A note on the sample complexity of multi-target detection

Amnon Balanov, Shay Kreymer, and Tamir Bendory

***Abstract*—This work studies the sample complexity of the multi-target detection (MTD) problem, which involves recovering a signal from a noisy measurement containing multiple instances of a target signal in unknown locations, each transformed by a random group element. This problem is primarily motivated by single-particle cryo-electron microscopy (cryo-EM), a groundbreaking technology for determining the structures of biological molecules. We establish upper and lower bounds for various MTD models in the high-noise regime as a function of the group, the distribution over the group, and the arrangement of signal occurrences within the measurement. The lower bounds are established through a reduction to the related multi-reference alignment problem, while the upper bounds are derived from explicit recovery algorithms utilizing autocorrelation analysis. These findings provide fundamental insights into estimation limits in noisy environments and lay the groundwork for extending this analysis to more complex applications, such as cryo-EM.**

I. Introduction

We study the multi-target detection (MTD) problem, where multiple signals are embedded at unknown positions within a long, noisy observation y [8]. Let $x \in V$ be the target signal to be estimated, where V is a finite-dimensional vector space. To present the problem, we consider $V = \mathbb{R}^L$; however, we later extend the discussion to include band-limited images. Let $g_i \in G$ be a group element drawn i.i.d. from a distribution ρ over G, and let $x_i = g_i \cdot x \in \mathbb{R}^L$. An MTD observation $y \in \mathbb{R}^{LM}$ is given by

$$y = \sum_{i=0}^{N-1} s_i * x_i + \epsilon, \tag{I.1}$$

where $*$ denotes linear convolution, N is the number of signal occurrences, and $\epsilon \sim \mathcal{N}(0, \sigma^2 I_{LM\times LM})$. The unknown locations of the signal occurrences are represented by location indicators $\{s_i\}_{i=0}^{N-1}$. Each s_i is a one-hot vector, where $s_i[j] = 1$ specifies that the first entry of x_i is placed at entry j of the observation y. The locations $\{s_i\}_{i=0}^{N-1}$ are assumed to be deterministic, but unknown. Since the statistics of y is the same for x and $g \cdot x$ for any $g \in G$, we aim to recover the G-orbit of x $\{g \cdot x \,|\, g \in G\}$.

Motivation: Cryo-EM. A key application of the MTD model is single particle cryo-electron microscopy (cryo-EM) [23], [7], [25]. In cryo-EM, 3-D biological samples, such as macromolecules or viruses, are rapidly frozen in a thin layer of vitreous ice (Figure 1(a)). The specimen is then imaged using an electron microscope that captures 2D tomographic images, known as *micrographs*. Each micrograph contains multiple tomographic projections of the molecules, where the 3-D orientation and 2-D position of each molecule are unknown. Therefore, cryo-EM models can be described similarly to a 2-D version of (I.1), where each signal occurrence is of the form $x_i = \Pi(g_i \cdot x)$, where g_i is a 3-D rotation and Π is a tomographic projection. (Note that the tomographic projection is not a group action, and thus (I.1) does not fully capture the cryo-EM model.)

Traditional methods for reconstructing the 3D structure from micrographs typically involve a two-step process: detecting and extracting the projection images from the micrographs, followed by 3-D reconstruction from those images. However, reliably detecting individual particles is challenging in high-noise environments. This is especially important for small molecular structures that induce high noise levels. Therefore, it is typically claimed that recovering small molecular structures using cryo-EM is impossible [15]. In this context, the MTD model was introduced as a computational framework to circumvent particle picking by directly reconstructing the 3D structure from the micrograph. Consequently, the inability to perform particle picking does not imply the impossibility of recovering the molecular structure, thereby enabling the recovery of small molecular structures [9], [19]. This technique also mitigates the ubiquitous bias issues in particle picking [16], [3].

Contribution. The main goal of this study is to derive upper and lower bounds on the sample complexity of the MTD problem in the high noise regime, focusing on three specific instances. Section II provides some basic definitions, and Section III introduces autocorrelation analysis, which is required for deriving the upper bounds of sample complexity. Section IV presents the multi-reference alignment (MRA) model and its tight connec-

The authors are with the School of Electrical Engineering, Tel Aviv University, Israel. The research is supported by the BSF grant no. 2020159, the NSF-BSF grant no. 2019752, the ISF grant no. 1924/21, and a grant from The Center for AI and Data Science at Tel Aviv University (TAD).

