IMAGE RECOVERY FROM ROTATIONAL AND TRANSLATIONAL INVARIANTS

Nicholas F. Marshall[⋆] *Ti-Yen Lan*[†] *Tamir Bendory*[‡] *Amit Singer*[⋆†]

[⋆]Department of Mathematics, Princeton University, Princeton, NJ, USA
[†]The Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA
[‡]School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel

ABSTRACT

We introduce a framework for recovering an image from its rotationally and translationally invariant features based on autocorrelation analysis. This work is an instance of the multi-target detection statistical model, which is mainly used to study the mathematical and computational properties of single-particle reconstruction using cryo-electron microscopy (cryo-EM) at low signal-to-noise ratios. We demonstrate with synthetic numerical experiments that an image can be reconstructed from rotational and translational invariants and show that the reconstruction is robust to noise. These results constitute an important step towards the goal of structure determination of small biomolecules using cryo-EM.

Index Terms— autocorrelation analysis, multi-target detection, single-particle cryo-electron microscopy

1. INTRODUCTION

We consider the problem of estimating a target image from a large $m \times m$ noisy measurement M that contains p randomly rotated copies of the target. More precisely, suppose that $f : \mathbb{R}^2 \to \mathbb{R}$ is a target image supported on the open unit disc, and f_ϕ is the rotation of f by angle ϕ about the origin. Let $F_\phi : \mathbb{Z}^2 \to \mathbb{R}$ be a discrete version of f_ϕ defined by

$$F_\phi(\vec{x}) := f_\phi\left(\vec{x}/n\right), \quad \text{for } \vec{x} \in \mathbb{Z}^2,$$

where $n \ll m$ is a fixed positive integer. We assume the measurement $M : \{1, \dots, m\}^2 \to \mathbb{R}$ is of the form

$$M(\vec{x}) := \sum_{j=1}^{p} F_{\phi_j}(\vec{x} - \vec{x}_j) + \varepsilon(\vec{x}), \text{ for } \vec{x} \in \{1, \dots, m\}^2, \quad (1)$$

where $\phi_1, \dots, \phi_p \in [0, 2\pi)$ are uniformly random angles, $\vec{x}_1, \dots, \vec{x}_p \in \{n, \dots, m - n + 1\}^2$ are arbitrary translations, and $\varepsilon(\vec{x})$ is i.i.d. Gaussian noise with mean 0 and variance σ^2. For simplicity, we assume that images in the measurement are separated by at least one image diameter, see (4). Recent work on a related problem suggests that this separation restriction can be alleviated [13].

NFM was supported by NSF DMS-1903015. TYL and AS were supported in part by Award Number FA9550-17-1-0291 from AFOSR, Simons Foundation Math+X Investigator Award, the Moore Foundation Data-Driven Discovery Investigator Award, and NSF BIGDATA Award IIS-1837992. The authors thank Iris Rukshin for stimulating discussions.

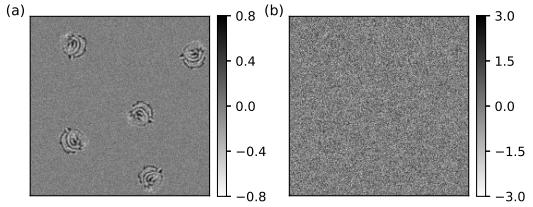

Fig. 1. Two measurements M with five target images and different noise levels: SNR = 10 (left), and SNR = 0.1 (right).

Our goal is to recover the image f from the noisy measurement M. With respect to this goal, the rotations $\phi_1, \dots, \phi_p$ and translations $\vec{x}_1, \dots, \vec{x}_p$ are nuisance parameters that do not necessarily need to be estimated. If the signal-to-noise ratio (SNR) is high as in Figure 1(a), then estimating these rotations and translations is straightforward, and the image can be recovered by aligning and averaging its different copies in the measurement. However, if the SNR is low as in Figure 1(b), which is the case of interest in this study, then this approach may be problematic as even detecting the image occurrences becomes challenging [1, 4].

We are interested in the following question: given a large enough measurement M, can we recover the target image f regardless of the level of noise? More precisely, if the number of target image occurrences p grows at a constant rate γ with the size m^2 of the measurement M, can we recover the image f for any fixed noise variance σ^2 as $m \to \infty$?

