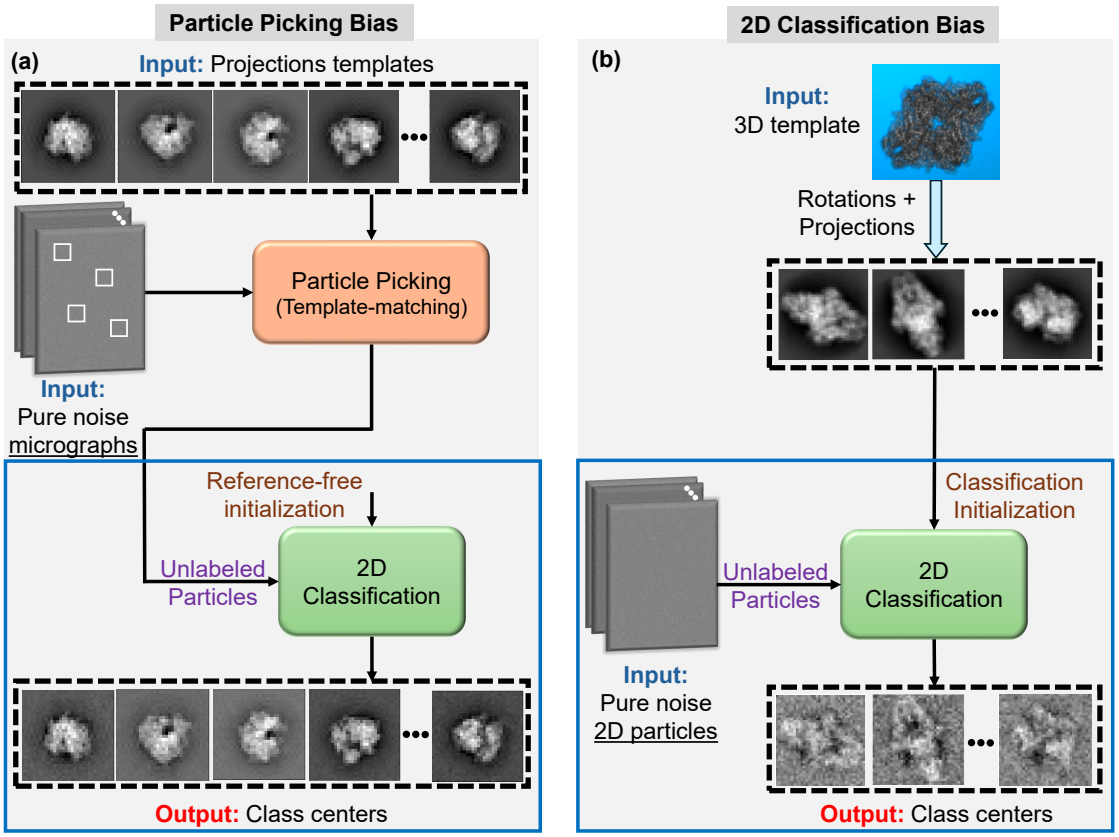

Figure 6: Comparison of confirmation bias in particle picking versus 2D classification. **(a)** Confirmation bias during the particle-picking stage, the primary bias mechanism analyzed in this work. **(b)** Confirmation bias during the 2D classification stage. Here, the input consists solely of pure noise particles, and the classification is initialized using predefined templates. Despite the absence of signal in the data, the resulting class averages exhibit a bias toward the initial templates, indicating a positive correlation (i.e., confirmation bias) [2]. Notably, the class centers in Panel (a) appear to more closely resemble the initial templates than those in Panel (b).

4.2 Strategies for mitigating confirmation bias in particle picking

Previous studies have explored the detection and mitigation of various forms of bias in cryo-EM. For example, recent work [49] has proposed several practical strategies to reduce reconstruction bias, including identifying bias at the level of estimated parameters, performing multiple independent reconstructions, averaging parameter estimates to reduce variance, applying phase randomization at the image level, and validating algorithms across diverse preprocessing and reconstruction settings. Despite these advances, the particle-picking stage remains a particularly vulnerable and relatively understudied source of bias. Here, we focus specifically on computational strategies to mitigate bias during particle picking, incorporating both practical considerations and new theoretically grounded approaches.

Bypassing particle picking. A promising strategy is to eliminate the particle-picking step altogether. Recent methods for direct structure determination, such as those based on the method of moments [9] or approximations to expectation-maximization [27], aim

to reconstruct 3D volumes directly from raw micrographs, thereby avoiding any particle selection. These approaches offer a principled way to avoid selection bias entirely [5], though their current resolution and computational feasibility remain limited.

Statistical control of selection thresholds. All particle-picking methods inevitably produce both false positives (non-particles mistakenly identified as particles) and false negatives (true particles that are missed). A common practice today is to tune picking parameters to favor high recall, that is, to select as many candidate particles as possible. This strategy aims to maximize the extraction of structural information from the micrographs [49].

In particular, in template-matching based approaches (see Algorithm 1), the selection threshold T plays a crucial role in this process and has a significant impact on the extent of confirmation bias introduced during particle selection. A high threshold offers certain advantages: (1) it reduces false positives by selecting only the strongest matches to the templates, increasing the likelihood that picked particles represent true biological structures; and (2) it streamlines downstream processing by limiting the number of particles. However, this approach comes with trade-offs. First, as demonstrated in Section 2, a high threshold can amplify confirmation bias, skewing reconstructions toward the input templates even when the data consists of pure noise. Second, it reduces the number of selected particles, potentially weakening the statistical power of the reconstruction and limiting achievable resolution.

Despite its importance, the threshold T is typically chosen heuristically or manually, without rigorous statistical justification. A more principled alternative is to set thresholds using multiple hypothesis testing frameworks that control the false discovery rate (FDR)[10]. These techniques guarantee that the expected proportion of false positives remains below a predefined level, while typically achieving high statistical power. As a result, they help reduce the selection of spurious particles that arise due to random correlations with the templates. However, these methods are often specifically tailored to well-specified statistical models and may not generalize easily to more complex or poorly understood settings [19]. Recent developments, such as knockoff-based procedures, offer a promising direction for extending FDR control to more flexible and model-agnostic scenarios by introducing synthetic variables that mimic the data's dependence structure [6].

Template filtering and design. Several tools in current cryo-EM pipelines aim to reduce bias by preprocessing the templates themselves [42]. For example, low-pass filtering the templates reduces their high-frequency features, which are most susceptible to overfitting by noise. This encourages the selection of broader structural motifs rather than fine-grained details, mitigating high-frequency confirmation bias. Our theoretical model (see Section 2) shows that this practice leads to reconstructions that converge toward low-pass-filtered versions of the templates. While not a complete remedy, such filtering is routinely employed in practice and is supported by empirical evidence, including the analysis presented in [3].

Template-free picking and alternative similarity measures. In addition to template-based methods, several template-free approaches have been proposed that aim to detect candidate particles in an unsupervised manner, without relying on external references [21]. An

alternative strategy within the template-based paradigm is to refine the similarity measures used during matching. Instead of standard cross-correlation, more robust alternatives, such as cross-entropy-based criteria [48], have been suggested. While promising, the development and validation of such metrics remain largely open challenges.

Taken together, these approaches represent complementary strategies for reducing confirmation bias in particle picking, ranging from algorithmic adjustments to theoretical diagnostics. Continued development of such tools will be essential for ensuring the reliability of cryo-EM reconstructions, especially in low-SNR regimes.

4.3 Future work

While the theoretical framework developed here lays the groundwork for understanding confirmation bias in single-particle reconstruction, several important directions remain open for future research.

First, a natural extension of our model involves incorporating in-plane alignment, both rotational and translational, as it is an essential component of modern 2D classification algorithms. This would enable a more comprehensive analysis of bias introduced during iterative alignment procedures. Second, real-world 2D classification often has a mismatch between the number of templ, and and the number of classes, leading to centroid merging and other nonlinear effects not captured by current models. Third, our current analysis assumes white Gaussian noise, whereas actual cryo-EM micrographs are dominated by colored and structured noise stemming from factors such as the microscope's point spread function, and background ice contamination. Finally, an important future work direction is to theoretically analyze the case where true particles are embedded in low-SNR micrographs. Overall, modeling such settings more accurately would enhance the applicability of our framework to real experimental conditions.

Extension to 3D reconstruction in cryo-EM. The theoretical framework presented here regarding the cryo-EM process applies to a 2D classification process modeled as a GMM. It is important to note that in cryo-EM, the different projections (or class centers) exhibit a specific algebraic structure. This structure arises from the statistical generative model of the observations:

$$y_i = \Pi\left(R_i \cdot V\right) + \epsilon, \tag{4.1}$$

where Π denotes the tomographic projection operator, R_i represents a random 3D rotation operator, such that $R_i \cdot V \triangleq V\left(R_i^{-1}x\right)$, V is the underlying 3D volume, and ϵ accounts for additive Gaussian noise. In other words, all observations are noisy projections of randomly rotated versions of the same volume V.

An important extension of the theoretical framework is to analyze the full 3D reconstruction problem, in which the objective is to recover the volume V, from the extracted particles $\{y_i\}$ defined in (4.1). This leads to a central question for future research: To what extent does the estimated volume reflect the actual structure present in the data versus the template volume V_{template} that was used during particle picking? Addressing this question would provide deeper insight into how early-stage biases introduced by template-based picking propagate through the full cryo-EM pipeline and affect final structural interpretations.

Acknowledgment

T.B. is supported in part by BSF under Grant 2020159, in part by NSF-BSF under Grant 2019752, in part by ISF under Grant 1924/21, and in part by a grant from The Center for AI and Data Science at Tel Aviv University (TAD). We thank Amitay Eldar for insightful discussions and for his contributions to the initial formulation of the framework. T.B. thanks Amit Singer and Marc Gilles for bringing the phenomenon of hallucinated heterogeneity in conformational analysis to his attention.

References

[1] Xiao-Chen Bai, Greg McMullan, and Sjors HW Scheres. How cryo-EM is revolutionizing structural biology. *Trends in biochemical sciences*, 40(1):49–57, 2015.

[2] Amnon Balanov, Tamir Bendory, and Wasim Huleihel. Confirmation bias in Gaussian mixture models. *arXiv preprint arXiv:2408.09718*, 2024.

[3] Amnon Balanov, Wasim Huleihel, and Tamir Bendory. Einstein from noise: Statistical analysis. *arXiv preprint arXiv:2407.05277*, 2024.

[4] Amnon Balanov, Wasim Huleihel, and Tamir Bendory. Expectation-maximization for multi-reference alignment: Two pitfalls and one remedy. *arXiv preprint arXiv:2505.21435*, 2025.

[5] Amnon Balanov, Shay Kreymer, and Tamir Bendory. A note on the sample complexity of multi-target detection. *arXiv preprint arXiv:2501.11980*, 2025.

[6] Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. *The Annals of statistics*, pages 2055–2085, 2015.

[7] Alberto Bartesaghi, Doreen Matthies, Soojay Banerjee, Alan Merk, and Sriram Subramaniam. Structure of β-galactosidase at 3.2-Å resolution obtained by cryo-electron microscopy. *Proceedings of the National Academy of Sciences*, 111(32):11709–11714, 2014.

[8] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE signal processing magazine*, 37(2):58–76, 2020.

[9] Tamir Bendory, Nicolas Boumal, William Leeb, Eitan Levin, and Amit Singer. Toward single particle reconstruction without particle picking: Breaking the detection limit. *SIAM Journal on Imaging Sciences*, 16(2):886–910, 2023.

[10] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.

[11] Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J Noble, and Bonnie Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nature methods*, 16(11):1153–1160, 2019.

[12] Jochen Böhm, Achilleas S Frangakis, Reiner Hegerl, Stephan Nickell, Dieter Typke, and Wolfgang Baumeister. Toward detecting and identifying macromolecules in a cellular context: template matching applied to electron tomograms. *Proceedings of the National Academy of Sciences*, 97(26):14245–14250, 2000.

[13] Muyuan Chen, James M Bell, Xiaodong Shi, Stella Y Sun, Zhao Wang, and Steven J Ludtke. A complete data processing workflow for cryo-ET and subtomogram averaging. *Nature methods*, 16(11):1161–1168, 2019.

[14] Sergio Cruz-León, Tomáš Majtner, Patrick Hoffmann, Jan P Kreysing, Maarten W Tuijtel, Stefan L Schaefer, Katharina Geißler, Martin Beck, Beata Turoňová, and Gerhard Hummer. High-confidence 3D template matching for cryo-electron tomography. *Biophysical Journal*, 123(3):183a, 2024.

[15] Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.

[16] Kingma Diederik. Adam: A method for stochastic optimization. *(No Title)*, 2014.

[17] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.

[18] Amitay Eldar, Boris Landa, and Yoel Shkolnisky. KLT picker: Particle picking using data-driven optimal templates. *Journal of structural biology*, 210(2):107473, 2020.

[19] Amitay Eldar, Keren Mor Waknin, Samuel Davenport, Tamir Bendory, Armin Schwartzman, and Yoel Shkolnisky. Object detection under the linear subspace model with application to cryo-EM images. *arXiv preprint arXiv:2405.00364*, 2024.

[20] Marc Aurèle Gilles and Amit Singer. Cryo-EM heterogeneity analysis using regularized covariance estimation and kernel regression. *Proceedings of the National Academy of Sciences*, 122(9):e2419140122, 2025.

[21] 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.

[22] Richard Henderson. Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise. *Proceedings of the National Academy of Sciences*, 110(45):18037–18041, 2013.

[23] Zhong Huang and Pawel A Penczek. Application of template matching technique to particle detection in electron micrographs. *Journal of Structural Biology*, 145(1-2):29–40, 2004.

[24] Saul M Kassin, Itiel E Dror, and Jeff Kukucka. The forensic confirmation bias: Problems, perspectives, and proposed solutions. *Journal of applied research in memory and cognition*, 2(1):42–52, 2013.