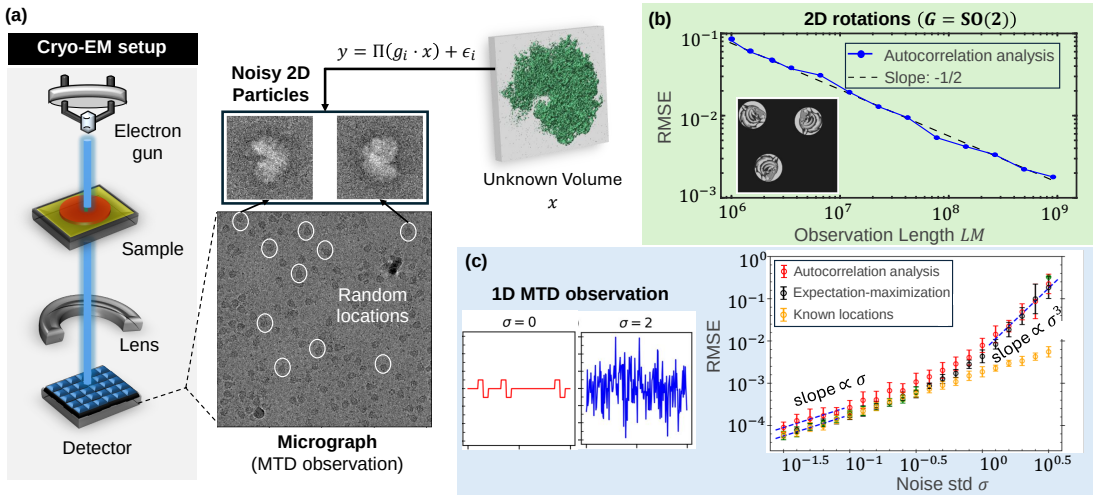

Fig. 1: **(a)** Single-particle electron microscopy reconstructs 3D structures from 2D projections [7]. Particles in vitrified ice form a *micrograph*, modeled using the MTD framework (I.1), where $x_i = \Pi(g_i \cdot x)$, with g_i a 3D rotation and Π as a tomographic projection. In high noise levels, direct particle detection is infeasible, but 3D reconstruction might be feasible by directly processing the micrographs [9]. **(b)** The estimation error (RMSE) versus observation length for MTD with 2D rotations based on autocorrelation analysis (taken from [17]). As seen empirically, the image can be estimated using the third-order autocorrelation. **(c)** An MTD observation of 1D signals in which identical signals $x_i = x$ are located at unknown positions in the observation y. In the high noise regime, their locations cannot be estimated, but the signal x can be estimated accurately. The estimation error (RMSE) scales as σ^3 in high noise levels, consistent across autocorrelation analysis and expectation-maximization (taken from [20]).

tion to the MTD model, which is the cornerstone of the sample complexity lower bounds. Section V integrates these findings, focusing on three specific MTD models, and Section VI delineates how these techniques could be extended to derive the sample complexity of the cryo-EM technology, the ultimate goal of this research. Due to space constraints, the proofs are omitted from this manuscript and are available in [4].

II. Preliminaries

A. Assumptions on the MTD model

We consider the asymptotic case of $\sigma, N, M \to \infty$, such that the density $\gamma = N/M < 1$ remains fixed. We assume that γ and the noise variance σ^2 are known.

Throughout, we assume that the signal occurrences do not overlap. This, in turn, means that their starting indices are separated by at least L entries so that their supports do not intersect, and the corresponding signal instances do not interfere with each other within the observation y. In some cases, we consider a stronger assumption that the minimal separation is doubled and refer to this model as the well-separated case. In this case, the starting positions of any two occurrences must be separated by at least $2L$ positions, ensuring that their endpoints are separated by at least L signal-free entries (although still noisy) in the data.

