Provable algorithms for multi-reference alignment over $\mathsf{SO}(2)$

Gil Drozatz, Tamir Bendory, *Senior Member, IEEE,* and Nir Sharon

***Abstract*—The multi-reference alignment (MRA) problem involves reconstructing a signal from multiple noisy observations, each transformed by a random group element. In this paper, we focus on the group $\mathrm{SO}(2)$ of in-plane rotations and propose two computationally efficient algorithms with theoretical guarantees for accurate signal recovery under a non-uniform distribution over the group. The first algorithm exploits the spectral properties of the second moment of the data, while the second utilizes the frequency marching principle. Both algorithms achieve the optimal estimation rate in high-noise regimes, marking a significant advancement in the development of computationally efficient and statistically optimal methods for estimation problems over groups.**

***Index Terms*—Multi-reference alignment, spectral algorithm, frequency marching, method of moments**

I. Introduction

Let G be the group of 2-D rotations, $\mathsf{SO}(2)$, acting on a finite-dimensional vector space V. We aim to recover a signal $x \in V$ from n observations $y_1, \dots, y_n$ of the form:

$$y_i = g_i \cdot x + \varepsilon_i, \quad i = 1, \dots, n, \tag{I.1}$$

where $\cdot$ denotes the group action, $g_i \in G$ are unknown random group elements, and ε_i is a noise term. We assume that the group elements are drawn from a nonuniform distribution over $\mathsf{SO}(2)$. This paper studies the action of the group $\mathsf{SO}(2)$ on two vector spaces: bandlimited 1-D signals and bandlimited 2-D images, as detailed in Sections II and III. We also assume that $\varepsilon_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2 I)$ over the vector of coefficients in the appropriate basis of V (as detailed later).

The studied model is a special case of the multi-reference alignment (MRA) problem, in which $\mathsf{SO}(2)$ is replaced by other compact groups acting on finite-dimensional vector spaces, as introduced in [2], [3], [4], [5]. The primary motivation for studying the MRA model is the transformative technology of single-particle cryo-electron microscopy (cryo-EM) to elucidate the spatial structure of biological molecules [6], [7], [8]. Since the observations are invariant under an intrinsic group of symmetries, it is impossible to distinguish x from $g \cdot x$ for any fixed $g \in G$. Thus, the goal is to estimate the G-orbit of the signal: $\{g \cdot x \mid g \in G\}$.

G. Drozatz and T. Bendory are with the School of Electrical and Computer Engineering, Tel Aviv University, Israel. N. Sharon is with the Department of Applied Mathematics, Tel Aviv University, Israel. The research was supported in part by NSF-BSF under Grant 2019752. T.B. is also supported in part by BSF under Grant 2020159, 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). N.S. is also partially supported by the DFG award 514588180.

To estimate the signal, we propose a two-stage framework based on the classical method of moments. First, we approximate the first two population moments using the empirical moments of the observations. This approximation is accurate if the number of observations is sufficiently large $n \gg \sigma^4$. Next, we aim to determine the parameters that represent the signal based on the estimated moments. Previous works have demonstrated that in the high-noise regime, the sample complexity of the MRA problem is governed by the lowest-order moment that uniquely determines the signal's orbit [9], [10]. Since recovery from the first moment is impossible, recovery from the second moment implies that the sample complexity of the model in (I.1) is proportional to σ^4.

The main computational challenge in the method of moments is estimating the signal from the moments. To address this, we develop two methods. The first method employs a frequency marching approach, extending a similar technique used in discrete settings [3]. We prove that exact recovery is possible from the first two population moments, which implies the sample complexity of the model. This proof is constructive, as it introduces an explicit computationally efficient algorithm. The second method utilizes the spectral decomposition of an approximation of the second-moment matrix. This algorithm generalizes the spectral method used in discrete 1-D MRA [9], and we refer to it as the spectral algorithm. This algorithm *approximates* the solution based on the spectral properties of the second moment matrix; see Theorem II.4. The algorithms are detailed in Section II and Section III for the 1-D and 2-D cases, respectively. We support these findings with numerical experiments in Section IV. Due to space constraints, the proofs for the claims presented in this paper can be found in the full version [1].

Contribution. The uniqueness of recovering a signal from its MRA moments has been extensively studied in recent years (see, for example, [4], [5], [10]). While most of these works focus on uniform distributions over the group, non-uniform distributions have also been explored [9], [11], [12]. However, computationally efficient provable algorithms have been developed only for the simplest case, where the group of circular shifts acts on 1-D discrete signals [3], [9], [10]. This paper introduces the first provable algorithms for a continuous group $\mathsf{SO}(2)$ that acts on signals and images, marking a significant milestone in the development of provable algorithms for more complex MRA models. In Section V, we discuss the potential implications for cryo-EM. This is particularly crucial since existing algorithms in cryo-EM rely on heuristics without robust validation measures.