[25] Dari Kimanius, Liyi Dong, Grigory Sharov, Takanori Nakane, and Sjors HW Scheres. New tools for automated cryo-EM single-particle analysis in RELION-4.0. *Biochemical Journal*, 478(24):4169–4185, 2021.

[26] Joshua Klayman. Varieties of confirmation bias. *Psychology of learning and motivation*, 32:385–418, 1995.

[27] Shay Kreymer, Amit Singer, and Tamir Bendory. A stochastic approximate expectation-maximization for structure determination directly from cryo-EM micrographs. *arXiv preprint arXiv:2303.02157*, 2023.

[28] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[29] Youdong Mao, Luis R Castillo-Menendez, and Joseph G Sodroski. Reply to subramaniam, van heel, and henderson: Validity of the cryo-electron microscopy structures of the HIV-1 envelope glycoprotein complex. *Proceedings of the National Academy of Sciences*, 110(45):E4178–E4182, 2013.

[30] Youdong Mao, Liping Wang, Christopher Gu, Alon Herschhorn, Anik Désormeaux, Andrés Finzi, Shi-Hua Xiang, and Joseph G Sodroski. Molecular architecture of the uncleaved HIV-1 envelope glycoprotein trimer. *Proceedings of the National Academy of Sciences*, 110(30):12438–12443, 2013.

[31] Antonio Martinez-Sanchez. Template matching and machine learning for cryo-electron tomography. *Current Opinion in Structural Biology*, 93:103058, 2025.

[32] Jacqueline LS Milne, Mario J Borgnia, Alberto Bartesaghi, Erin EH Tran, Lesley A Earl, David M Schauder, Jeffrey Lengyel, Jason Pierson, Ardan Patwardhan, and Sriram Subramaniam. Cryo-electron microscopy–a primer for the non-microscopist. *The FEBS journal*, 280(1):28–45, 2013.

[33] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.

[34] Clifford R Mynatt, Michael E Doherty, and Ryan D Tweney. Confirmation bias in a simulated research environment: An experimental study of scientific inference. *Quarterly Journal of Experimental Psychology*, 29(1):85–95, 1977.

[35] Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.

[36] Tao Ni, Thomas Frosio, Luiza Mendonça, Yuewen Sheng, Daniel Clare, Benjamin A Himes, and Peijun Zhang. High-resolution in situ structure determination by cryo-electron tomography and subtomogram averaging using emclarity. *Nature protocols*, 17(2):421–444, 2022.

[37] Eva Nogales. The development of cryo-EM into a mainstream structural biology technique. *Nature methods*, 13(1):24–27, 2016.

[38] Ali Punjani and David J Fleet. 3DFlex: determining structure and motion of flexible proteins from cryo-EM. *Nature Methods*, 20(6):860–870, 2023.

[39] Jean-Paul Renaud, Ashwin Chari, Claudio Ciferri, Wen-ti Liu, Hervé-William Rémigy, Holger Stark, and Christian Wiesmann. Cryo-EM in drug discovery: achievements, limitations and prospects. *Nature reviews Drug discovery*, 17(7):471–492, 2018.

[40] Douglas A Reynolds et al. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663), 2009.

[41] Miroslava Schaffer, Stefan Pfeffer, Julia Mahamid, Stephan Kleindiek, Tim Laugks, Sahradha Albert, Benjamin D Engel, Andreas Rummel, Andrew J Smith, Wolfgang Baumeister, et al. A cryo-FIB lift-out technique enables molecular-resolution cryo-ET within native Caenorhabditis elegans tissue. *Nature methods*, 16(8):757–762, 2019.

[42] Sjors HW Scheres. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *Journal of structural biology*, 180(3):519–530, 2012.

[43] Sjors HW Scheres. Semi-automated selection of cryo-EM particles in RELION-1.3. *Journal of structural biology*, 189(2):114–122, 2015.

[44] Sjors HW Scheres and Shaoxia Chen. Prevention of overfitting in cryo-EM structure determination. *Nature methods*, 9(9):853–854, 2012.

[45] Maxim Shatsky, Richard J Hall, Steven E Brenner, and Robert M Glaeser. A method for the alignment of heterogeneous macromolecules from electron microscopy. *Journal of structural biology*, 166(1):67–78, 2009.

[46] Fred J Sigworth. A maximum-likelihood approach to single-particle image refinement. *Journal of structural biology*, 122(3):328–339, 1998.

[47] Amit Singer and Fred J Sigworth. Computational methods for single-particle electron cryomicroscopy. *Annual review of biomedical data science*, 3:163–190, 2020.

[48] Carlos Oscar S Sorzano, JR Bilbao-Castro, Y Shkolnisky, M Alcorlo, R Melero, G Caffarena-Fernández, M Li, G Xu, R Marabini, and JM Carazo. A clustering approach to multireference alignment of single-particle projections in electron microscopy. *Journal of structural biology*, 171(2):197–206, 2010.

[49] COS Sorzano, Amaya Jiménez-Moreno, David Maluenda, Marta Martínez, Erney Ramírez-Aportela, James Krieger, Roberto Melero, Ana Cuervo, Javier Conesa, J Filipovic, et al. On bias, variance, overfitting, gold standard and consensus in single-particle analysis by cryo-electron microscopy. *Biological Crystallography*, 78(4):410–423, 2022.

[50] Sriram Subramaniam. Structure of trimeric HIV-1 envelope glycoproteins. *Proceedings of the National Academy of Sciences*, 110(45):E4172–E4174, 2013.

[51] Martin Turk and Wolfgang Baumeister. The promise and the challenges of cryo-electron tomography. *FEBS letters*, 594(20):3243–3261, 2020.

[52] Aad W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.

[53] Marin van Heel. Finding trimeric HIV-1 envelope glycoproteins in random noise. *Proceedings of the National Academy of Sciences*, 110(45):E4175–E4177, 2013.

[54] Thorsten Wagner, Felipe Merino, Markus Stabrin, Toshio Moriya, Claudia Antoni, Amir Apelbaum, Philine Hagel, Oleg Sitsel, Tobias Raisch, Daniel Prumbaum, et al. SPHIRE-crYOLO is a fast and accurate fully automated particle picker for cryo-EM. *Communications biology*, 2(1):218, 2019.

[55] Abraham Wald. Note on the consistency of the maximum likelihood estimate. *The Annals of Mathematical Statistics*, 20(4):595–601, 1949.

[56] Abigail JI Watson and Alberto Bartesaghi. Advances in cryo-ET data processing: meeting the demands of visual proteomics. *Current Opinion in Structural Biology*, 87:102861, 2024.

[57] Wilson Wong, Xiao-chen Bai, Alan Brown, Israel S Fernandez, Eric Hanssen, Melanie Condron, Yan Hong Tan, Jake Baum, and Sjors HW Scheres. Cryo-EM structure of the Plasmodium falciparum 80S ribosome bound to the anti-protozoan drug emetine. *Elife*, 3:e03080, 2014.

[58] Ka Man Yip, Niels Fischer, Elham Paknia, Ashwin Chari, and Holger Stark. Atomic-resolution protein structure determination by cryo-em. *Nature*, 587(7832):157–161, 2020.

[59] Peijun Zhang. Advances in cryo-electron tomography and subtomogram averaging and classification. *Current opinion in structural biology*, 58:249–258, 2019.

[60] Ellen D Zhong, Tristan Bepler, Bonnie Berger, and Joseph H Davis. CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks. *Nature methods*, 18(2):176–185, 2021.

Appendix

Appendix organization. In Appendix A, we detail the empirical setup and procedures used for the simulations presented in Section 3. Appendix B provides a thorough description of the theoretical model underlying our analytical study of confirmation bias. In Appendix C, we state the main formal theoretical results of this work. Finally, Appendix D contains the detailed proofs of all the theorems.

Notations. We denote by $\mathsf{SO}(3)$ the special orthogonal group representing 3D rotations. The notation $[L] = \{0, 1, \ldots, L-1\}$ is used to represent discrete index set. All norms $\|\cdot\|$ are Frobenius norms. $\mathbb{1}$ represents the indicator function.

A Setup of the numerical experiments

In the theoretical model presented in Section 2, candidate particles $\{y_i\}_{i=0}^{N-1}$ were assumed to be drawn i.i.d. from an underlying distribution. However, in practical particle picking from micrographs, candidates are not sampled i.i.d.; instead, they are extracted using a sliding-window approach that introduces spatial overlap between neighboring regions and induces statistical dependencies. Furthermore, when selecting particles, it is desirable to avoid redundancy by ensuring that the selected patches do not overlap. To achieve this, the particle detection procedure must enforce a non-overlapping constraint: only one candidate patch should be selected from any set of overlapping regions. The algorithm used to implement this constraint is described in detail in Algorithm 2.

This algorithm identifies candidate particles in a collection of 2D micrographs $\{Y_i\}_{i=0}^{R-1} \subset \mathbb{R}^{k\times k}$ by computing their cross-correlation with a set of template images $\{x_\ell\}_{\ell=0}^{L-1} \subset \mathbb{R}^{d\times d}$. For each micrograph-template pair, a 2D circular cross-correlation (assuming periodic boundary conditions) is computed to produce a heatmap $C_{\ell,i}$, where each pixel indicates the similarity between a local patch of the micrograph and the template. Each $C_{\ell,i}$ has the same spatial resolution as the micrograph. To consolidate the response across templates, the algorithm computes a combined heat map C_i by taking the maximum pixel over all L templates. This results in a score map where each entry reflects the strongest template match at that location. The pixels in C_i are then flattened and sorted in descending order by score. The algorithm then iterates through the ranked pixels and selects particle candidates that (i) exceed a predefined threshold T, and (ii) do not spatially overlap with any previously selected patch. A binary mask is maintained to track used regions and enforce non-overlap. Only the highest-scoring, non-overlapping patches are added to the output set $\mathcal{A}$. This non-overlapping template matching strategy is commonly used in cryo-EM; see, for example, [19].

B Statistical models

B.1 Preliminaries

The particle picker's role is to classify each patch as either a particle or noise assuming no overlap between patches. Each micrograph patch, or equivalently, candidate particle is

Algorithm 2 `Template-Matching Particle Picker with Overlap`

Input: Templates $\{x_\ell\}_{\ell=0}^{L-1} \subset \mathbb{R}^{d\times d}$, Micrographs $\{Y_i\}_{i=0}^{R-1} \subset \mathbb{R}^{k\times k}$, Threshold $T \in \mathbb{R}$
Output: Set $\mathcal{A}$ of extracted particle patches

$\mathcal{A} \leftarrow \emptyset$
for $i = 0$ to $R-1$ **do** ▷ Loop over micrographs
  $C_i \leftarrow$ zero matrix in $\mathbb{R}^{k\times k}$ ▷ Heatmap storing the maximum cross-correlation
  for $\ell = 0$ to $L-1$ **do** ▷ Loop over templates
    $C_{\ell,i} \leftarrow x_\ell \star Y_i$ ▷ 2D circular cross-correlation
    for all pixels (u, v) **do**
      $C_i[u, v] \leftarrow \max(C_i[u, v], C_{\ell,i}[u, v])$ ▷ Maximum of the heatmap
    end for
  end for
  $(S_i, I_i) \leftarrow$ SortDescend(Flatten(C_i)) ▷ Sort the heatmap
  $M_i \leftarrow$ zero mask in $\{0, 1\}^{k\times k}$ ▷ Binary map tracking already selected regions
  for $m = 0$ to $k^2 - 1$ **do** ▷ Loop over pixels
    $(u, v) \leftarrow$ IndexToCoords($I_i[m], k$) ▷ Convert flat index to 2D coordinates
    if $S_i[m] > T$ **and** patch around (u, v) of size $d \times d$ does not overlap in M_i **then**
      $P \leftarrow Y_i[u - d/2 : u + d/2,\ v - d/2 : v + d/2]$
      $\mathcal{A} \leftarrow \mathcal{A} \cup \{P\}$
      Fill with ones the region of size $d \times d$ around (u, v) in M_i
    end if
  end for
end for
return $\mathcal{A}$

represented as $\{y_i\}_{i=0}^{N-1} \in \mathbb{R}^d$, where N is the number of candidate particles, and d is their dimension (which may correspond to flattened 2D or 3D image patches). Accordingly, the particle-picking task is defined by a binary function $f : \mathbb{R}^d \to \{0, 1\}$:

$$f(y_i) = \begin{cases} 1, & \text{if } y_i \text{ is a particle,} \\ 0, & \text{otherwise.} \end{cases} \tag{B.1}$$

To analyze confirmation bias in particle picking, we model the candidate particles $\{y_i\}_{i=0}^{N-1}$ as pure noise: $\{y_i\}_{i=0}^{N-1} \sim \mathcal{N}(0, \sigma^2 I_{d\times d})$, where $\mathcal{N}(0, \sigma^2 I_{d\times d})$ denotes multivariate Gaussian distribution with mean 0, and diagonal covariance matrix with variance σ^2. We consider a particle picker based on *template matching*, as outlined in Algorithm 1. This picker operates by cross-correlating each observation $\{y_i\}_{i=0}^{N-1}$ with L predefined templates, denoted $\{x_\ell\}_{\ell=0}^{L-1}$. Templates are assumed to be normalized so that $\|x_\ell\|_2 = 1$ for all $\ell \in \{0, \ldots, L-1\}$. If the correlation with at least one template exceeds a threshold T, the observation is classified as a particle and added to the set of selected particles, $\mathcal{A}$. We denote by $M = |\mathcal{A}|$ the number of selected particles.