Definition II.1 (Separation)**.** *Consider the following separation condition:*

$$\text{If } s_{i_1}[j_1] = 1 \text{ and } s_{i_2}[j_2] = 1 \text{ for } i_1 \neq i_2, \quad \text{then } |j_1 - j_2| \geq L_{sep}. \tag{II.1}$$

If the MTD model satisfies (II.1) with $L_{sep} \geq L$, we say it is the non-overlapping; if $L_{sep} \geq 2L$, we say it is well-separated.

Remark II.2. *The results of this paper regarding the well-separated model are valid for $L_{sep} \geq 2L - 1$. Nevertheless, for the sake of simplifying the expressions throughout the text, we assume a separation of $2L$.*

B. Sample complexity definition

The objective of this work is to determine the sample complexity, i.e., the minimum number of measurements required to achieve a desired mean-squared error (MSE) between the true signal x and its estimator $\hat{x} = \hat{x}(y)$. Importantly, as the estimation of x is considered up to a group action on x, the MSE is defined as [1]:

$$\mathrm{MSE}\left(\hat{x}, x\right) = \frac{1}{\|x\|_F^2} \mathbb{E}\left[\min_{g \in G} \|g \cdot \hat{x} - x\|_F^2\right]. \tag{II.2}$$

Then, the sample complexity is defined as follows:

Definition II.3 (Sample complexity)**.** *Suppose that all the parameters of the MTD model (I.1) are fixed except for N, M, and σ. We define the smallest MSE attainable by any estimator as*

$$\mathsf{MSE}^*_{\mathsf{MTD}}(\sigma^2, N) := \inf_{\hat{x}} \mathbb{E}\left[\mathsf{MSE}\left(\hat{x}(y), x\right)\right], \tag{II.3}$$

and the sample complexity as,

$$N^*_{\mathsf{MTD}}\left(\sigma^2, \epsilon\right) := \min\left\{N : \mathsf{MSE}^*_{\mathsf{MTD}}(\sigma^2, N) \leq \epsilon\right\}. \tag{II.4}$$

In the above definition, ϵ represents a desired error, which in this paper approaches zero.

III. Autocorrelation analysis

In this section, we present the autocorrelation analysis method and provide several results that will be essential in establishing the upper bounds for sample complexity. To address the difficulty of locating signals at high noise levels [2], [13], previous works suggested the use of features that are invariant to unknown locations and possibly to the unknown group actions [8], [9], [17], [20], [12].

Definition III.1 (Empirical autocorrelations)**.** *Let $z \in \mathbb{R}^{LM}$ be an observation following the MTD model in (I.1). The empirical d-order autocorrelation of z is defined by:*

$$a_z^{(d)}[\ell_1, \ell_2, ..., \ell_{d-1}] = \frac{1}{LM} \sum_{j=0}^{LM-1} \tilde{z}[j]\tilde{z}[j+\ell_1]...\tilde{z}[j+\ell_{d-1}], \quad \text{(III.1)}$$

for $0 \le \ell_1, \ell_2, ..., \ell_{d-1} \le L-1$, where $\tilde{z}$ is the padding with zeros of z to a length of $(M+1)L$.

Definition III.2 (Autocorrelation ensemble mean)**.** *Let $X_G = g \cdot x \in \mathbb{R}^L$, where $g \in G$ is a random group element drawn from the distribution $g \sim \rho$. Define*

$$Y = \tilde{X}_G + \epsilon \in \mathbb{R}^{3L}, \quad \text{(III.2)}$$

where $\tilde{X}_G = [\mathbf{0}_L, X_G, \mathbf{0}_L]$ is the padding of X_G with zeros from left and right by L zeros each, and $\epsilon \sim \mathcal{N}\left(0, \sigma^2 I_{3L\times 3L}\right)$. Then, the d-order autocorrelation ensemble mean of Y is defined by,

$$\bar{a}_{Y,\rho}^{(d)}[\ell_1, \ell_2, ..., \ell_{d-1}] = \frac{1}{2L} \sum_{j=0}^{2L-1} \mathbb{E}_{G\sim\rho,\epsilon}\left\{Y[j]Y[j+\ell_1]...Y[j+\ell_{d-1}]\right\}, \quad \text{(III.3)}$$

for $0 \le \ell_1, \ell_2, ..., \ell_{d-1} \le L-1$.