1.1. Motivation

We are motivated by challenges in single-particle cryo-electron microscopy (cryo-EM). The measurements in cryo-EM consist of 2-D tomographic projections of identical biomolecules of unknown 3-D orientations embedded in a large noisy image, called a micrograph. In the current analysis workflow of cryo-EM data [2, 7, 8, 17, 19, 20], the 2-D projections are first detected and extracted from the micro-

graph, and later rotationally and translationally aligned to reconstruct the 3-D molecular structure. This approach is problematic for small molecules, which are difficult to detect due to their lower SNR. This detection difficulty sets a lower bound on the usable molecule size in the current cryo-EM analysis workflow [9]. To circumvent this fundamental barrier, recent papers [3, 4] suggest to directly estimate the 3-D structure from the micrograph, without an intermediate detection stage; this approach was inspired by Kam [11] who introduced autocorrelation analysis to cryo-EM.

The estimation problem described in Section 1 is a simplified version of the cryo-EM reconstruction problem: the tomographic projection operator is omitted and we observe the same 2-D image multiple times with random in-plane rotations. This image recovery problem is an instance of the *multi-target detection* statistical model, in which a set of signals appear multiple times at unknown locations in a noisy measurement [3, 4, 13]. Here, we extend previous works by taking in-plane rotations into account, which forms an important step towards the analysis of the full cryo-EM problem.

1.2. Main contribution

The principal contribution of this paper is to demonstrate that a target image f can be recovered from a measurement M of the form (1) without estimating the rotations and translations of the images in the measurement. As a consequence, limits on estimating these nuisance parameters do not impose limits on estimating the target image f. In particular, we empirically demonstrate that the estimation of f is feasible at any noise level given a large enough measurement. From a computational perspective, this approach is highly efficient as it requires only one pass over the measurement to calculate the autocorrelations; this is in contrast to the likelihood-based techniques that may become prohibitive as the data size increases. For a more detailed discussion, see [13].

2. INVARIANT FEATURES

We begin by studying invariance in the continuous setting. The simplest invariant to both rotations and translations is

$$s_1 := \int_{\mathbb{R}^2} f(\vec{x}) d\vec{x},$$

which is the mean of the image; however, clearly more information is needed to recover the image. Motivated by autocorrelation analysis we next consider the rotationally-averaged second-order autocorrelation $s_2 : \mathbb{R}^2 \to \mathbb{R}$ by

$$s_2(\vec{x}_1) := \int_0^{2\pi} \int_{\mathbb{R}^2} f_\phi(\vec{x}) f_\phi(\vec{x}+\vec{x}_1) d\vec{x} d\phi.$$

By a change of variables of integration, it is clear that s_2 is only a function of $|\vec{x}_1|$, that is, it is a 1-D function, and thus does not contain enough information to encode a 2-D image. Thus, we proceed to consider the rotationally-averaged third-order autocorrelation $s_3 : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ defined by

$$s_3(\vec{x}_1, \vec{x}_2) := \int_0^{2\pi} \int_{\mathbb{R}^2} f_\phi(\vec{x}) f_\phi(\vec{x}+\vec{x}_1) f_\phi(\vec{x}+\vec{x}_2) d\vec{x} d\phi.$$

It is straightforward to verify that s_3 is a function of three parameters: $|\vec{x}_1|$, $|\vec{x}_2|$ and $\theta(\vec{x}_1, \vec{x}_2)$, where $\theta(\vec{x}_1, \vec{x}_2)$ denotes the angle between the vectors $\vec{x}_1, \vec{x}_2 \in \mathbb{R}^2$. In this work, we show empirically that f can be recovered from $S_3 : \mathbb{Z}^2 \times \mathbb{Z}^2 \to \mathbb{R}$, which is a discrete version of s_3 defined by

$$S_3(\vec{x}_1, \vec{x}_2) := \int_0^{2\pi} \sum_{\vec{x} \in \mathbb{Z}^2} F_\phi(\vec{x}) F_\phi(\vec{x}+\vec{x}_1) F_\phi(\vec{x}+\vec{x}_2) d\phi. \quad (2)$$