II. 1-D MRA over SO(2)

We first consider a vector space V of 1-D bandlimited signals. Let $x(\theta) = \sum_{k=-B}^{B} \hat{x}[k] e^{\iota k\theta}$, be a B-band-limited signal on the circle $\theta \in [0, 2\pi)$, where $\iota = \sqrt{-1}$. Here, $\hat{x} \in \mathbb{C}^{2B+1}$ denotes the Fourier coefficients of x. In Fourier space, the observations are given by

$$\hat{y}_i[k] = g_i \cdot \hat{x}[k] + \hat{\varepsilon}_i[k], \quad i = 1, \ldots, n, \tag{II.1}$$

where $\hat{y}[k]$ is the k-th Fourier coefficient of the i-th observation, $g \cdot \hat{x}[k] = \hat{x}[k] e^{-\iota k\phi}$, where ϕ is the angle associated with g, $\hat{\varepsilon}_i[0] \sim \mathcal{N}(0, \sigma^2)$, and for $k \geq 0$ the real and imaginary part of $\hat{\varepsilon}_i[k]$ are i.i.d. normal variables with zero mean and variance $\sigma^2/2$ that obey the conjugation rule $\hat{\varepsilon}_i[k] = \hat{\varepsilon}_i^*[-k]$. As we show next, the high frequencies of the distribution are annihilated by the moments, and thus, we assume, without loss of generality, that the Fourier coefficients of the distribution of group elements, $\hat{\rho}$, is 2B-bandlimited.

Recall that the first two population moments are given by

$$\begin{aligned} M_1 &= \mathbb{E}\left[\hat{y}_i\right] \in \mathbb{C}^{(2B+1)}, \\ M_2 &= \mathbb{E}\left[\hat{y}_i \hat{y}_i^*\right] \in \mathbb{C}^{(2B+1)^2}. \end{aligned} \tag{II.2}$$

The population moments can be estimated in practice by the empirical moments of the observations by averaging over the observable moments

$$\begin{aligned} M_{1,\text{est}} &\triangleq \frac{1}{n} \sum_{i=1}^{n} \hat{y}_i, \\ M_{2,\text{est}} &\triangleq \frac{1}{n} \sum_{i=1}^{n} \hat{y}_i \hat{y}_i^*. \end{aligned} \tag{II.3}$$

The population moments can be accurately estimated when $n/\sigma^4 \to \infty$. This is our assumption for the rest of the paper unless stated otherwise.

Let $T \in \mathbb{C}^{(2B+1)\times(2B+1)}$ be a Toeplitz matrix with elements defined by $T_\rho[k_1, k_2] \triangleq \hat{\rho}[k_1 - k_2]$, $D_{\hat{x}}$ is the diagonal matrix of $\hat{x}$, $\mathbb{I}$ is the identity matrix, and $\odot$ represents the Hadamard product. The moments are given explicitly in the following lemma.

Lemma II.1. *Consider the model* (I.1)*. Then,*

$$\begin{aligned} M_1 &= 2\pi \hat{x} \odot \hat{\rho}, \\ M_2 &= 2\pi D_{\hat{x}} T_\rho D_{\hat{x}}^* + \sigma^2 \mathbb{I}_{2B+1}. \end{aligned} \tag{II.4}$$

A. Frequency Marching Algorithm

We begin by studying a frequency marching algorithm that successively recovers the high frequencies based on the low frequencies. In particular, we show that given the population moments M_1 and M_2, this algorithm recovers the signal and the distribution ρ uniquely. The underlying idea of the algorithm is to reformulate M_2 as a function solely of $\hat{\rho}$, using the information of the first moment M_1. This allows iteratively recovering $\hat{\rho}$. Then, $\hat{x}$ is recovered from M_1 and $\hat{\rho}$. By saying that $\hat{x}$ is *non-vanishing*, we mean that all the Fourier coefficients are non-zero.

Algorithm 1 A frequency marching algorithm for the 1-D model

Input: $M_{1,\text{est}}$, $M_{2,\text{est}}$, and σ
Output: $\hat{x}_{\text{est}}$, $\hat{\rho}_{\text{est}}$

1) $M_{2,\text{est}} \leftarrow M_{2,\text{est}} - \sigma^2 \mathbb{I}_{2B+1}$ (debiasing)
2) $S \leftarrow 2\pi D_{M_{1,\text{est}}}^{-1} M_{2,\text{est}} D_{M_{1,\text{est}}^*}^{-1}$
3) $\hat{\rho}_{\text{est}}[0] \leftarrow \frac{1}{2\pi}$
4) $\hat{\rho}_{\text{est}}[1] \leftarrow \sqrt{1/2\pi S[1,1]}$
5) **for** $2 \leq k \leq B$ **do**
6) $\hat{\rho}_{\text{est}}[k] \leftarrow \frac{\hat{\rho}_{\text{est}}[1]}{S[k,k-1]\hat{\rho}_{\text{est}}^*[k-1]}$
7) **end for**
8) **for** $B+1 \leq k \leq 2B$ **do**
9) $\hat{\rho}_{\text{est}}[k] \leftarrow S[k-B, -B]\hat{\rho}_{\text{est}}[k-B]\hat{\rho}_{\text{est}}[B]$
10) **end for**
11) **for** $-2B \leq k \leq -1$ **do**
12) $\hat{\rho}_{\text{est}}[k] = \hat{\rho}_{\text{est}}^*[-k]$
13) **end for**
14) $\hat{x}_{\text{est}} \leftarrow \frac{M_{1,\text{est}}}{2\pi\hat{\rho}_{\text{est}}}$