The following proposition shows, for the well-separated case, that the empirical autocorrelation converges almost surely to the autocorrelation ensemble mean. This result is an extension of [8, Appendix A].

Proposition III.3. *Let MTD_G be a well-separated MTD model defined in (I.1) with a compact group G acting on $\mathbb{R}^L$, with group elements drawn according to a distribution ρ. Let $a_z^{(d)}$ be the empirical d-order autocorrelation as defined in (III.1), and $\bar{a}_{Y,\rho}^{(d)}$ be the d-order autocorrelation ensemble mean as defined in (III.3). Then, for every $0 \le \ell_1, \ell_2, ..., \ell_{d-1} \le L-1$,*

$$\lim_{N,M\to\infty} a_z^{(d)}[\ell_1, \ell_2, ..., \ell_{d-1}] \overset{\text{a.s.}}{=} 2\gamma \bar{a}_{Y,\rho}^{(d)}[\ell_1, \ell_2, ..., \ell_{d-1}] + (1-2\gamma)\,\chi[\ell_1, \ell_2, ..., \ell_{d-1}], \quad \text{(III.4)}$$

where

$$\chi[\ell_1, \ell_2, ..., \ell_{d-1}] = \mathbb{E}\left[\epsilon[0]\,\epsilon[\ell_1], ...\epsilon[\ell_{d-1}]\right].$$

Next, we show that if the autocorrelations up to order $\bar{d}$ uniquely define an orbit of the signal, then the sample complexity is bounded above by $\omega\left(\sigma^{2\bar{d}}\right)$[1].

Proposition III.4. *Assume the same conditions as in Proposition III.3. In addition, assume the parameter space Θ of the unknown signals $x \in \Theta$ is compact. Then, if the autocorrelation ensemble mean up to order $\bar{d}$, as defined in (III.3), uniquely determines the orbit of the signal, then the sample complexity of recovering the orbit is upper bounded by $\omega\left(\sigma^{2\bar{d}}\right)$.*

IV. The multi-reference alignment model

A. The multi-reference alignment model

To establish the lower bounds for the MTD models, we first introduce the multi-reference alignment (MRA) model [6]. Formally, the MRA model is defined as,

$$z_i = g_i \cdot x + \epsilon_i, \quad g_i \in G, \quad \text{(IV.1)}$$

where G is a known compact group, the group elements are drawn from some distribution ρ over G, and $\epsilon_i \sim \mathcal{N}(0, \sigma^2 I)$. The objective is to estimate $x \in \mathbb{R}^L$ from N observations $\{z_i\}_{i=0}^{N-1}$. From the perspective of this paper, the MRA problem can be interpreted as a simplified model of the MTD problem, when the signal locations are known. The sample complexity of the MRA problem can be defined analogously to the sample complexity of Definition II.3, where the number of observations N replaces the number of signal occurrences, and the observed data is now $\{z_i\}_{i=0}^{N-1}$. We denote the sample complexity of MRA by $N^*_{\mathsf{MRA}_G}\left(\sigma^2, \epsilon\right)$.

B. Mapping between the MRA model and the associated MTD model.

Comparing the MRA model (IV.1) with the MTD model (I.1) shows that the MTD model has an additional degree of uncertainty due to the unknown locations. The following proposition shows that the sample complexity of the corresponding MRA problem lower bounds the sample complexity of the MTD model.

Proposition IV.1. *Consider an MRA model (IV.1), with a distribution ρ defined on a group G. Consider a corresponding MTD model (I.1) with the same group G and distribution ρ, with N signal occurrences located at unknown positions $\{s_i\}_{i=0}^{N-1}$, assuming the non-overlapping case (Definition II.1). Then, the sample*

[1] $n = \omega(\sigma^{2d})$ means that $n/\sigma^{2d} \to \infty$ as $n, \sigma \to \infty$.

complexity of the MTD model (I.1) is lower bounded by the sample complexity of the associated MRA model (IV.1), i.e.,

$$N^*_{\mathsf{MTD}_G}\left(\sigma^2,\epsilon\right) \geq N^*_{\mathsf{MRA}_G}\left(\sigma^2,\epsilon\right). \tag{IV.2}$$

V. Sample complexity bounds

We now establish lower and upper bounds on the sample complexity for three specific MTD instances. The lower bounds are derived by reducing the problem to the corresponding MRA models and leveraging existing results. The upper bounds are obtained through explicit recovery techniques based on autocorrelation analysis.