We claim that the function S_3 can be approximated from the measurement M by computing the third-order autocorrelation $A_3 : \mathbb{Z}^2 \times \mathbb{Z}^2 \to \mathbb{R}$ defined by

$$A_3(\vec{x}_1, \vec{x}_2) := \frac{1}{m^2} \sum_{\vec{x} \in \mathbb{Z}^2} M(\vec{x}) M(\vec{x}+\vec{x}_1) M(\vec{x}+\vec{x}_2). \quad (3)$$

Indeed, recall that the images are assumed to be separated in the measurement by one image diameter:

$$|\vec{x}_{j_1} - \vec{x}_{j_2}| > 4n, \quad \text{for} \quad j_1 \neq j_2, \quad (4)$$

and that the rotations in the measurement are chosen uniformly at random. Under these assumptions it is straightforward to show, see for example [3], that if $p/m^2 \to \gamma$, then

$$A_3(\vec{x}_1, \vec{x}_2) \to \frac{\gamma}{2\pi} S_3(\vec{x}_1, \vec{x}_2) + \sigma^2 \frac{\gamma S_1}{2\pi} \big(\delta(\vec{x_1}) + \delta(\vec{x_2}) + \delta(\vec{x_1} - \vec{x_2})\big), \quad (5)$$

as $m \to \infty$ for any fixed noise level σ^2, where S_1 is the discrete analog of s_1, and $\delta(\vec{x}) = 1$ if $\vec{x} = \vec{0}$ and is zero otherwise. In practice, σ^2 and γS_1 can be estimated from M, and thus, S_3 can be estimated from A_3 up to a constant factor.

We demonstrate empirically that the image f can be recovered from S_3. For simplicity of exposition, we describe a method of recovering f that just involves S_3, although we note that S_1 and S_2, which represent the discrete analogs of s_1 and s_2, may be used to aid in the reconstruction process.

The problem of estimating f from M, as described in Section 1, should not be confused with the related problem of recovering an image from a set of measurements $\{M_j\}_{j=1}^n$, each of which has exactly one shifted and rotated observation. Several moment-based techniques have been proposed to address this problem at high noise levels, see for instance [12, 15, 16, 18], but none can be applied directly to our problem where only a single, large observation M is available.

3. IMAGE RECOVERY FROM INVARIANTS

Recall that the target image $f : \mathbb{R}^2 \to \mathbb{R}$ is supported on the unit disc. We assume that f is bandlimited in the basis of Dirichlet Laplacian eigenfunctions on the unit disk. In polar coordinates (r, θ), these eigenfunctions are of the form

$$\psi_{\nu,q}(r,\theta) := J_\nu\left(\lambda_{\nu,q} r\right) e^{i\nu\theta}, \quad \text{for } (\nu, q) \in \mathbb{Z} \times \mathbb{Z}_{>0}, \quad (6)$$

where J_ν is the ν-th order Bessel function of the first kind, and $\lambda_{\nu,q} > 0$ is the q-th positive root of J_ν. The eigenvalue associated with $\psi_{\nu,q}$ is $\lambda^2_{\nu,q}$, and thus, the assumption that f is bandlimited can be written as

$$f(r,\theta) = \sum_{(\nu,q):\lambda_{\nu,q}\le\lambda} \alpha_{\nu,q}\psi_{\nu,q}(r,\theta), \quad \text{for } r \le 1, \quad (7)$$

where $\lambda > 0$ is the bandlimit frequency, and $\alpha_{\nu,q}$ are the expansion coefficients. A key property of this basis is that it is steerable in the sense that it diagonalizes the rotation operator:

$$f(r,\theta+\phi) = \sum_{(\nu,q):\lambda_{\nu,q}\le\lambda} \alpha_{\nu,q}\psi_{\nu,q}(r,\theta)e^{i\nu\phi}. \quad (8)$$