To examine the confirmation bias introduced during the particle selection stage, it is necessary to establish a metric to compare the extracted particles with the initial templates used in the template selection process. Specifically, we analyze two approaches applied to the

particles extracted during this stage: one is based on the 2D classification output (relevant for cryo-EM) and the other is based on the 3D reconstruction process (relevant for subtomogram averaging in cryo-ET). In Section B.2, we focus on the 2D classification output, modeled using Gaussian mixture models (GMM). We analyze the relationship between the maximum-likelihood estimates of the GMM means and the predefined templates employed during particle picking. In Section B.3, we turn to cryo-ET and investigate the connection between the 3D template volume used in the particle-picking stage and the maximum-likelihood estimator derived from the 3D reconstruction process output. This analysis aims to uncover how the initial templates shape the reconstruction results and whether they contribute to biases in the reconstruction workflow.

B.2 2D classification of the particles extracted in cryo-EM

In cryo-EM, following the particle-picking stage, the extracted particles are typically subjected to a 2D classification process. This step groups particles into clusters based on their similarity, with each cluster represented by a centroid, which is the average of the observations assigned to that cluster. These centroids serve two primary purposes: they help assess the overall quality of the dataset and provide a basis for iteratively sorting out false positives or poorly aligned particles. In practice, the appearance of particle-like features in the centroids is often interpreted by researchers as evidence that the selected particles contain meaningful structural information.

B.2.1 2D Classification based on Gaussian Mixture Models

During the 2D classification process, particles are often modeled using a GMM, a widely used and effective probabilistic framework for describing particle heterogeneity. While alternative models exist, we focus on the GMM due to its analytical tractability. In this framework, the observed data are assumed to be generated from a mixture of Gaussian distributions, with each component representing a distinct subpopulation within the dataset. Formally, the realizations of the extracted particles are modeled as:

$$z_0, z_1, \ldots, z_{M-1} \overset{\text{i.i.d.}}{\sim} \sum_{\ell=0}^{L-1} w_\ell \cdot \mathcal{N}(\mu_\ell, \Sigma_\ell), \tag{B.2}$$

where $z_0, z_1, \ldots, z_{M-1}$ are the realizations of the *postulated* GMM distribution of the extracted particles, w_ℓ, μ_ℓ, and Σ_ℓ are the weight, mean, and covariance matrix of the ℓ-th Gaussian component, respectively. The weights $\{w_\ell\}_{\ell=0}^{L-1}$ satisfy $w_\ell \geq 0$, and $\sum_{\ell=0}^{L-1} w_\ell = 1$, and are assumed to be known. For the purpose of theoretical modeling, we assume that the number of templates used in the particle-picking process is equal to the number of classes in the GMM, denoted by L. In Section 3, we provide empirical results that demonstrate the impact of relaxing this assumption. This assumption is formally stated below as Assumption B.1.

Assumption B.1. *The number of templates used in the particle extraction algorithm (Algorithm 1) and the number of classes assumed in the GMM during the 2D classification process* (B.2) *are equal to* L.

Remark B.2. *Throughout the Appendix, we use the notation $\{y_i\}_{i=0}^{N-1}$ to denote the candidate particles provided as input to the template-matching particle picker. The resulting set of selected particles, which serves as the input for downstream tasks such as 2D classification and 3D reconstruction, is denoted by $\{z_i\}_{i=0}^{M-1}$.*

B.2.2 The log-likelihood function of GMM

In cryo-EM, it is common to assume that all components of the GMM share the same covariance matrix. Furthermore, we assume that this covariance is diagonal and identical across all L components, i.e., $\Sigma_\ell = \sigma^2_{\mathsf{GMM}} I_{d\times d}$ for all $\ell \in [L]$. The objective of the 2D classification process is to estimate the maximum-likelihood centroids $\{\mu_\ell\}_{\ell=0}^{L-1}$ in (B.2). In this context, the likelihood of the GMM specified in (B.2) is given by:

$$\begin{aligned}\mathcal{L}\left(\{\mu_\ell\}_{\ell\in[L]}\,;\mathcal{A}=\{z_i\}_{i\in[M]}\right) &= \prod_{i=0}^{M-1} f\left(z_i \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) \\ &= \prod_{i=0}^{M-1}\sum_{\ell=0}^{L-1} w_\ell \cdot \mathcal{N}\left(z_i \mid \mu_\ell, \Sigma_\ell = \sigma^2_{\mathsf{GMM}} I_{d\times d}\right), \end{aligned} \tag{B.3}$$

where we have defined,

$$f\left(z \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) \triangleq \sum_{\ell=0}^{L-1} w_\ell f_\ell(z) = \sum_{\ell=0}^{L-1} w_\ell \cdot \mathcal{N}\left(z \mid \mu_\ell, \Sigma_\ell = \sigma^2_{\mathsf{GMM}} I_{d\times d}\right). \tag{B.4}$$

$\mathcal{N}(z_i|\mu_\ell, \Sigma_\ell)$ represents the multivariate Gaussian density for the ℓ-th component, and the weights w_ℓ are the unknown mixture probabilities for each component. The objective of the 2D classification process is to estimate the centroids $\{\mu_\ell\}_{\ell=0}^{L-1}$, using the maximum likelihood estimator that maximizes the likelihood (B.3). The maximum-likelihood centroids, denoted by $\hat{\mu}_\ell$, are formally defined as

$$\{\hat{\mu}_\ell\}_{\ell=0}^{L-1} = \underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\max}\, \mathcal{L}(\{\mu_\ell\}_{\ell=0}^{L-1}\,;\mathcal{A}). \tag{B.5}$$

Thus, the aim of the 2D classification is to estimate the centroids of the GMM, defined in (B.2) using the extracted unlabeled particles $\mathcal{A} = \{z_i\}_{i=0}^{M-1}$ (i.e., using the output particles of Algorithm 1). The main question we address in this subsection is how the estimated centroids $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ relate to the templates $\{x_\ell\}_{\ell=0}^{L-1}$ used in the particle-picker template-matching (Algorithm 1), when the input micrograph is pure noise.

Remark B.3. *We use the notation σ^2_{GMM} to denote the variance of the postulated GMM model described in (B.2), and σ^2 to denote the variance of the noise in the input to the particle picker, as described in Algorithm 1. These two variances are not necessarily equal in our analysis.*

B.2.3 The extracted particles form a non-Gaussian mixture model

Although the extracted particles are often assumed to follow a Gaussian distribution, as defined in (B.2), the statistical properties of the particle picking process described in Algorithm 1 deviate from this assumption. Formally, the probability density function of the particles extracted by Algorithm 1 forms a mixture model that can be expressed as follows:

$$g(z) = \sum_{\ell=0}^{L-1} \pi_\ell \cdot g_\ell(z), \tag{B.6}$$

where L is the number of components (i.e., the number of templates), π_ℓ are the mixing coefficients satisfying $\sum_{\ell=0}^{L-1} \pi_\ell = 1$ and $\pi_\ell \geq 0$. Here, $\{g_\ell\}_{\ell=0}^{L-1} : \mathbb{R}^d \to \mathbb{R}$ are the probability density functions of the individual template-matching components, with $z \in \mathbb{R}^d$, given by:

$$g_\ell(z) = C_\ell(T) \cdot \exp\left(-\frac{\|z\|_F^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, x_\ell \rangle \geq T}. \tag{B.7}$$

where $C_\ell(T) \triangleq 1/\mathbb{P}\left[\langle z, x_\ell \rangle \geq T\right]$ is a normalization factor satisfying $\int_{z \in \mathbb{R}^d} dz \; g_\ell(z) = 1$. In other words, g_ℓ represents the distribution from which a particle associated with the ℓ-th template is selected, meaning its correlation with the template x_ℓ exceeds the threshold T. Clearly, $g_\ell(z)$ does not follow a Gaussian distribution as it is strictly positive only for $z \in \mathbb{R}^d$ values which satisfy $\langle z, x_\ell \rangle \geq T$. As shown in Proposition C.1, the normalization constants are identical across all templates, i.e., $C_{\ell_1}(T) = C_{\ell_2}(T)$ for any $\ell_1, \ell_2 \in [L]$, so we simplify the notation by writing $C_T = C_\ell(T)$. Figure 7 illustrates the probability density function of this truncated Gaussian distribution in one and two dimensions.

Misspecification of true statistics and the GMM model. It is important to recognize that the statistical assumptions underlying the particles, specifically, the GMM statistics defined in (B.2), do not accurately reflect the true statistics of the particle picking process based on template matching, as described in (B.6) and (B.7). Consequently, *the models are misspecified*, and this must be considered when estimating the centroids.

B.2.4 Relationship between 2D classification centers and templates

The model presented here raises two central questions:

1. What is the relationship between the mean of each component g_ℓ in the true mixture model (B.7) and the corresponding template x_ℓ when the number of particles $M \to \infty$? (answered in Proposition C.1)

2. In the 2D classification process based on a GMM model, how do the templates $\{x_\ell\}_{\ell=0}^{L-1}$ relate to the GMM maximum likelihood estimators of the means $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$, as specified in (B.5)? (answered in Theorem C.2)

Proposition C.1 examines the relationship between the mean of a single component g_ℓ in the mixture model (B.7) and the corresponding template x_ℓ. In other words, it relates

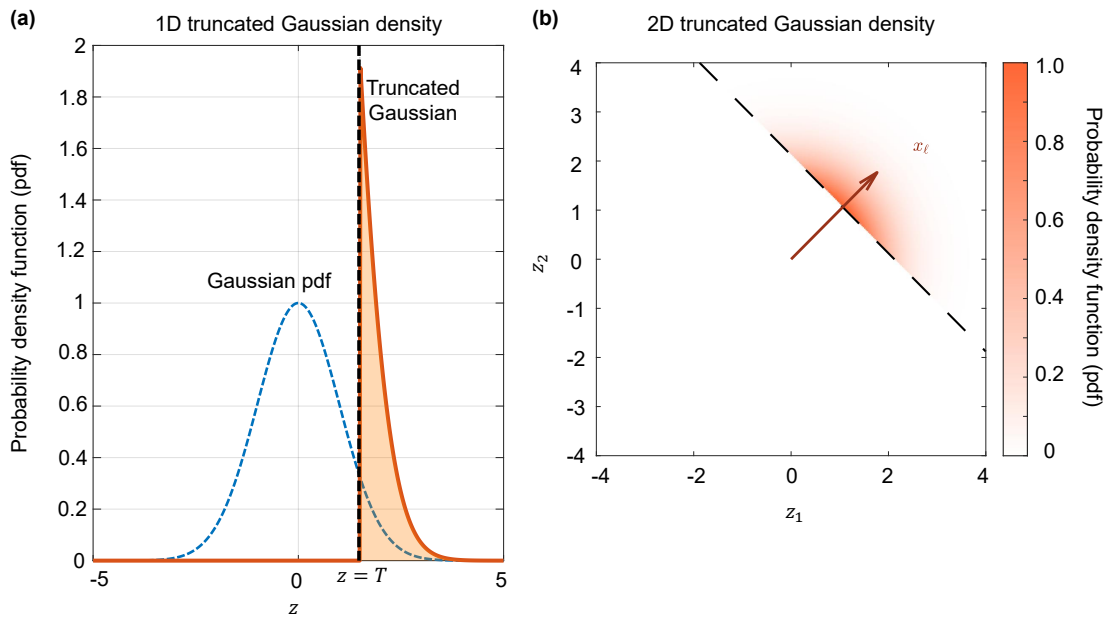

Figure 7: Template-matching-induced truncation of a Gaussian distribution. (a) In one dimension, the original Gaussian density (dashed) is truncated by a hard threshold T, resulting in a re-normalized density $g_\ell(z)$ supported only on the region where $z \geq T$. **(b)** In two dimensions, the original isotropic Gaussian is truncated to the half-space defined by $\langle z, x_\ell \rangle \geq T$, where x_ℓ is the template direction (shown as a red arrow) and the dashed line denotes the threshold boundary.

between the mean of all extracted particles whose correlation with the ℓ-th template x_ℓ exceeds the threshold T, and the template x_ℓ itself. We show that, in this case, as the number of particles increases ($M \to \infty$), the mean of these particles converges to the template x_ℓ, scaled by a factor $\alpha \geq T$, which is bounded below by the threshold used in the template-matching process. Notably, this convergence holds without imposing any specific constraints on the threshold T in Algorithm 1. Furthermore, in the second part of the proposition, we show that for asymptotically large thresholds T, the ratio $\alpha/T \to 1$, indicating that α approaches the lower bound.

The second result, Theorem C.2, demonstrates that the output of the 2D classification process reproduces the templates used during the particle-picking stage under two conditions: a large number of particles ($M \to \infty$) and a sufficiently large template-matching threshold T. The necessity of a high threshold is illustrated in Figure 3, and discussed in the main text.

B.3 3D reconstruction of the particles extracted in cryo-ET

Subtomogram averaging in cryo-ET. The cryo-ET model assumes a 3D tomogram composed of 3D subtomograms located at unknown locations, where each subtomogram is a noisy, randomly rotated version of the same underlying volume V. Specifically, each subtomogram z_i, is modeled as:

$$z_i = R_i \cdot V + \epsilon_i, \tag{B.8}$$

where V is the underlying 3D volume, $\{R_i\}_i \in \mathsf{SO}(3)$ are random 3D rotations, and $\{\epsilon_i\}_{i=0}^{M-1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_V^2 I_{d\times d})$ represents additive centered Gaussian noise with variance σ_V^2. In the cryo-ET computational pipeline, subtomograms are first detected and extracted from the tomogram during the particle-picking stage, followed by a 3D reconstruction process used to estimate the volume V from the extracted subtomograms $\{z_i\}_{i=0}^{M-1}$. Below, we describe the typical particle-picking and 3D reconstruction processes, along with the assumptions regarding their relationship.