A. One-dimensional MTD with circular translations

The first MTD model considers $x \in \mathbb{R}^L$, G is the group of circular translations $G = \mathbb{Z}_L$, and ρ is a uniform distribution. We demonstrate that both the upper and lower bounds for the sample complexity are $\omega\left(\sigma^6\right)$.

Proposition V.1. *Consider the MTD model with a uniform distribution ρ of the group $G = \mathbb{Z}_L$ elements, acting on $x \in \mathbb{R}^L$. Assume that the signal x has a non-vanishing DFT. Then, as $\sigma, N \to \infty$, we have,*

1) *The sample complexity of the well-separated case is bounded from above by $\omega\left(\sigma^6\right)$.*
2) *The sample complexity of the non-overlapping case is bounded from below by $\omega\left(\sigma^6\right)$.*

We remark that if ρ is non-uniform, then the upper bound remains $\omega\left(\sigma^6\right)$ and the lower bound is $\omega\left(\sigma^4\right)$ [1].

B. MTD with two-dimensional rotations

Next, we consider the MTD problem of estimating a two-dimensional *band-limited* target image x acted upon by SO(2) rotations, drawn from a uniform distribution. By bandlimited images, we mean it can be represented by finitely many Fourier-Bessel coefficients or alternative steerable basis coefficients [22]. Importantly, although the results from the previous sections pertain to signals in $\mathbb{R}^L$, they can be naturally extended to this case with suitable modifications.

Proposition V.2. *Consider the MTD model of estimating non-overlapping images, acted upon by SO(2) rotations drawn from a uniform distribution. Assume that the image x has a finite expansion in a steerable basis (e.g., Fourier-Bessel) and all the coefficients are non-zero. Then, as $\sigma, N \to \infty$, the sample complexity is lower bounded by $\omega\left(\sigma^6\right)$.*

Previous empirical findings suggest that the third autocorrelation moment is sufficient to reconstruct the two-dimensional image [21], [17]; see Figure 1(b). Another study [18] showed that an approximate expectation-maximization method achieves comparable results. Based on this, we conjecture an upper bound of $\omega\left(\sigma^6\right)$.

C. MTD with one-dimensional single signal

We now consider the one-dimensional MTD model with no group action, $G = I$. In this case, the sample complexity is upper bounded by $\omega\left(\sigma^6\right)$.

Proposition V.3. *Consider the MTD model for $x \in \mathbb{R}^L$ and $G = I$. Assume that the signal x has non-vanishing entries. Then, as $\sigma, N \to \infty$, the sample complexity is upper bounded by $\omega\left(\sigma^6\right)$.*

The lower bound for this model has not yet been proven. Based on previous empirical results [20] (see Figure 1(c)), we conjecture a lower bound of $\omega(\sigma^6)$.

VI. Outlook

This work is a first step toward analyzing the sample complexity of the MTD model, where the ultimate goal is to understand the sample complexity of the cryo-EM.

Cryo-EM and related models. The lower bounds in this work are derived by reducing the MTD model to a simpler MRA model. Consequently, any findings on the sample complexity of the MRA problem directly translate into lower bounds for the MTD model. In particular, in [5], the authors derived the sample complexity of many different MRA models and cryo-EM for generic signals in some specific parameter regimes. In particular, this paper implies a lower bound of $\omega\left(\sigma^6\right)$ on the sample complexity of the cryo-EM problem, at least in some parameter regime. Similarly, for signals in $\mathbb{R}^L$ acted upon by elements of the dihedral group, the results of [11], [14] imply a lower bound of $\omega\left(\sigma^6\right)$ if the distribution over the group is uniform, and $\omega\left(\sigma^4\right)$ if the distribution is non-uniform.