3.1. Computing invariants

Recall that $F_\phi : \mathbb{Z}^2 \to \mathbb{R}$ is the discretely sampled version of f_ϕ, which is defined by $F_\phi(\vec{x}) = f_\phi(\vec{x}/n)$ for $\vec{x} \in \mathbb{Z}^2$. In the following, we consider F_ϕ as a function on $\mathcal{J} := \{-2n, \dots, 2n-1\}^2 \subset \mathbb{Z}^2$. Since f_ϕ is supported on the open unit disc, it follows that $\text{supp}(F_\phi) \subset \{\vec{x} \in \mathcal{J} : |\vec{x}| < n\}$. Let $\hat{F}_\phi : \mathcal{J} \to \mathbb{C}$ by

$$\hat{F}_\phi(\vec{k}) = \sum_{\vec{x}\in\mathcal{J}} F_\phi(\vec{x})e^{-2\pi i\vec{k}\cdot\vec{x}/4n}, \quad \text{for } \vec{k} \in \mathcal{J},$$

denote the discrete Fourier transform (DFT) of F_ϕ. Similarly, we can consider S_3 as a function on $\mathcal{J}\times\mathcal{J}$ and define its DFT $\hat{S}_3 : \mathcal{J} \times \mathcal{J} \to \mathbb{C}$ by

$$\hat{S}_3(\vec{k}_1,\vec{k}_2) := \sum_{\vec{x}_1,\vec{x}_2\in\mathcal{J}} S_3(\vec{x}_1,\vec{x}_2)e^{-2\pi i(\vec{k}_1\cdot\vec{x}_1+\vec{k}_2\cdot\vec{x}_2)/4n}, \quad (9)$$

for $\vec{k}_1, \vec{k}_2 \in \mathcal{J}$. By substituting (2) into (9), we have

$$\hat{S}_3(\vec{k}_1,\vec{k}_2) = \int_0^{2\pi} \hat{F}_\phi(\vec{k}_1)\hat{F}_\phi(\vec{k}_2)\hat{F}_\phi(-\vec{k}_1-\vec{k}_2)d\phi. \quad (10)$$

The triple product in (10) is the Fourier transform of the third-order autocorrelation, and is called the bispectrum [21]; many of its analytical and computational properties have been studied; see for instance [5, 18]. The bispectrum can be generalized to more involved operations, such as 2-D rotations [14], 3-D rotations [12], and general compact groups [10].

Define $\Psi_{\nu,q} : \mathcal{J} \to \mathbb{C}$ as the discrete samples of the Dirichlet Laplacian eigenfunctions: $\Psi_{\nu,q}(\vec{x}) = \psi_{\nu,q}(\vec{x}/n)$, for $\vec{x} \in \mathcal{J}$, where $\psi_{\nu,q}$ is considered as a function supported on the unit disc. With this notation, we can express F_ϕ as

$$F_\phi(\vec{x}) = \sum_{(\nu,q):\lambda_{\nu,q}\le\lambda} \alpha_{\nu,q}\Psi_{\nu,q}(\vec{x})e^{i\nu\phi}$$

If $\nu_{\max} := \max\{\nu : \lambda_{\nu,1} \le \lambda\}$, then the products $F_\phi(\vec{x})F_\phi(\vec{x}+\vec{x}_1)F_\phi(\vec{x}+\vec{x}_2)$ that appear in (2) are bandlimited by $3\nu_{\max}$ with respect to ϕ, and the integral over ϕ in (10) can be replaced by a summation over angles at the Nyquist rate:

$$\hat{S}_3(\vec{k}_1,\vec{k}_2) = \sum_{\nu=0}^{6\nu_{\max}-1} \hat{F}_{\phi_\nu}(\vec{k}_1)\hat{F}_{\phi_\nu}(\vec{k}_2)\hat{F}_{\phi_\nu}(-\vec{k}_1-\vec{k}_2),$$

where $\phi_\nu := 2\pi\nu/(6\nu_{\max})$. By linearity of the DFT, we have

$$\hat{F}_\phi(\vec{k}) = \sum_{(\nu,q):\lambda_{\nu,q}\le\lambda} \alpha_{\nu,q}\hat{\Psi}_{\nu,q}(\vec{k})e^{i\nu\phi},$$

where $\hat{\Psi}_{\nu,q} : \mathcal{J} \to \mathbb{C}$ denotes the DFT of $\Psi_{\nu,q} : \mathcal{J} \to \mathbb{C}$.