Remark B.4. *We use the notation σ_V^2 to denote the variance of the postulated model described in (B.8), and σ^2 to denote the variance of the noise in the input to the particle picker, as described in Algorithm 1. These two variances are not necessarily equal in our analysis.*

Templates generation in cryo-ET. During the particle-picking stage in cryo-ET, L samples from the group of 3D rotations $\mathsf{SO}(3)$ are used to generate the templates from an underlying 3D volume template V_{template}, i.e., $x_\ell = R_\ell \cdot V_{\mathsf{template}}$, for $\ell \in [L]$. These templates are then used as inputs to Algorithm 1. The statistics of the output set $\mathcal{A}$ of the template-matching particle picker algorithm (Algorithm 1), when applied to pure Gaussian noise tomograms, follow the distribution

$$g_{\mathsf{picker}}(z) \triangleq \sum_{\ell=0}^{L-1} \pi(R_\ell) \cdot g_\ell(z) = C_T \sum_{\ell=0}^{L-1} \pi(R_\ell) \cdot \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, R_\ell \cdot V_{\mathsf{template}}\rangle \geq T}, \tag{B.9}$$

where $\pi(R_\ell)$ is the actual prior distribution of the samples $\{z_i\}_{i=0}^{M-1}$ for the L different rotations $\{R_\ell\}_{\ell=0}^{L-1}$, $g_\ell(z)$ is defined in (B.7), and $C_T = 1/\mathbb{P}\left[\langle z, R_\ell \cdot V_{\mathsf{template}}\rangle \geq T\right]$ is a normalization factor.

The 3D reconstruction process. The aim of the 3D reconstruction process is to estimate the volume V from the observations (subtomograms) $\{z_i\}_{i\in[M]}$ specified in (B.8). Below, we derive the explicit likelihood function.

Let $\mathcal{A} = \{z_i\}_{i\in[M]}$ be subtomograms follows the model in (B.8). Let V denote the underlying volume given in the process defined in (B.8). Let $f_{\mathsf{ET}}(z; V)$ represent the probability distribution function of the model in (B.8):

$$f_{\mathsf{ET}}(z; V) = \frac{1}{(2\pi\sigma_V^2)^{d/2}} \int_{R\in\mathsf{SO}(3)} p(R) \cdot \exp\left(-\frac{\|z - (R\cdot V)\|^2}{2\sigma_V^2}\right) dR, \tag{B.10}$$

where $p(R)$ represents the prior probability distribution of the rotations $R \in \mathsf{SO}(3)$. Then, the likelihood function of the model (B.8) is given by:

$$\begin{aligned}\mathcal{L}_{\mathsf{postulated}}\left(V; \mathcal{A} = \{z_i\}_{i\in[M]}\right) &= \prod_{i=0}^{M-1} f_{\mathsf{ET}}(z_i; V) \\ &= \frac{1}{(2\pi\sigma_V^2)^{dM/2}} \prod_{i=0}^{M-1} \int_{R\in\mathsf{SO}(3)} p(R) \cdot \exp\left(-\frac{\|z_i - (R\cdot V)\|^2}{2\sigma_V^2}\right) dR. \end{aligned} \tag{B.11}$$

The 3D reconstruction process typically involves iterative methods (such as expectation-maximization) to estimate the volume V from the 3D extracted particles (B.8). During these iterative methods, a discretization of $\mathsf{SO}(3)$ is performed, which we denote by $\{R_n\}_n \in \mathsf{SO}(3)$. Then, we approximate the likelihood function of (B.11),

$$\bar{\mathcal{L}}_{\mathsf{postulated}}\left(V; \mathcal{A} = \{z_i\}_{i=0}^{M-1}\right) = \frac{1}{(2\pi\sigma_V^2)^{dM/2}} \prod_{i=0}^{M-1} \sum_{R_n \in \mathsf{SO}(3)} w(R_n) \cdot \exp\left(-\frac{\|z_i - (R_n \cdot V)\|^2}{2\sigma_V^2}\right). \tag{B.12}$$

In addition, we represent the corresponding probability density function associated with this discretization, following from (B.10) by,

$$\bar{f}_{\mathsf{ET}}(z; V) = \frac{1}{(2\pi\sigma_V^2)^{d/2}} \sum_{R_n \in \mathsf{SO}(3)} w(R_n) \cdot \exp\left(-\frac{\|z - (R_n \cdot V)\|^2}{2\sigma_V^2}\right). \tag{B.13}$$

The maximum likelihood estimator in this case is given by:

$$\hat{V} = \arg\max_V \bar{\mathcal{L}}_{\mathsf{postulated}}(V; \mathcal{A}). \tag{B.14}$$

It is important to note that the *postulated* statistics of the subtomograms, as specified in (B.10), differ from the *actual* statistics of the output from the particle-picking stage, as given in (B.9). Therefore, the model is misspecified, similar to the GMM case discussed in the previous section.

The following assumption states that the same set of SO(3) samples is used both in the 3D reconstruction process and to generate the templates in the particle-picking stage. Specifically, the rotations $\{R_\ell\}_{\ell \in \{0,\ldots,L-1\}}$ used during particle picking, i.e., $x_\ell = R_\ell \cdot V_{\mathsf{template}}$, are identical to the discretization $\{R_n\}_n$ of SO(3) employed in the reconstruction procedure, as defined in (B.13).

Assumption B.5. *The* SO(3) *samples used for generating the templates in the particle-picking stage are the same as the samples employed during the 3D reconstruction process.*

The primary question we address concerns the relationship between V_{template}, the 3D template volume used during the particle-picking stage, and $\hat{V}$, the output of the 3D reconstruction process, represented as the maximum likelihood estimator $\hat{V}$ defined in (B.14). At first glance, one might expect the reconstructed volume to converge to zero since the underlying candidate particles are zero-mean Gaussian noise. However, Theorem C.3 reveals a surprising result: as the number of extracted particles $M \to \infty$ and for asymptotic threshold $T \to \infty$, the reconstructed 3D volume $\hat{V}$ converges to the template 3D volume V_{template}, up to a scaling factor.

C Theoretical results

C.1 Particle picking based on a single template matching

We begin by presenting a theoretical result that establishes the relationship between the mean of the ℓ-th component of the mixture model, as defined in (B.7), and the corresponding

template x_ℓ.

The centers of the template matching process. Using the definitions in Algorithm 1, let $m_\ell^{(T)}$ be defined as:

$$m_\ell^{(T)} = \frac{1}{|\mathcal{A}_\ell^{(T)}|} \sum_{z_i \in \mathcal{A}_\ell^{(T)}} z_i, \tag{C.1}$$

where

$$\mathcal{A}_\ell^{(T)} = \{z_i \in \mathcal{A} \mid \langle z_i, x_\ell \rangle \geq T\}. \tag{C.2}$$

Here, $m_\ell^{(T)}$ represents the centroid of the particles whose correlation with the template x_ℓ exceeds the threshold T. The key question we address is: what is the relationship between $m_\ell^{(T)}$ and the corresponding template x_ℓ?

To this end, it would be beneficial to define, for every $\ell \in [L]$, the following sets:

$$\mathcal{V}_\ell^{(T)} \triangleq \left\{ v \in \mathbb{R}^d : \langle v, x_\ell \rangle \geq T \right\}. \tag{C.3}$$

Note that $\left\{ \mathcal{V}_\ell^{(T)} \right\}_{\ell \in [L]}$ are deterministic, and are not necessarily disjoint sets.

Proposition C.1. *Let $\{y_i\}_{i=0}^{N-1} \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2 I_{d\times d})$ represent the input candidate particles to Algorithm 1. Let $\{x_\ell\}_{\ell=0}^{L-1}$ denote the normalized templates, T the threshold, and $\mathcal{A}$ the output set of selected particles in Algorithm 1. Define $m_\ell^{(T)}$ as given in (C.1), and $\mathcal{A}_\ell^{(T)}$ as the corresponding set, specified in (C.2). Then,*

1. *For $N \to \infty$,*

$$m_\ell^{(T)} = \frac{1}{|\mathcal{A}_\ell^{(T)}|} \sum_{y_i \in \mathcal{A}_\ell^{(T)}} y_i \xrightarrow{\text{a.s.}} \alpha \cdot x_\ell, \tag{C.4}$$

   for a constant $\alpha \geq T$, independent of ℓ.

2. *For $N \to \infty$, and $T \to \infty$,*

$$\lim_{T\to\infty} \lim_{N\to\infty} \frac{m_\ell^{(T)}}{T} = x_\ell, \tag{C.5}$$

   where the convergence is almost surely, and for every $\ell \in [L]$.

This result proves that averaging the noise observations $\{y_i\}_{i=0}^{N-1}$, whose correlation with the template x_ℓ exceeds a specified threshold, converges almost surely to the template x_ℓ, scaled by a factor α. The scaling factor α is guaranteed to be at least as large as the threshold used in the template-matching process. Furthermore, the second part of the result shows that for asymptotically large thresholds T, the ratio α/T approaches 1, meaning α attains the lower bound asymptotically in T. In essence, Proposition C.1 establishes that the centers of the GMM defined in (B.2) resemble the templates, up to a scaling factor. This differs from the "Einstein from noise" phenomenon and the associated confirmation bias in 2D classifications, where the resulting estimates are positively correlated with the initial templates but do not fully replicate them [3].

C.2 Particle picking based on multiple templates

In the previous sections, we demonstrated that the particle-picking process generates a mixture model with L components, where L corresponds to the number of templates. The centers of these mixtures reproduce the templates, up to a scaling factor. However, this result does not guarantee that the output of the 2D classification process (based on the maximum likelihood estimation of the GMM defined in (B.5)) will directly yield these centroids. The primary goal of this section is to analyze the relationship between the templates $\{x_\ell\}_{\ell=0}^{L-1}$ and the 2D classification outputs $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$, as defined in (B.5).

In general, the sequence of maximum likelihood estimators converges in probability to the true value being estimated, provided that the assumed model distribution aligns with the underlying distribution from which the samples are generated [35]. Therefore, if the underlying distribution of the extracted particles were GMM, the maximum-likelihood centroid estimators $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$, as defined in (B.5), would converge to the centroids $\{m_\ell\}_{\ell=0}^{L-1}$, as defined in (C.1), assuming the 2D classification process follows a GMM.

However, as discussed in Section B, in our case, while the assumed model is a GMM, the actual statistics of the particle-picking output do not follow Gaussian statistics. Specifically, the distribution of the particle-picking output is described by (B.6) and (B.7). Therefore, additional conditions are required for the maximum-likelihood centroid estimators of the assumed GMM model $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$, to converge to the centroids of the true mixture model, $\{m_\ell\}_{\ell=0}^{L-1}$. These conditions are outlined in Theorem C.2.

Theorem C.2 (Convergence of 2D class averages to particle-picking templates)**.** *Fix $L \geq 2$, and denote the input templates $\{x_\ell\}_{\ell=0}^{L-1}$, input threshold T, and input observations $\{y_i\}_{i=0}^{N-1} \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2 I_{d\times d})$ to Algorithm 1. Let $\mathcal{A}$ be the output of Algorithm 1, such that $|\mathcal{A}| = M$. Assume that $x_{\ell_1} \neq x_{\ell_2}$ for every $\ell_1 \neq \ell_2$. Let $\{\hat{\mu}_\ell\}_{l=0}^{L-1}$ be the maximum-likelihood estimators of the 2D classification process, applied on the extracted particles set $\mathcal{A}$, as defined in (B.5). Then, under Assumption B.1, there exists a permutation function $\pi : [L] \to [L]$, such that for every $0 \leq \ell \leq L-1$,*

$$\lim_{T\to\infty} \lim_{N\to\infty} \frac{\hat{\mu}_{\pi(\ell)}}{T} = x_\ell, \tag{C.6}$$

where the convergence is in probability.

Theorem C.2 demonstrates that, for an asymptotic number of extracted particles M and an asymptotic template-matching threshold T, the output of the 2D classification process reproduces the templates used during the particle-picking stage. This reproduction occurs up to a permutation of the output centers, meaning the output centroids are aligned with the templates, but their sequencing order may differ. The requirement for a sufficiently large threshold is crucial, as shown in Figure 3.

The relation between Proposition C.1 and Theorem C.2 is shown in Figure 8, which illustrates the difference between the labeled setting (Proposition C.1), where particle-template associations are known and mean vectors converge directly to scaled templates, and the unlabeled setting (Theorem C.2), where convergence to templates occurs for a high threshold.