Priors. The analysis of the lower bounds on the sample complexity can be extended to include a prior on the signal. For instance, in [24], the sample complexity of high-dimensional MRA under a Gaussian prior has been derived. Similarly, in [10], the sample complexity of signals that reside in a low-dimensional semi-algebraic set was studied. This generalization underscores the robustness of the derived bounds, demonstrating their applicability even when additional prior information about x is included in the model.

Upper bounds via autocorrelation analysis. The upper bounds are based on explicit algorithms for signal recovery from autocorrelations. To date, very few provable MTD algorithms have been designed. Designing such algorithms—besides their importance for validation, reproducibility, and consistency—will immediately imply sample complexity upper bounds.

References

[1] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *IEEE Transactions on Information Theory*, 65(6):3565–3584, 2018.

[2] Cecilia Aguerrebere, Mauricio Delbracio, Alberto Bartesaghi, and Guillermo Sapiro. Fundamental limits in multi-image alignment. *IEEE Transactions on Signal Processing*, 64(21):5707–5722, 2016.

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

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

[5] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Jonathan Niles-Weed, Amelia Perry, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *Applied and Computational Harmonic Analysis*, 66:236–319, 2023.

[6] Afonso S Bandeira, Moses Charikar, Amit Singer, and Andy Zhu. Multireference alignment using semidefinite programming. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 459–470, 2014.

[7] 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.

[8] Tamir Bendory, Nicolas Boumal, William Leeb, Eitan Levin, and Amit Singer. Multi-target detection with application to cryo-electron microscopy. *Inverse Problems*, 35(10):104003, 2019.

[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] Tamir Bendory, Nadav Dym, Dan Edidin, and Arun Suresh. A transversality theorem for semi-algebraic sets with application to signal recovery from the second moment and cryo-EM. *arXiv preprint arXiv:2405.04354*, 2024.

[11] Tamir Bendory, Dan Edidin, William Leeb, and Nir Sharon. Dihedral multi-reference alignment. *IEEE Transactions on Information Theory*, 68(5):3489–3499, 2022.

[12] Tamir Bendory, Ti-Yen Lan, Nicholas F Marshall, Iris Rukshin, and Amit Singer. Multi-target detection with rotations. *Inverse problems and imaging (Springfield, Mo.)*, 17(2):362, 2023.

[13] Marom Dadon, Wasim Huleihel, and Tamir Bendory. Detection and recovery of hidden submatrices. *IEEE Transactions on Signal and Information Processing over Networks*, 2024.

[14] Dan Edidin and Josh Katz. Generic orbit recovery from invariants of very low degree. *arXiv preprint arXiv:2408.09599*, 2024.

[15] Richard Henderson. The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological molecules. *Quarterly reviews of biophysics*, 28(2):171–193, 1995.

[16] 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.

[17] Shay Kreymer and Tamir Bendory. Two-dimensional multi-target detection: An autocorrelation analysis approach. *IEEE Transactions on Signal Processing*, 70:835–849, 2022.

[18] Shay Kreymer, Amit Singer, and Tamir Bendory. An approximate expectation-maximization for two-dimensional multi-target detection. *IEEE signal processing letters*, 29:1087–1091, 2022.

[19] 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.

[20] Ti-Yen Lan, Tamir Bendory, Nicolas Boumal, and Amit Singer. Multi-target detection with an arbitrary spacing distribution. *IEEE Transactions on Signal Processing*, 68:1589–1601, 2020.

[21] Chao Ma, Tamir Bendory, Nicolas Boumal, Fred Sigworth, and Amit Singer. Heterogeneous multireference alignment for images with application to 2D classification in single particle reconstruction. *IEEE Transactions on Image Processing*, 29:1699–1710, 2019.

[22] Nicholas F Marshall, Oscar Mickelin, and Amit Singer. Fast expansion into harmonics on the disk: a steerable basis with fast radial convolutions. *SIAM Journal on Scientific Computing*, 45(5):A2431–A2457, 2023.

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

[24] Elad Romanov, Tamir Bendory, and Or Ordentlich. Multi-reference alignment in high dimensions: sample complexity and phase transition. *SIAM Journal on Mathematics of Data Science*, 3(2):494–523, 2021.

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