Let $\mathcal{V}$ denote the set of all the pairs (ν, q) in the expansion above. We define the column vector $w_{\vec{k},\phi} \in \mathbb{C}^{\mathcal{V}}$ by $(w_{\vec{k},\phi})_{\nu,q} = \hat{\Psi}_{\nu,q}(\vec{k})e^{i\nu\phi}$, and the column vector $z \in \mathbb{C}^{\mathcal{V}}$ by $(z)_{\nu,q} = \alpha_{\nu,q}$; the latter encodes the parameters that describe the target image. With this notation, $\hat{F}_\phi(\vec{k})$ can be expressed compactly as $\hat{F}_\phi(\vec{k}) = z^\top w_{\vec{k},\phi}$, and it follows that

$$\hat{S}_3^z(\vec{k}_1,\vec{k}_2) = \sum_{\nu=0}^{6\nu_{\max}-1} \left(z^\top w_{\vec{k}_1,\phi_\nu}\right)\left(z^\top w_{\vec{k}_2,\phi_\nu}\right)\left(z^\top w_{-\vec{k}_1-\vec{k}_2,\phi_\nu}\right), \quad (11)$$

where we write $\hat{S}_3^z$ to emphasize the dependence on z. This expression for $\hat{S}_3^z$ is particularly convenient for computational purposes. In particular, the gradient of $\hat{S}_3^z(\vec{k}_1,\vec{k}_2)$ with respect to z is easy to compute, which is important for solving the optimization problem defined in Section 3.3.

3.2. Leveraging symmetries

Recall that $S_3(\vec{x}_1,\vec{x}_2)$ is a discrete version of $s_3(\vec{x}_1,\vec{x}_2)$, which only depends on the three parameters: $|\vec{x}_1|$, $|\vec{x}_2|$ and $\theta(\vec{x}_1,\vec{x}_2)$. Let S_3^* be the estimate of S_3 from the noisy measurement M via (5), that is, the de-biased and normalized autocorrelation of the measurement A_3 (3), and $\hat{S}_3^*$ denote the DFT of S_3^*. Since the noise is assumed to be i.i.d. Gaussian with mean 0, its Fourier transform is also i.i.d. Gaussian with mean 0 in Fourier space. Therefore, we can reduce the effect of noise by binning the entries of $\hat{S}_3^*$ with similar values of $|\vec{k}_1|$, $|\vec{k}_2|$ and $\theta(\vec{k}_1,\vec{k}_2)$. Numerically, the binning is done through the mapping $B : \mathcal{J} \times \mathcal{J} \to \mathcal{I}$ by

$$B(\vec{k}_1,\vec{k}_2) = \left(\left\lfloor b_1|\vec{k}_1|\right\rfloor, \left\lfloor b_1|\vec{k}_2|\right\rfloor, \left\lfloor b_2\theta(\vec{k}_1,\vec{k}_2)\right\rfloor\right),$$

where $\mathcal{I} \subset \mathbb{Z}^3_{\ge 0}$ represents the set of the bins, and $b_1, b_2 \in \mathbb{R}$ determine the bin sizes.

3.3. Optimization problem

Suppose that we are given an estimate of $\hat{S}_3^*$—the de-biased and normalized third-order autocorrelation of the measurement. For a fixed model and any coefficient vector $z \in \mathcal{C}^{\mathcal{V}}$, we can compute $\hat{S}_3^z$ via (11). In order to estimate z that is consistent with $\hat{S}_3^*$, we define the cost function

$$f(z) := \sum_{\vec{j} \in \mathcal{I}} \left(\sum_{(\vec{k}_1, \vec{k}_2) : B(\vec{k}_1, \vec{k}_2) = \vec{j}} \hat{S}_3^z(\vec{k}_1, \vec{k}_2) - \hat{S}_3^*(\vec{k}_1, \vec{k}_2) \right)^2. \quad (12)$$

Minimizing this cost function is a non-convex (polynomial of degree 6) least squares optimization problem, and thus, a priori, there is no reason to suspect that the global minimum of this problem can be attained; however, our numerical results indicate that standard gradient-based methods result in accurate and stable recovery of the parameters z. We note that each iteration in the optimization costs $\mathcal{O}(n^4 \nu_{\max})$ operations.