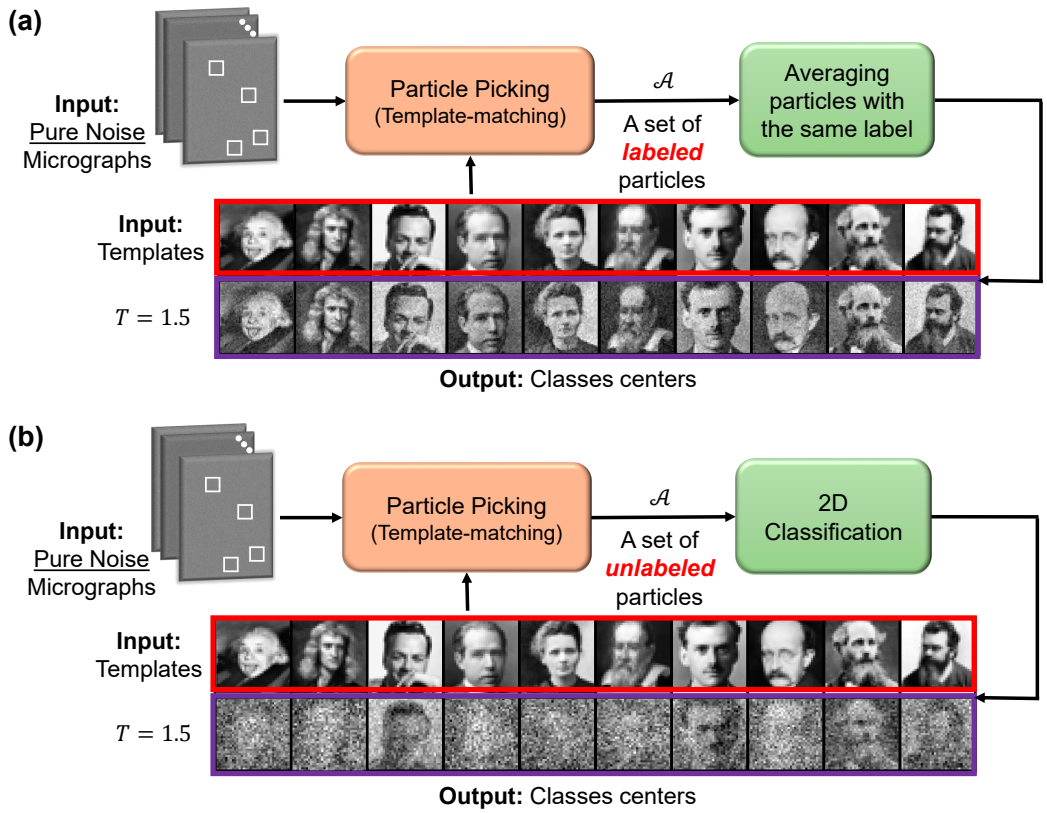

Figure 8: Illustration of the relationship between the means for labeled particles (Proposition C.1) and unlabeled particles (Theorem C.2). Starting from pure-noise observations, template-based particle picking selects highly correlated candidates. **(a)** When the particles are labeled, that is, each selected particle is assigned to a known template, Proposition C.1 characterizes the asymptotic behavior of the mean of particles within each label. These means converge to the corresponding templates up to a scaling factor, and this convergence holds for any threshold T. **(b)** When the particles are unlabeled, i.e., the problem involves latent variables, the 2D classification step treats the data as a GMM. Theorem C.2 shows that as the threshold increases, the estimated class centers increasingly align with the original templates (see Figure 3). However, at low thresholds, this convergence does not necessarily occur as can be seen in the example. This simulation uses $M = 10^5$ particles collected during the particle-picking process and image sizes of $d = 50 \times 50$. The data generation model for this experiment is the same as presented in Section 2.1.

C.3 Three-dimensional volume alignment

In the 3D reconstruction process, in contrast to the 2D classification process, it is assumed that the templates exhibit an algebraic structure governed by $\mathsf{SO}(3)$ rotations of the underlying volume. This raises the question: under what conditions does the maximum-likelihood estimator, as defined in (B.14), converge to the template volume used in the particle-picking procedure based on template matching? Similar to the 2D classification case discussed in the previous section, the conditions include an asymptotic number of extracted particles, an asymptotic threshold, and an additional assumption regarding the $\mathsf{SO}(3)$ samples, as outlined in Assumption B.5.

Theorem C.3. *Assume the input templates to Algorithm 1 are of the form* $\{x_\ell\}_{\ell=0}^{L-1} = \{R_\ell \cdot V_{\mathsf{template}}\}_{\ell=0}^{L-1}$*, where* V_{template} *is the input 3D template volume, and* $\{R_\ell\}_{\ell=0}^{L-1} \in \mathsf{SO}(3)$*. Assume that that* $R_{\ell_1} \cdot V_{\mathsf{template}} \neq R_{\ell_2} \cdot V_{\mathsf{template}}$ *for every* $\ell_1 \neq \ell_2$*. Let* $\hat{V}$ *be the maximum-likelihood 3D volume estimator, as defined in (B.14). Let* T *be the threshold of the template matching algorithm as specified in Algorithm 1. Then, under Assumption B.5,*

$$\lim_{T\to\infty} \lim_{N\to\infty} \frac{\hat{V}}{T} = R \cdot V_{\mathsf{template}}, \tag{C.7}$$

for some $R \in \mathsf{SO}(3)$*, where the convergence is in probability.*

C.4 The impact of the number of observations and signal dimension

In the previous sections, we analyzed the asymptotic behavior of the convergence of 2D classes in cryo-EM and the reconstructed 3D structure in cryo-ET relative to the underlying 3D template. This analysis was carried out under the asymptotic regimes of an infinite number of extracted particles ($M \to \infty$) and an infinite threshold ($T \to \infty$). We now shift focus to the convergence rates of maximum-likelihood estimators (MLEs), specifically examining how these rates depend on the signal dimension and the number of samples.

Sample complexity for estimating the mean of a Gaussian distribution. We begin with the sample complexity for estimating the mean of a d-dimensional multivariate Gaussian distribution. Consider i.i.d. observations $X_i \sim \mathcal{N}(\mu, \Sigma = \sigma^2 I_{d\times d})$ in $\mathbb{R}^d$, and let $\hat{\mu}_{\mathrm{ML}}$ denote the MLE of μ based on M samples $\{X_i\}_{i=0}^{M-1}$. By the asymptotic efficiency of MLEs [35], as $n \to \infty$,

$$\sqrt{M}\,(\hat{\mu}_{\mathrm{ML}} - \mu) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \mathcal{I}^{-1}), \tag{C.8}$$

where $\mathcal{I}$ is the Fisher information matrix. In our case, since $\Sigma = \sigma^2 I_{d\times d}$, the limiting distribution becomes

$$\sqrt{M}\,(\hat{\mu}_{\mathrm{ML}} - \mu) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma^2 I_{d\times d}). \tag{C.9}$$

Thus, for sufficiently large n, the mean squared error (MSE) between $\hat{\mu}_{\mathrm{ML}}$ and μ scales as

$$\|\hat{\mu}_{\mathrm{ML}} - \mu\|^2 \propto \frac{d\sigma^2}{M}. \tag{C.10}$$

Sample complexity for estimating the means in a Gaussian mixture model (GMM). Now consider a GMM with L components, where each component has mean μ_ℓ and covariance matrix $\Sigma_\ell = \sigma^2 I_{d\times d}$. Let $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ be the MLEs of the component means. As in the single-Gaussian case, the efficiency result of (C.8) applies, but the convergence behavior is more intricate due to the latent structure.

In the general case, without assumptions on the separation between components, the MSE satisfies the following pessimistic bound [33]:

$$\sum_{\ell=0}^{L-1} \|\hat{\mu}_\ell - \mu_\ell\|^2 = \mathcal{O}\left(\frac{\sigma^{O(L)}}{M}\right). \tag{C.11}$$

However, if the component means are well-separated, specifically, if $\|\mu_{\ell_1} - \mu_{\ell_2}\| \gg d\sigma^2$ for all $\ell_1 \neq \ell_2$, then the MSE improves significantly, scaling as [15]

$$\sum_{\ell=0}^{L-1} \|\hat{\mu}_\ell - \mu_\ell\|^2 \propto Ld\frac{\sigma^2}{M}. \tag{C.12}$$

Sample complexity for the convergence of the particle-picking mixture to templates. The distribution of extracted particles, defined in (B.7), does not follow a standard GMM but rather a *truncated Gaussian mixture*, due to the selection bias induced by the particle-picking mechanism. For such non-Gaussian models, consistency and asymptotic efficiency of MLEs can still hold under suitable regularity conditions [52].

In particular, we state the following known result for truncated Gaussian distributions.

Lemma C.4 ([17])**.** *Let $X \sim \mathcal{N}(0, \sigma^2)$. Then, for any threshold $T \in \mathbb{R}$,*

$$\mathbb{E}[X \mid X \geq T] = \sigma \cdot \frac{\varphi(T/\sigma)}{Q(T/\sigma)}, \tag{C.13}$$

where $\varphi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ is the standard normal density function, $Q(z) = \int_z^\infty \varphi(t)\, dt = 1 - \Phi(z)$ is the standard Gaussian upper tail probability, and Φ is the CDF of the standard normal distribution. In addition, the variance of X conditioned on $X \geq T$ is given by:

$$\mathrm{Var}(X \mid X \geq T) = \sigma^2 \left(1 + \frac{T}{\sigma} \cdot \frac{\varphi(T/\sigma)}{Q(T/\sigma)} - \left(\frac{\varphi(T/\sigma)}{Q(T/\sigma)}\right)^2\right). \tag{C.14}$$

Asymptotically, as $T \to \infty$, we have, [3]

$$\mathbb{E}[X \mid X \geq T] \sim T + \frac{\sigma^2}{T}, \tag{C.15}$$

$$\mathrm{Var}(X \mid X \geq T) \sim \sigma^2 \left(1 - \frac{\sigma^2}{T^2}\right). \tag{C.16}$$

In our setting, the random variable X corresponds to the projection of the noisy observation onto the direction of the template, and Lemma C.4 characterizes the distribution of this projection after truncation by the particle-picking threshold (see Lemmas D.2 and D.3 for further details). We thus define the effective variance of each truncated component by

$$\sigma_T^2 = \sigma^2 \left(1 - \frac{\sigma^2}{T^2}\right). \tag{C.17}$$

[3] Here, the symbol $f(t) \sim g(t)$ denotes asymptotic equivalence, i.e., $\lim_{t\to\infty} \frac{f(t)}{g(t)} = 1$.

In the regime where the threshold $T \to \infty$, the extracted particle mixture becomes well-separated, assuming the templates are separated, since the the mean of each component concentrates near $m_\ell^{(T)} \approx T \cdot x_\ell$, and the variance tends toward $\sigma_T^2 \approx \sigma^2$. Thus, they satisfy the conditions under which the rate in (C.12) applies.

Let M denote the number of extracted particles selected by the thresholding mechanism. Then, the MSE of the estimated means from the extracted mixture satisfies:

$$\sum_{\ell=0}^{L-1} \|\hat{\mu}_\ell - m_\ell^{(T)}\|^2 \propto \frac{Ld\sigma_T^2}{M}, \tag{C.18}$$

where σ_T^2 is defined in (C.17), as $M, T \to \infty$. This scaling indicates that a smaller signal dimension d or a larger number of observations M leads to faster convergence of the estimated means to the true ones. In the context of cryo-EM or cryo-ET particle picking, this implies that smaller particle sizes and a greater number of extracted particles amplify the bias toward the underlying templates.

D Proofs

D.1 Proof of Proposition C.1

Let $\{y_i\}_{i=0}^{N-1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}_d, \sigma^2 I_{d\times d})$, be the input observations to Algorithm 1, and $\mathcal{A} = \{z_i\}_{i=0}^{M-1}$ be the output set of the extracted particles of Algorithm 1, and recall the definitions of $\mathcal{A}_\ell^{(T)}$ (C.1), and $m_\ell^{(T)}$ (C.2). In addition, recall the definition of $\mathcal{V}_\ell^{(T)} \subset \mathbb{R}^d$ (C.3), which defines the subset of $\mathbb{R}^d$, where the inner-product between the particles and the ℓ-th template x_ℓ exceed the threshold T.

Let $\mathcal{P} : \mathbb{R}^d \to \mathbb{R}$ denote the d-dimensional Gaussian density function:

$$\mathcal{P}(y) \triangleq \mathcal{N}\left(y; \mu = 0, \Sigma = \sigma^2 I_{d\times d}\right) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left\{-\frac{\|y\|^2}{2\sigma^2}\right\}. \tag{D.1}$$

The input to Algorithm 1 follows the distribution $\mathcal{P}$ (by the assumption on the distribution of the input observations $\{y_i\}_{i=0}^{N-1} \sim \mathcal{N}(\mathbf{0}_d, \Sigma = \sigma^2 I_{d\times d})$), while the output follows the mixture model defined in (B.6) and (B.7). In particular, the extracted particles in the set $\mathcal{A}_\ell^{(T)}$ follows the distribution $g_\ell(z)$, as defined in (B.7), i.e., $g_\ell(z) = C_T \cdot \mathcal{P}(z)\, \mathbb{1}_{\langle z, x_\ell\rangle \geq T}$, where C_T is a normalization factor.

By the Strong Law of Large Numbers (SLLN), for $N \to \infty$, we have

$$\begin{aligned} m_\ell^{(T)} &= \frac{1}{|\mathcal{A}_\ell^{(T)}|} \sum_{\langle z_i, x_\ell\rangle \geq T} z_i \\ &= \frac{\frac{1}{N}\sum_{\langle y_i, x_\ell\rangle \geq T} y_i}{|\mathcal{A}_\ell^{(T)}|/N} \xrightarrow{\text{a.s.}} \frac{\mathbb{E}\left[y_1 \mathbb{1}\left\{y_1 \in \mathcal{V}_\ell^{(T)}\right\}\right]}{\mathbb{P}\left[y_1 \in \mathcal{V}_\ell^{(T)}\right]}. \end{aligned} \tag{D.2}$$

The numerator term at the right-hand-side of (D.2) is simply the expected value of all vectors which their inner product with the template x_ℓ is greater than the threshold T. In addition, the denominator term $\mathbb{P}\left[y_1 \in \mathcal{V}_\ell^{(T)}\right] = \int_{y\in\mathcal{V}_\ell^{(T)}} \mathcal{P}(y)\,dy > 0$; thus, the right-hand-side of (D.2) is well-defined for every $T < \infty$. By the definition of g_ℓ (B.7), and following (D.2), we have,

$$m_\ell^{(T)} \xrightarrow{\text{a.s.}} \frac{\mathbb{E}\left[y_1 \mathbb{1}\left\{y_1 \in \mathcal{V}_\ell^{(T)}\right\}\right]}{\mathbb{P}\left[y_1 \in \mathcal{V}_\ell^{(T)}\right]} = \frac{\int_{y\in\mathcal{V}_\ell^{(T)}} y\cdot\mathcal{P}(y)\,dy}{\int_{y\in\mathcal{V}_\ell^{(T)}} \mathcal{P}(y)\,dy}$$
$$= C_T \int_{\mathbb{R}^d} y\cdot\mathcal{P}(y)\,\mathbb{1}_{\langle y,x_\ell\rangle\geq T}\,dy \triangleq \int_{\mathbb{R}^d} z\cdot g_\ell(z)\,dz = \mathbb{E}_{g_\ell}[z]. \tag{D.3}$$

where $C_T \triangleq 1/\mathbb{P}\left[y_1 \in \mathcal{V}_\ell^{(T)}\right]$.