4. NUMERICAL RESULTS

We first consider the problem of recovering a model image from a noiseless set of $\hat{S}_3^*$—namely, from the exact rotationally-averaged third-order autocorrelation. This reflects the case where the noise level is fixed and $m \to \infty$. The model image is generated by expanding a 65×65 image of a tiger (Figure 2) into a linear combination of the first 600 Dirichlet Laplacian eigenfunctions, sorted by the eigenvalues in ascending order. We denote this model image by F_0. Since $\hat{S}_3^*$ is noiseless in this case, we accelerate the optimization by only choosing one entry from each bin in $\mathcal{I}$ instead of summing all the entries to calculate the cost and gradient, which reduces the computational cost by $\mathcal{O}(n)$.

Let F_z be the image formed by the coefficients z that were recovered by minimizing the least squares (12). Since $\hat{S}_3^*$ is invariant under in-plane rotations of the model image, we only expect to reconstruct the image F_0 up to some arbitrary rotation ϕ. As a result, we define the reconstruction error by

$$\text{error}_{\text{recon}} := \inf_{\phi \in [0, 2\pi)} \frac{\|F_0 - F_z^\phi\|_2}{\|F_0\|_2},$$

where F_z^ϕ is the rotation of F_z by an angle ϕ. Using the BFGS optimization algorithm, we recovered the image in Figure 2 with $\text{error}_{\text{recon}} = 5 \times 10^{-12}$. The optimization took 6.5×10^4 seconds parallelized over 100 CPUs in total.

In the second experiment, we study the robustness of the reconstruction to noise. Limited by the computational resources, we downsample the tiger image to 17×17 pixels, and expand it in the first 100 Dirichlet Laplacian eigenfunctions. From the expansion we compute S_3, which is later contaminate by i.i.d. Gaussian noise with mean 0 and variance σ^2. A total of 10 different values of σ are used to model the noisy

Fig. 2. A 65×65 image of a tiger expanded in the first 600 eigenfunctions, which can be recovered from a noiseless set of $\hat{S}_3^*$ to high precision.

counterparts S_3^*. The relative error of S_3^* is quantified by

$$\text{error}_{S_3^*} := \frac{\|S_3 - S_3^*\|_2}{\|S_3\|_2}.$$

The relative errors of the reconstructed images from the 10 sets of $\hat{S}_3^*$ are shown in Figure 3. We observe that the images can be reliably recovered over a wide range of noise levels. The high correlation between $\text{error}_{S_3^*}$ and $\text{error}_{\text{recon}}$ might indicate that the optimization landscape is benign.

$\text{error}_{S_3^*}$	$\text{error}_{\text{recon}}$
1.9×10^{-4}	1.6×10^{-3}
3.9×10^{-4}	3.8×10^{-3}
7.7×10^{-4}	4.9×10^{-3}
1.5×10^{-3}	1.0×10^{-2}
3.1×10^{-3}	2.3×10^{-2}
6.2×10^{-3}	4.5×10^{-2}
1.2×10^{-2}	1.1×10^{-1}
2.4×10^{-2}	1.5×10^{-1}
5.0×10^{-2}	3.2×10^{-1}
9.9×10^{-2}	6.9×10^{-1}

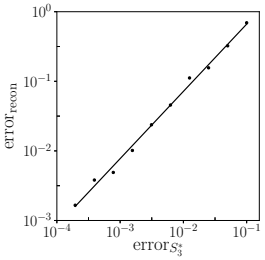

Fig. 3. We consider S_3^* with 10 different levels of Gaussian noise quantified by $\text{error}_{S_3^*}$. For each level of noise, we run the optimization five times with random initializations; we report $\text{error}_{\text{recon}}$ corresponding to the trial that minimizes the cost function (12). The slope of the best-fit line is 0.97.

5. DISCUSSION

This work serves as a proof of concept for the feasibility of estimating a target image from its rotational and translational invariants. We have demonstrated the reconstruction of a target image from its noiseless invariants, and showed that our algorithmic approach is robust to noise. In future work, we intend to extend the framework to include the recovery of the target image from the observation M, to mitigate the separation condition [13], and to allow the recovery of multiple images simultaneously, in a similar fashion to [4, 6]. From a theoretical perspective, we wish to complement the empirical results of this work by proving that indeed a generic image f is determined uniquely by its rotationally-averaged third-order autocorrelation.

6. REFERENCES

[1] C. Aguerrebere, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Transactions on Signal Processing*, vol. 64, no. 21, pp. 5707–5722, 2016.

[2] T. Bendory, A. Bartesaghi, and A. Singer, "Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities", *arXiv preprint arXiv:1908.00574*, 2019.

[3] T. Bendory, N. Boumal, W. Leeb, E. Levin, and A. Singer, "Toward single particle reconstruction without particle picking: Breaking the detection limit," *arXiv preprint arXiv:1810.00226*, 2018.

[4] T. Bendory, N. Boumal, W. Leeb, E. Levin, and A. Singer, "Multi-target detection with application to cryo-electron microscopy," *Inverse Problems*, 2019.

[5] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer, "Bispectrum inversion with application to multireference alignment," *IEEE Transactions on Signal Processing*, vol. 66, no. 4, pp. 1037–1050, 2017.

[6] N. Boumal, T. Bendory, R. R. Lederman, and A. Singer, "Heterogeneous multireference alignment: A single pass approach," in *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*. pp. 1–6, IEEE, 2018.

[7] J. Frank, *Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state*. Oxford University Press, 2006.

[8] T. Grant, A. Rohou, and N. Grigorieff, "cisTEM, user-friendly software for single-particle image processing," *Elife*, vol. 7, p. e35383, 2018.

[9] R. Henderson, "The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological molecules," *Quarterly reviews of biophysics*, vol. 28, no. 2, pp. 171–193, 1995.

[10] R. Kakarala, "The bispectrum as a source of phase-sensitive invariants for Fourier descriptors: a group-theoretic approach," *Journal of Mathematical Imaging and Vision*, vol. 44, no. 3, pp. 341–353, 2012.

[11] Z. Kam, "The reconstruction of structure from electron micrographs of randomly oriented particles," *Journal of Theoretical Biology*, vol. 82, no. 1, pp. 15–39, 1980.

[12] R. Kondor, "A novel set of rotationally and translationally invariant features for images based on the non-commutative bispectrum," *arXiv preprint arXiv:0701127*, 2007.

[13] T.-Y. Lan, T. Bendory, N. Boumal, and A. Singer, "Multi-target detection with an arbitrary spacing distribution," *arXiv preprint arXiv:1905.03176*, 2019.

[14] C. Ma, T. Bendory, N. Boumal, F. Sigworth, and A. Singer, "Heterogeneous multireference alignment for images with application to 2-D classification in single particle reconstruction," *To appear in IEEE Transactions on Image Processing*, 2019.

[15] S. Mallat, "Group invariant scattering", *Communications on Pure and Applied Mathematics* vol. 65, no. 10, pp. 1331–1398, 2012.

[16] O. Özyeşil, N. Sharon, and A. Singer, "Synchronization over cartan motion groups via contraction", *SIAM Journal on Applied Algebra and Geometry*, vol. 2, no. 2, pp. 207–241, 2018.

[17] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, "cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination," *Nature methods*, vol. 14, no. 3, p. 290, 2017.

[18] B. M. Sadler and G. B. Giannakis, "Shift-and rotation-invariant object reconstruction using the bispectrum," *JOSA A*, vol. 9, no. 1, pp. 57–69, 1992.

[19] S. H. Scheres, "RELION: implementation of a Bayesian approach to cryo-EM structure determination," *Journal of structural biology*, vol. 180, no. 3, pp. 519–530, 2012.

[20] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, "EMAN2: an extensible image processing suite for electron microscopy," *Journal of structural biology*, vol. 157, no. 1, pp. 38–46, 2007.

[21] J. Tukey, "The spectral representation and transformation properties of the higher moments of stationary time series," *Reprinted in The Collected Works of John W. Tukey*, vol. 1, pp. 165–184, 1953.