Part (1). First, we prove that $m_\ell^{(T)} \xrightarrow{\text{a.s.}} \alpha_\ell \cdot x_\ell$, in the limit as $N \to \infty$, where α_ℓ is a constant. This condition can be equivalently expressed algebraically: for any vector $u \in \mathbb{R}^d$ orthogonal to x_ℓ (i.e., $\langle u, x_\ell\rangle = 0$), the mean $m_\ell^{(T)}$ must also be orthogonal to u (i.e., $\langle u, m_\ell^{(T)}\rangle = 0$). Consequently, proving this is equivalent to showing that for every $u \in \mathbb{R}^d$ orthogonal to the template x_ℓ ($\langle u, x_\ell\rangle = 0$), the mean $m_\ell^{(T)}$ satisfies the condition:

$$\langle m_\ell^{(T)}, u\rangle = \frac{1}{|\mathcal{A}_\ell^{(T)}|} \sum_{\langle z_i, x_\ell\rangle\geq T} \langle z_i, u\rangle \xrightarrow{\text{a.s.}} 0, \tag{D.4}$$

where $\mathcal{A}_\ell^{(T)}$ is the set of indices for which $\langle z_i, x_\ell\rangle \geq T$.

To prove the right-hand-side of (D.4), we substitute (D.3) into the left-hand-side of (D.4), and use the linearity of the inner-product to obtain,

$$\langle m_\ell^{(T)}, u\rangle \xrightarrow{\text{a.s.}} C_T \int_{\mathbb{R}^d} \langle z, u\rangle \mathcal{P}(z)\,\mathbb{1}_{\langle z,x_\ell\rangle\geq T}\,dz. \tag{D.5}$$

We prove in Lemma D.1, stated below and proved in Section D.1.1, that the right-hand size of (D.5) is 0 from symmetry properties, which proves that $m_\ell^{(T)} \xrightarrow{\text{a.s.}} \alpha_\ell \cdot x_\ell$.

Lemma D.1. *Let $\mathcal{P}(z)$ be the d-dimensional multivariate Gaussian distribution defined in (D.1). Assume $u \in \mathbb{R}^d$, satisfying $\langle u, x_\ell\rangle = 0$. Then, for every T,*

$$\int_{\mathbb{R}^d} \langle z, u\rangle \mathcal{P}(z)\,\mathbb{1}_{\langle z,x_\ell\rangle\geq T}\,dz = 0. \tag{D.6}$$

In Lemma D.2, stated below, and proved in Section D.1.2, we prove that $\alpha_\ell \geq T$, and that α_ℓ is independent of ℓ, i.e., $\alpha_{\ell_1} = \alpha_{\ell_2}$, for every $\ell_1, \ell_2 \in [L]$, which completes the proof of the first part of the proposition.

Lemma D.2. *Recall the definition of $m_\ell^{(T)}$ in (C.1). For every $\ell_1, \ell_2 \in [L]$, the following holds as $N \to \infty$,*

$$\alpha_{\ell_1} = \left\|\langle m_{\ell_1}^{(T)}, x_{\ell_1}\rangle\right\| = \left\|\langle m_{\ell_2}^{(T)}, x_{\ell_2}\rangle\right\| = \alpha_{\ell_2} \geq T. \tag{D.7}$$

Part (2). In this part, we prove that asymptotically in T, the factor α grows linearly with T, attaining the lower bound from the previous part. To prove this result, we use Lemma D.3, which is stated below and proved in Section D.1.3.

Lemma D.3. *Let $\mathbb{E}_{g_\ell}[z] = \int_{\mathbb{R}^d} z \cdot g_\ell(z)\, dz$, and let C_T be the normalization factor of g_ℓ, for every $T \in \mathbb{R}^+$. Then, for $T \to \infty$,*

$$\lim_{T\to\infty} \left\| \frac{\mathbb{E}_{g_\ell}[z]}{T} \right\| = \lim_{T\to\infty} \left\| \frac{1}{T} \int_{\mathbb{R}^d} C_T \cdot y \cdot \mathcal{P}(y)\, \mathbb{1}_{\langle y, x_\ell\rangle \geq T}\, dy \right\| = 1. \tag{D.8}$$

Then, combining Lemma D.1, Lemma D.2, and Lemma D.3, and the fact that $\|x_\ell\| = 1$, for every $\ell \in [L]$, results

$$\lim_{N,T\to\infty} \left\| \frac{\alpha_\ell^{(T)}}{T} \right\| = \lim_{N,T\to\infty} \left\| \frac{m_\ell^{(T)}}{T} \right\| \overset{\text{a.s.}}{=} \lim_{T\to\infty} \left\| \frac{\mathbb{E}_{g_\ell}[z]}{T} \right\| = 1. \tag{D.9}$$

where the first almost sure convergence follows from (D.3), and the second convergence follows by applying Lemma D.3, which completes the proof of the second part of the proposition.

D.1.1 Proof of Lemma D.1

Fix $u \in \mathbb{R}^d$, such that $\langle u, x_\ell\rangle = 0$. First, we show that for every $v \in \mathcal{V}_\ell^{(T)}$, there is a unique corresponding $v' \in \mathcal{V}_\ell^{(T)}$ such that $\langle u, v\rangle = -\langle u, v'\rangle$, and $\mathcal{P}(v) = \mathcal{P}(v')$. Every $v \in \mathbb{R}^d$ can be decomposed uniquely by $v = \gamma_v u + w$, where $w \in \mathbb{R}^d$ orthogonal to u, $\langle u, w\rangle = 0$. Define $v' = -\gamma_v u + w$. Clearly, the map $v \to v'$ is unique, and due to the orthogonality between u and w, we have, $\langle u, v\rangle = \gamma_v$, and $\langle u, v'\rangle = -\gamma_v$; thus, $\langle u, v\rangle = -\langle u, v'\rangle$. In addition, $v \in \mathcal{V}_\ell^{(T)}$ if and only if $v' \in \mathcal{V}_\ell^{(T)}$, since,

$$\begin{aligned} v \in \mathcal{V}_\ell^{(T)} &\iff \langle v, x_\ell\rangle \geq T \iff \langle \gamma_v u + w, x_\ell\rangle \geq T \\ &\iff \langle w, x_\ell\rangle \geq T \\ &\iff \langle -\gamma_v u + w, x_\ell\rangle \geq T \iff \langle v', x_\ell\rangle \geq T \iff v' \in \mathcal{V}_\ell^{(T)}, \end{aligned} \tag{D.10}$$

where the transition in (D.10) follows from the orthogonality between u and x_ℓ.

Next, we show that $\mathcal{P}(v) = \mathcal{P}(v')$. We note that $\mathcal{P}$ (D.1) is an isotropic function that depends solely on the norm of z. Thus, we prove below that $\|v\| = \|v'\|$, which would prove that $\mathcal{P}(v) = \mathcal{P}(v')$:

$$\|v\|^2 = \|\gamma_v u + w\|^2 = \gamma_v^2 \|u\|^2 + \|w\|^2 = \|-\gamma_v u + w\|^2 = \|v'\|^2, \tag{D.11}$$

where in the second equality we have used the orthogonality between u and w. Therefore, $\|v\| = \|v'\|$, which shows that $\mathcal{P}(v) = \mathcal{P}(v')$.

Next, note that the right-hand-side of (D.3) can be rewritten as follows,

$$\int_{\mathbb{R}^d} \langle z, u\rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell\rangle \geq T}\, dz = \int_{z \in \mathcal{V}_\ell^{(T)}} \langle z, u\rangle \mathcal{P}(z)\, dz$$

$$= \int_{z \in \mathcal{V}_\ell^{(T)} : \langle z,u \rangle > 0} \langle z, u \rangle \mathcal{P}(z) \; dz + \int_{z \in \mathcal{V}_\ell^{(T)} : \langle z,u \rangle < 0} \langle z, u \rangle \mathcal{P}(z) \; dz. \tag{D.12}$$

Thus, since for every $v \in \mathcal{V}_\ell^{(T)}$ there is a unique corresponding $v' \in \mathcal{V}_\ell^{(T)}$ such that $\langle v, u_i \rangle = -\langle v', u_i \rangle$ and $\mathcal{P}(v) = \mathcal{P}(v')$, it is clear that the r.h.s. of (D.12) is zero, which completes the proof.

D.1.2 Proof of Lemma D.2

First, we show that for every $\ell_1, \ell_2 \in [L]$, $\langle x_{\ell_1}, m_{\ell_1}^{(T)} \rangle = \langle x_{\ell_2}, m_{\ell_2}^{(T)} \rangle$, which proves that α is independent of ℓ. As the templates $\{x_\ell\}_{\ell=0}^{L-1}$ are normalized, there is a orthonormal matrix O_{ℓ_1,ℓ_2}, such that $O_{\ell_1,\ell_2} x_{\ell_1} = x_{\ell_2}$. Then, following from the right-hand-side of (D.3), we have,

$$\begin{aligned}
\langle x_{\ell_1}, m_{\ell_1}^{(T)} \rangle &= \int_{\mathbb{R}^d} \langle z, x_{\ell_1} \rangle \mathcal{P}(z) \, \mathbb{1}_{\langle z, x_{\ell_1} \rangle \geq T} \; dz && \text{(D.13)} \\
&= \int_{\mathbb{R}^d} \langle z, O_{\ell_2,\ell_1} x_{\ell_2} \rangle \mathcal{P}(z) \, \mathbb{1}_{\langle z, O_{\ell_2,\ell_1} x_{\ell_2} \rangle \geq T} \; dz && \text{(D.14)} \\
&= \int_{\mathbb{R}^d} \langle O_{\ell_2,\ell_1} z, O_{\ell_2,\ell_1} x_{\ell_2} \rangle \mathcal{P}(O_{\ell_2,\ell_1} z) \, \mathbb{1}_{\langle O_{\ell_2,\ell_1} z, O_{\ell_2,\ell_1} x_{\ell_2} \rangle \geq T} \; dz && \text{(D.15)} \\
&= \int_{\mathbb{R}^d} \langle z, x_{\ell_2} \rangle \mathcal{P}(O_{\ell_2,\ell_1} z) \, \mathbb{1}_{\langle z, x_{\ell_2} \rangle \geq T} \; dz && \text{(D.16)} \\
&= \int_{\mathbb{R}^d} \langle z, x_{\ell_2} \rangle \mathcal{P}(z) \, \mathbb{1}_{\langle z, x_{\ell_2} \rangle \geq T} \; dz && \text{(D.17)} \\
&= \langle x_{\ell_2}, m_{\ell_2}^{(T)} \rangle, && \text{(D.18)}
\end{aligned}$$

where (D.13) follows from the dentition of $m_\ell^{(T)}$; (D.14) follows from the orthonormal operator definition of O_{ℓ_2,ℓ_1}; (D.15) follows from substitution of variables $z \to O_{\ell_2,O_{\ell_1}} \cdot z$; (D.16) follows from the property of orthonormal operator O, which satisfies, $\langle O \cdot a, O \cdot b \rangle = \langle a, b \rangle$ for every $a, b \in \mathbb{R}^d$; (D.17) follows from the isotropic property of $\mathcal{P}(z) = \mathcal{P}(O \cdot z)$, that is for every orthonormal operator O and for every $z \in \mathbb{R}^d$, $\|O \cdot z\| = \|z\|$; and (D.18) follows from the dentition of $m_\ell^{(T)}$.

Next, we show that $\alpha_\ell \geq T$. As the templates $\{x_\ell\}_{\ell=0}^{L-1}$ are normalized, $\|x_\ell\| = 1$, for every $\ell \in [L]$, then,

$$\alpha_\ell = \|\langle m_\ell^{(T)}, x_\ell \rangle\| = \|m_\ell^{(T)}\|. \tag{D.19}$$

Then, following the definition of $m_\ell^{(T)}$ (C.1) and the SLLN (D.3), we have,

$$\begin{aligned}
\alpha_\ell = \left\| \langle m_\ell^{(T)}, x_\ell \rangle \right\| &= \left\| \frac{1}{|\mathcal{A}_\ell^{(T)}|} \sum_{z_i \in \mathcal{A}_\ell} \langle z_i, x_\ell \rangle \right\| \\
&\xrightarrow{\text{a.s.}} \left\| C_T \int_{\mathbb{R}^d} \langle z, x_\ell \rangle \mathcal{P}(z) \, \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \; dz \right\| && \text{(D.20)}
\end{aligned}$$

Similar to Lemma D.1, for every $z \in \mathbb{R}^d$, we can decompose it into $z = \beta_z x_\ell + \gamma_z w_z$, where $\langle w_z, x_\ell \rangle = 0$, and $\beta_z \geq T$, by the definition of the domain of the integral. Thus, we can decompose the integral in (D.20) as follows:

$$\int_{\mathbb{R}^d} \langle z, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz = \int_{\mathbb{R}^d} \langle \beta_z x_\ell, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz + \int_{\mathbb{R}^d} \langle \gamma_z w_z, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz \tag{D.21}$$

By Lemma D.1, there exists $z' = \beta_z x_\ell - \gamma_z w_z$, such that $\|z'\| = \|z\|$, and $\langle z', x_\ell \rangle \geq T$. Thus, as $\mathcal{P}$ is isotropic, that is, $\mathcal{P}(z) = \mathcal{P}(z')$ for $\|z\| = \|z'\|$, the last term in (D.21) vanishes:

$$\int_{\mathbb{R}^d} \langle \gamma_z w_z, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz = 0. \tag{D.22}$$

Substituting (D.22) into (D.21) results,

$$\begin{aligned} \int_{\mathbb{R}^d} \langle z, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz &= \int_{\mathbb{R}^d} \langle \beta_z x_\ell, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz \\ &= \int_{\mathbb{R}^d} \beta_z \|x_\ell\|^2 \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz \\ &= \int_{\mathbb{R}^d} \beta_z \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz, \end{aligned} \tag{D.23}$$

where we have used $\|x_\ell\| = 1$. As $\beta_z \geq T$ for every $z \in \mathcal{V}_\ell^{(T)}$, we have,

$$C_T \int_{\mathbb{R}^d} \beta_z \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz \geq T. \tag{D.24}$$

Substituting (D.24) into (D.20) results,

$$\alpha_\ell \xrightarrow{\text{a.s.}} \left\| C_T \int_{\mathbb{R}^d} \langle z, x_\ell \rangle \mathcal{P}(z) \mathbb{1}_{\langle z, x_\ell \rangle \geq T} \, dz \right\| \geq T, \tag{D.25}$$

which completes the proof of the lemma.

D.1.3 Proof of Lemma D.3

Assume without loss of generality that $x_\ell = e_1$, the first standard basis vector. This can always be arranged by rotating coordinates, due to the isotropy of the Gaussian distribution (see Lemma D.2). Let $z = (z_1, z_2, \ldots, z_d) \sim \mathcal{N}(0, \sigma^2 I_{d \times d})$. Then, $z_1 \sim \mathcal{N}(0, \sigma^2)$, independent of $z_2, \ldots, z_d$.

The distribution $g_\ell(z)$, defined in (B.7), can be represented by

$$g_\ell(z) = \frac{1}{\mathbb{P}(z_1 \geq T)} \cdot \mathcal{P}(z) \cdot \mathbb{1}_{\{z_1 \geq T\}}, \tag{D.26}$$

where $\mathcal{P}(z) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right)$ is the isotropic Gaussian density with mean 0 and covariance $\sigma^2 I_{d\times d}$. Thus, the expectation becomes:

$$\mathbb{E}_{g_\ell}[z] = \frac{1}{\mathbb{P}(z_1 \geq T)} \int_{z_1 \geq T} z \cdot \mathcal{P}(z)\, dz. \tag{D.27}$$

Because the Gaussian is symmetric and the truncation is only in the z_1 direction, the only nonzero component of the expectation is in the e_1 direction. That is,

$$\mathbb{E}_{g_\ell}[z] = \mathbb{E}_{g_\ell}[z_1] \cdot e_1, \tag{D.28}$$

and thus we have,

$$\left\| \frac{\mathbb{E}_{g_\ell}[z]}{T} \right\| = \frac{\mathbb{E}_{g_\ell}[z_1]}{T}. \tag{D.29}$$

Our aim is to compute the right-hand-side of (D.29) for $T \to \infty$,

$$\lim_{T\to\infty} \frac{1}{T} \cdot \frac{1}{\mathbb{P}(z_1 \geq T)} \int_T^\infty z_1 \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{z_1^2}{2\sigma^2}\right) dz_1. \tag{D.30}$$

Let us define the one-dimensional Gaussian probability density function:

$$\varphi_\sigma(z) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-z^2/(2\sigma^2)}, \tag{D.31}$$

and the upper right tail:

$$Q_\sigma(T) = \int_T^\infty \varphi_\sigma(z)\, dz. \tag{D.32}$$

Then, substituting (D.31) and (D.32) into (D.27) results,

$$\mathbb{E}_{g_\ell}[z_1] = \frac{1}{Q_\sigma(T)} \int_T^\infty z_1 \varphi_\sigma(z_1)\, dz_1. \tag{D.33}$$

This expression corresponds to the conditional mean of a Gaussian random variable truncated from below. For $X \sim \mathcal{N}(0, \sigma^2)$ it is known that [17]

$$\mathbb{E}[X \mid X \geq T] = \sigma \cdot \frac{\varphi(T/\sigma)}{Q(T/\sigma)}, \tag{D.34}$$

where φ is the standard normal density and Q is the standard upper tail probability.

Using the following known asymptotic relation [17],

$$\frac{\varphi(t)}{Q(t)} \sim t + \frac{1}{t}, \quad \text{as } t \to \infty. \tag{D.35}$$

and by (D.34), we have,

$$\mathbb{E}[X \mid X \geq T] = \sigma \cdot \frac{\varphi(T/\sigma)}{Q(T/\sigma)} \sim T + \frac{\sigma^2}{T}, \tag{D.36}$$

as $T \to \infty$. Therefore,

$$\frac{\mathbb{E}_{g_\ell}[z_1]}{T} = \frac{\mathbb{E}[X \mid X \geq T]}{T} \sim \frac{T + \sigma^2/T}{T} = 1 + \frac{\sigma^2}{T^2}. \tag{D.37}$$

Hence, combining (D.29), (D.30), and (D.37) results,

$$\lim_{T\to\infty} \left\| \frac{\mathbb{E}_{g_\ell}[z]}{T} \right\| = \lim_{T\to\infty} \frac{\mathbb{E}_{g_\ell}[z_1]}{T} = 1, \tag{D.38}$$

completing the proof of the lemma.

D.2 Proof of Theorem C.2

Recall the definitions of the *postulated* GMM model, given in (B.2), and the *actual* statistics of the extracted particles, given in (B.6). Their respective probability density functions are $f, g : \mathbb{R}^d \to \mathbb{R}$, each consists of L components, and are defined by,

$$f\left(z; \{\mu_\ell\}_{\ell\in[L]}\right) = \sum_{\ell=0}^{L-1} w_\ell f_\ell(z) = \sum_{\ell=0}^{L-1} w_\ell \cdot \mathcal{N}\left(z; \mu_\ell, \Sigma_\ell = \sigma^2_{\mathsf{GMM}} I_{d\times d}\right), \tag{D.39}$$

and,

$$g\left(z; \{m_\ell\}_{\ell\in[L]}\right) = \sum_{\ell=0}^{L-1} \pi_\ell g_\ell(z) = C_T \sum_{\ell=0}^{L-1} \pi_\ell \cdot \exp\left(-\frac{\|z\|_2^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, x_\ell\rangle \geq T}, \tag{D.40}$$

where g_ℓ is defined by (B.7), $\{m_\ell\}_{\ell=0}^{L-1}$ are the means of the mixture component distributions $\{g_\ell\}_{\ell=0}^{L-1}$, and $\{\mu_\ell\}_{\ell=0}^{L-1}$ are the means of the Gaussian mixture component distributions $\{f_\ell\}_{\ell=0}^{L-1}$. By definition,

$$m_\ell = \mathbb{E}_{g_\ell}[z] = \int_{\mathbb{R}^d} z \cdot g_\ell(z) \; dz. \tag{D.41}$$

Note that $f, g \in C^0(\mathbb{R}^d)$ are continuous almost everywhere.

The output of Algorithm 1 is the set $\mathcal{A} = \{z_i\}_{i=0}^{M-1} \in \mathbb{R}^d$, which follows the distribution g. However, the maximum-likelihood estimators of the centers $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ are given with respect to the likelihood function of the postulated model f:

$$\{\hat{\mu}_\ell\}_{\ell=0}^{L-1} = \underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\max} \prod_{i=0}^{M-1} f\left(z_i; \{\mu_\ell\}_{\ell=0}^{L-1}\right). \tag{D.42}$$

The next lemma relates the maximum-likelihood estimators given in (D.42) to the true statistics of the extracted particles, g. The lemma is proved in Section D.2.1.

Lemma D.4. *Let $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ be the maximum-likelihood estimators given by (D.42), based on the postulated GMM model with probability density function f, given in (D.39). Let g represent the actual statistics of the extracted particles, given by (D.40). Then, as $N \to \infty$,*

$$\begin{aligned}\{\hat{\mu}_\ell\}_{\ell=0}^{L-1} &\xrightarrow{\text{a.s.}} \underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\min} \int dz\ g\left(z; \{m_\ell\}_{\ell=0}^{L-1}\right) \log \frac{g\left(z; \{m_\ell\}_{\ell=0}^{L-1}\right)}{f\left(z; \{\mu_\ell\}_{\ell=0}^{L-1}\right)} \\ &= \underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\min} D_{\mathsf{KL}}\left(g || f\left(\{\mu_\ell\}_{\ell=0}^{L-1}\right)\right), \end{aligned} \tag{D.43}$$

where D_{KL} is the Kullback–Leibler divergence.

In other words, the maximum-likelihood estimator of the means $\{\hat{\mu}_\ell\}_{\ell\in[L]}$, given by (D.42), converges to the estimators that minimize the Kullback–Leibler (KL) divergence between the expected GMM statistics f, and the true probability distribution g, from which the extracted particles were generated. Thus, the maximum-likelihood estimator provides the closest distribution (within the constraints of the GMM model, dependent on $\{\hat{\mu}_\ell\}_{\ell\in[L]}$) to the true distribution g.

We substitute the expression of f, g following from (D.39), and (D.40) into (D.43), and obtain,

$$\begin{aligned}\{\hat{\mu}_\ell\}_{\ell=0}^{L-1} &\xrightarrow{\text{a.s.}} \underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\max} \int dz\ g\left(z; \{m_\ell\}_{\ell=0}^{L-1}\right) \log f\left(z \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) \\ &= \underset{\{\mu_k\}_{k=0}^{L-1}}{\arg\max} \int dz \sum_{\ell=0}^{L-1} \pi_\ell g_\ell(z) \log\left[\sum_{k=0}^{L-1} w_k f_k(z; \mu_k)\right] \\ &= \underset{\{\mu_k\}_{k=0}^{L-1}}{\arg\max} \sum_{\ell=0}^{L-1} \pi_\ell \int dz\ g_\ell(z) \log\left[\sum_{k=0}^{L-1} w_k f_k(z; \mu_k)\right]. \end{aligned} \tag{D.44}$$

We define the expression within the arg max term in (D.44) by:

$$\mathcal{H}\left(\{\mu_\ell\}_{\ell=0}^{L-1}\right) \triangleq \sum_{\ell=0}^{L-1} \pi_\ell \int dz\ g_\ell(z) \log\left[\sum_{k=0}^{L-1} w_k f_k(z; \mu_k)\right]. \tag{D.45}$$

The next lemma shows that as $T \to \infty$ the means $\{\mu_k\}_{k\in[L]}$ that maximize $\mathcal{H}$ converge to the means $\{m_\ell\}_{\ell\in[L]}$ of the true probability distribution g. The lemma is proved in Section D.2.2.

Lemma D.5. *Let $\mathcal{H}$ denote the function defined in (D.45). Then, as $T \to \infty$, the $\{\mu_k\}_{k\in[L]}$ that maximize $\mathcal{H}$ are given by,*

$$\lim_{T\to\infty}\left[\underset{\{\mu_\ell\}_{\ell=0}^{L-1}}{\arg\max}\left(\frac{\mathcal{H}\left(\{\mu_\ell\}_{\ell=0}^{L-1}\right)}{T}\right)\right] = \{x_\ell\}_{\ell=0}^{L-1}. \tag{D.46}$$

The global maximum is unique, up to a permutation between the centers of the likelihood function, i.e., there exists a permutation function $\pi : [L] \to [L]$, such that for every $0 \leq \ell \leq L-1$,

$$\lim_{T\to\infty} \frac{\hat{\mu}_{\pi(\ell)}}{m_\ell} = 1. \tag{D.47}$$

Combining (D.42), Lemma D.4, (D.45), and Lemma D.5 completes the proof.

D.2.1 Proof of Lemma D.4

Rearranging (D.42) results,

$$\begin{aligned}
\{\hat{\mu}_\ell\}_{\ell=0}^{L-1} &= \arg\max_{\{\mu_\ell\}_{\ell=0}^{L-1}} \prod_{i=0}^{M-1} f\left(z_i \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) \\
&= \arg\max_{\{\mu_\ell\}_{\ell=0}^{L-1}} \sum_{i=0}^{M-1} \log f\left(z_i \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) \\
&= \arg\max_{\{\mu_\ell\}_{\ell=0}^{L-1}} \frac{1}{M}\sum_{i=0}^{M-1} \log f\left(z_i \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right) - \frac{1}{M}\sum_{i=0}^{M-1} \log g\left(z_i; \{m_\ell\}_{\ell=0}^{L-1}\right) \\
&= \arg\max_{\{\mu_\ell\}_{\ell=0}^{L-1}} \frac{1}{M}\sum_{i=0}^{M-1} \log \frac{f\left(z_i \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right)}{g\left(z_i; \{m_\ell\}_{\ell=0}^{L-1}\right)} \\
&\xrightarrow{\text{a.s.}} \arg\min_{\{\mu_\ell\}_{\ell=0}^{L-1}} \int dz\ g\left(z; \{m_\ell\}_{\ell=0}^{L-1}\right) \log \frac{g\left(z; \{m_\ell\}_{\ell=0}^{L-1}\right)}{f\left(z \mid \{\mu_\ell\}_{\ell=0}^{L-1}\right)} \\
&= \arg\min_{\{\mu_\ell\}_{\ell=0}^{L-1}} D_{\mathsf{KL}}\left(g || f\left(\{\mu_\ell\}_{\ell=0}^{L-1}\right)\right),
\end{aligned} \tag{D.48}$$

where the last transition follows from the strong law of large number (SLLN), and D_{KL} represents the Kullback-Leibler divergence [28].

D.2.2 Proof of Lemma D.5

To obtain the maximum argument of (D.45), we derive it with respect to each one of the means $\{\mu_\ell\}_{\ell=0}^{L-1}$. Recall the definition of g_ℓ in (B.7) and of the GMM distribution f in (D.39). Then,

$$\frac{\partial}{\partial \mu_k}\mathcal{H} = \frac{1}{\sigma_{\mathsf{GMM}}^2}\sum_{\ell=0}^{L-1}\pi_\ell \int dz\ g_\ell(z)\frac{w_k f_k(z;\mu_k)}{\sum_{r=0}^{L-1} w_r f_r(z;\mu_r)}(z-\mu_k) \tag{D.49}$$

$$= \frac{C_T}{\sigma_{\mathsf{GMM}}^2}\sum_{\ell=0}^{L-1}\pi_\ell \int dz\ \exp\left(-\frac{\|z\|^2}{2}\right)\mathbb{1}_{\langle z, x_\ell\rangle \geq T}\frac{w_k f_k(z;\mu_k)}{\sum_{r=0}^{L-1} w_r f_r(z;\mu_r)}(z-\mu_k) \tag{D.50}$$

To obtain the maximum of $\mathcal{H}$, the following condition is necessary:

$$\frac{\partial}{\partial \mu_k}\mathcal{H} = \sum_{\ell=0}^{L-1} \pi_\ell \int dz\ g_\ell(z) \frac{w_k f_k(z;\mu_k)}{\sum_{r=0}^{L-1} w_r f_r(z;\mu_r)} (z-\mu_k) = 0, \tag{D.51}$$

for every $k \in [L]$.

Our aim is to find sufficient conditions such that $\{\hat{\mu}_\ell\}_{\ell\in[L]} = \{m_\ell^{(T)}\}_{\ell\in[L]}$, i.e., that the maximum-likelihood estimators $\{\hat{\mu}_\ell\}_{\ell=0}^{L-1}$ would converge to the means of the true distribution of the particles $\{g_\ell\}_{\ell=0}^{L-1}$. Thus, the following should hold by substituting $\{\mu_\ell\}_{\ell=0}^{L-1} = \{m_\ell^{(T)}\}_{\ell\in[L]}$ into (D.51):

$$\sum_{\ell=0}^{L-1} \pi_\ell \int dz\ g_\ell(z) \frac{w_k f_k\left(z; m_k^{(T)}\right)}{\sum_{r=0}^{L-1} w_r f_r\left(z; m_r^{(T)}\right)} \left(z - m_k^{(T)}\right) = 0, \tag{D.52}$$

where (D.52) must hold for every $k \in [L]$. For notation simplicity, we remove the T transcript in $m_k^{(T)}$.

We rearrange the terms related to the distribution f in (D.52), following the definition of the GMM f from (D.39):

$$\begin{aligned}
\frac{w_k f_k(z;m_k)}{\sum_{r=0}^{L-1} w_r f_r(z;m_r)} &= \frac{w_k \exp\left(-\frac{\|z-m_k\|^2}{2\sigma_{\mathsf{GMM}}^2}\right)}{\sum_{r=0}^{L-1} w_r \exp\left(-\frac{\|z-m_r\|^2}{2\sigma_{\mathsf{GMM}}^2}\right)} \\
&= \frac{w_k \exp\left(-\frac{\|m_k\|^2 - 2\langle z, m_k\rangle}{2\sigma_{\mathsf{GMM}}^2}\right)}{\sum_{r=0}^{L-1} w_r \exp\left(-\frac{\|m_r\|^2 - 2\langle z, m_r\rangle}{2\sigma_{\mathsf{GMM}}^2}\right)} \\
&= \frac{w_k \exp\left\{\langle z, m_k\rangle/\sigma_{\mathsf{GMM}}^2\right\}}{\sum_{r=0}^{L-1} w_r \exp\left\{\langle z, m_r\rangle/\sigma_{\mathsf{GMM}}^2\right\}} \\
&= \frac{w_k \exp\left\{\langle z, \alpha x_k\rangle/\sigma_{\mathsf{GMM}}^2\right\}}{\sum_{r=0}^{L-1} w_r \exp\left\{\langle z, \alpha x_r\rangle/\sigma_{\mathsf{GMM}}^2\right\}},
\end{aligned} \tag{D.53}$$

where the last two steps followed from Proposition C.1, i.e., $\|m_\ell\| = \|m_k\| = \alpha$, for every $k, \ell \in [L]$. Recall that $\alpha \geq T$. Thus, for $T \to \infty$, we have for $\{w_\ell\}_{\ell\in[L]} > 0$ [2],

$$\lim_{T\to\infty} \frac{w_k \exp\left\{\alpha\langle z, x_k\rangle/\sigma_{\mathsf{GMM}}^2\right\}}{\sum_{r=0}^{L-1} w_r \exp\left\{\alpha\langle z, x_r\rangle/\sigma_{\mathsf{GMM}}^2\right\}} = \mathbb{1}\left\{\langle z, x_k\rangle = \max_{r\in[L]} \langle z, x_r\rangle\right\}. \tag{D.54}$$

Recall the definition of $\mathcal{V}_\ell^{(T)} \subset \mathbb{R}^d$ (C.3). The result in (D.54) implies that the components are well-separated, i.e., that $\mathcal{V}_{\ell_1}^{(T)}$ and $\mathcal{V}_{\ell_2}^{(T)}$ are disjoint sets as $T \to \infty$. Taking the limit $T \to \infty$ of each term of the summation in (D.52) results,

$$\lim_{T\to\infty} \frac{1}{T} \int dz\ g_\ell(z) \frac{w_k f_k(z;m_k)}{\sum_{r=0}^{L-1} w_r f_r(z;m_r)} (z - m_k)$$

$$= \lim_{T\to\infty} \frac{C_T}{T} \int_{z\in\mathcal{V}_\ell^{(T)}} dz \ \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{r=0}^{L-1} w_r f_r\left(z;m_r\right)} \left(z-m_k\right)$$

$$= \lim_{T\to\infty} \frac{C_T}{T} \int_{z\in\mathcal{V}_\ell^{(T)}} dz \ \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{r=0}^{L-1} w_r f_r\left(z;m_r\right)} z -$$

$$- \lim_{T\to\infty} \frac{C_T}{T} \int_{z\in\mathcal{V}_\ell^{(T)}} dz \ \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{r=0}^{L-1} w_r f_r\left(z;m_r\right)} m_k \tag{D.55}$$

Following from part (2) in Proposition C.1, (D.54), and by applying the dominated convergence theorem, which allows to interchange the sequence between the limits and the integral, we have,

$$\lim_{T\to\infty} \frac{C_T}{T} \int_{z\in\mathcal{V}_\ell^{(T)}} z \ \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{\ell=0}^{L-1} w_\ell f_\ell\left(z;m_\ell\right)} dz = \delta_{k\ell} x_\ell. \tag{D.56}$$

In addition, following from (D.54), we have,

$$\lim_{T\to\infty} C_T \int_{z\in\mathcal{V}_\ell^{(T)}} \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{\ell=0}^{L-1} w_\ell f_\ell\left(z;m_\ell\right)} dz = \delta_{k\ell}. \tag{D.57}$$

Substituting (D.56), and (D.57) into (D.55) leads to,

$$\lim_{T\to\infty} \frac{1}{T} \int dz \ g_\ell\left(z\right) \frac{w_k f_k\left(z;m_k\right)}{\sum_{r=0}^{L-1} w_r f_r\left(z;m_r\right)} \left(z-m_k\right)$$

$$= \delta_{k\ell} x_\ell - \delta_{k\ell} \lim_{T\to\infty} \frac{m_\ell^{(T)}}{T}. \tag{D.58}$$

From part (2) of Proposition C.1,

$$\lim_{T\to\infty} \frac{m_\ell^{(T)}}{T} = x_\ell. \tag{D.59}$$

Substituting (D.59) into (D.58) leads to (D.52). Thus, as $T \to \infty$, the means which maximizes $\mathcal{H}$ are the means $\{m_\ell\}_{\ell\in[L]}$ of the true components g, completing the proof of the lemma.

D.3 Proof of Theorem C.3

Let $x_\ell = R_\ell \cdot V_{\mathsf{template}}$ denote the input templates used in Algorithm 1. By Assumption B.5, the SO(3) samples used for generating these templates during the particle-picking stage, denoted by $\{R_\ell\}_{\ell\in\{0,\dots,L-1\}}$, is the same as the samples used during the 3D reconstruction step. That is, the samples in (B.9) and (B.12) are identical $\{R_\ell\}_{\ell\in\{0,\dots,L-1\}} = \{R_n\}_n$. Recall the assumed probability density function used during 3D reconstruction, defined in (B.13):

$$\bar{f}_{\mathrm{ET}}(z;V) = \frac{1}{(2\pi\sigma_V^2)^{d/2}} \sum_{R_\ell\in\mathrm{SO}(3)} w(R_\ell) \cdot \exp\left(-\frac{\|z-(R_\ell\cdot V)\|^2}{2\sigma_V^2}\right)$$

$$= \sum_{\ell=0}^{L-1} w(R_\ell) \cdot \bar{f}_\ell(z; R_\ell \cdot V), \tag{D.60}$$

where we define

$$\bar{f}_\ell(z; R_\ell \cdot V) \triangleq \frac{1}{(2\pi\sigma_V^2)^{d/2}} \exp\left(-\frac{\|z - R_\ell \cdot V\|_F^2}{2\sigma_V^2}\right). \tag{D.61}$$

In contrast, the actual distribution of the particles extracted by the template-matching particle picker is given in (B.9):

$$\begin{aligned} g_{\text{picker}}(z) &\triangleq C_T \sum_{\ell=0}^{L-1} \pi(R_\ell) \cdot \exp\left(-\frac{\|z\|^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, R_\ell \cdot V_{\mathsf{template}} \rangle \geq T} \\ &= \sum_{\ell=0}^{L-1} \pi(R_\ell) \cdot g_\ell(z), \end{aligned} \tag{D.62}$$

where

$$g_\ell(z) \triangleq C_T \cdot \exp\left(-\frac{\|z\|_F^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, x_\ell \rangle \geq T}. \tag{D.63}$$

The maximum likelihood estimator $\hat{V}$ is defined by

$$\hat{V} = \arg\max_V \bar{\mathcal{L}}_{\text{assumed}}(V; \mathcal{A}), \tag{D.64}$$

where the objective $\bar{\mathcal{L}}_{\text{assumed}}$ is specified in (B.12). Following the logic of Lemma D.4 and Equation (D.45), we define the function

$$\mathcal{H}(V) \triangleq \sum_{\ell=0}^{L-1} \pi(R_\ell) \int dz\, g_\ell(z) \log\left(\sum_{k=0}^{L-1} w(R_k) \cdot \bar{f}_k(z; R_k \cdot V)\right). \tag{D.65}$$

Maximizing $\mathcal{H}$ yields the maximum likelihood estimator:

$$\hat{V} = \arg\max_V \mathcal{H}(V). \tag{D.66}$$

We now state the following lemma, whose proof is given below.

Lemma D.6. *Let $\mathcal{H}(V)$ be defined as in (D.65). Then, in the limit as $T \to \infty$, the maximizer of $\mathcal{H}(V)/T$ converges to the template (up to a global rotation $R \in \mathsf{SO}(3)$):*

$$\lim_{T\to\infty} \left\{ \arg\max_V \left(\frac{\mathcal{H}(V)}{T}\right) \right\} = R \cdot V_{\mathsf{template}}. \tag{D.67}$$

Combining Equation (D.66) with Lemma D.6 completes the proof.

Proof of Lemma D.6. Recall the objective function $\mathcal{H}(V)$ defined in (D.65):

$$\mathcal{H}(V) = \sum_{\ell=0}^{L-1} \pi(R_\ell) \int dz\, g_\ell(z) \log \left(\sum_{k=0}^{L-1} w(R_k) \cdot \bar{f}_k(z; R_k \cdot V) \right), \tag{D.68}$$

where $g_\ell(z) = C_T \cdot \exp\left(-\frac{\|z\|_F^2}{2\sigma^2}\right) \mathbb{1}_{\langle z, x_\ell \rangle \geq T}$ and $\bar{f}_k(z; R_k \cdot V)$ is the Gaussian density centered at $R_k \cdot V$ with variance σ_V^2. As $T \to \infty$, the support of $g_\ell(z)$ becomes increasingly concentrated on the set of z such that $\langle z, x_\ell \rangle \geq T$. Since the Gaussian term $\exp\left(-\frac{\|z\|^2}{2\sigma^2}\right)$ suppresses contributions from large-$\|z\|$ regions, the dominant contribution to the integral arises from values of z that are nearly colinear with x_ℓ, and of moderate norm. Specifically, $g_\ell(z)$ concentrates around the direction of $x_\ell = R_\ell \cdot V_{\text{template}}$ with increasing magnitude.

In the limit $T \to \infty$, the distribution $g_\ell(z)$ converges (in the weak sense) to a point mass supported at x_ℓ. That is,

$$g_\ell(z) \xrightarrow[T \to \infty]{} \delta(z - x_\ell), \tag{D.69}$$

up to normalization. Thus, the integral becomes

$$\int dz\, g_\ell(z) \log \left(\sum_{k=0}^{L-1} w(R_k) \cdot \bar{f}_k(z; R_k \cdot V) \right) \xrightarrow[T \to \infty]{} \log \left(\sum_{k=0}^{L-1} w(R_k) \cdot \bar{f}_k(x_\ell; R_k \cdot V) \right). \tag{D.70}$$

Plugging this into the expression for $\mathcal{H}(V)$, we find that:

$$\frac{\mathcal{H}(V)}{T} \xrightarrow[T \to \infty]{} \sum_{\ell=0}^{L-1} \pi(R_\ell) \log \left(\sum_{k=0}^{L-1} w(R_k) \cdot \bar{f}_k(x_\ell; R_k \cdot V) \right). \tag{D.71}$$

Now observe that for each ℓ, the function $\bar{f}_k(x_\ell; R_k \cdot V)$ is maximized when $R_k \cdot V = x_\ell = R_\ell \cdot V_{\text{template}}$, due to the form of the Gaussian kernel centered at $R_k \cdot V$. Therefore, the log-sum is maximized globally when $V = R \cdot V_{\text{template}}$, for some rotation $R \in \mathsf{SO}(3)$, since this makes $R_k \cdot V$ match x_k exactly for all k (up to a global rotation). Hence,

$$\lim_{T \to \infty} \left\{ \arg\max_V \left(\frac{\mathcal{H}(V)}{T} \right) \right\} = R \cdot V_{\text{template}}, \tag{D.72}$$

for $R \in \mathsf{SO}(3)$. □