Proposition II.2. *Assume that $\hat{x}$ and $\hat{\rho}$ are non-vanishing and $M_{1,\text{est}} = M_1$, and $M_{2,\text{est}} = M_2$, Then, Algorithm 1 recovers $\hat{x}$ and $\hat{\rho}$ exactly, up to a global rotation.*

Algorithm 1 relies on a single diagonal of M_2, ignoring others that carry valuable information. Utilizing additional diagonals could improve robustness and performance. In Section IV, we discuss the estimation error caused by using non-exact moments.

B. Spectral Algorithm

The second algorithm uses the spectral properties of M_2 and the close relation between circulant and Toeplitz matrices. Let $P_{\hat{x}} = |\hat{x}|^2$ be the power spectrum of $\hat{x}$, which is the diagonal of the second-moment matrix. The algorithm begins by conjugating (normalizing) the second-moment matrix by $D_{\frac{1}{\sqrt{P_{\hat{x}}}}}$ and then extracting an isolated eigenvector that contains an approximation of the Fourier phases of the signal. This eigenvector is then combined with its Fourier magnitudes. Note that the algorithm presented here is for a continuous case, generalizing a similar approach in an earlier discrete case [9].

The crux of the spectral algorithm is the similarity between Toeplitz and circulant matrices and the projection of a given Hermitian Toeplitz matrix into the space of circulant matrices, which are diagonalized by the DFT matrix. Consider the problem:

$$\underset{v \in \mathbb{C}^{2B+1}}{\operatorname{argmin}} \left\| T_\rho - C_v \right\|_F^2, \tag{II.5}$$

where v is a vector defining the circulant matrix C_v. This problem enjoys a closed-form solution [13], given by

$$v_{\text{opt}}[k] = \frac{k\hat{\rho}\left[-(2B+1-k)\right] + (2B+1-k)\hat{\rho}[k]}{2B+1}, \tag{II.6}$$

and the distance between T_ρ and the associated circulant matrix (the error) is given by

$$S_B(\hat{\rho}) = \sum_{k=1}^{2B} \left|\hat{\rho}[k] - \hat{\rho}[-(2B+1-k)]\right|^2 \left(\frac{k(2B+1-k)}{2B+1}\right). \tag{II.7}$$

Algorithm 2 A Spectral Algorithm for the 1-D model

Input: $M_{1,\text{est}}$, $M_{2,\text{est}}$, and σ
Output: $\hat{x}_{\text{est}}$, $\hat{\rho}_{\text{est}}$

1) $M_{2,\text{est}} \leftarrow M_{2,\text{est}} - \sigma^2 \mathbb{I}_{2B+1}$ (debiasing)
2) $P_{\hat{x}} \leftarrow \text{diag}(M_{2,\text{est}})$
3) $\mathcal{M}_2 \leftarrow D_{\frac{1}{\sqrt{P_{\hat{x}}}}} M_{2,\text{est}} D_{\frac{1}{\sqrt{P_{\hat{x}}}}}$ (conjugating)
4) Find the eigenvalue decomposition of $\mathcal{M}_2$: eigenvalues $\lambda_0 \geq \lambda_1 \geq \ldots \geq \lambda_{2B}$ with corresponding eigenvectors $v_0, v_1, \ldots, v_{2B}$.
5) $\kappa \leftarrow \text{argmax}_{0 \leq k \leq 2B} \min_{k' \neq k} |\lambda_{k'} - \lambda_k|$
6) $\tilde{x}_{\text{est}} \leftarrow \sqrt{2B+1} v_\kappa$
7) $\beta \leftarrow e^{\iota(\angle M_{1,\text{est}}[0] - \angle \tilde{x}_{\text{est}}[0])}$
8) $\tilde{x}_{\text{est}} \leftarrow \beta \tilde{x}_{\text{est}}$
9) $\hat{x}_{\text{est}} \leftarrow \sqrt{P_{\hat{x}}} \odot \tilde{x}_{\text{est}}$
10) $\hat{\rho}_{\text{est}} \leftarrow \frac{M_{1,\text{est}}}{2\pi \hat{x}}$

The following theorem shows that given the ideal assumption of T_ρ being a circulant matrix, Algorithm 2 recovers the signal and the first B Fourier coefficients of ρ.

Proposition II.3. *Assume that $S_B(\hat{\rho}) = 0$, $\hat{x}$ is non-vanishing and that $\hat{v}_{opt}$ has an isolated entry. In addition, assume that $M_{1,\text{est}} = M_1$, $M_{2,\text{est}} = M_2$. Then, Algorithm 2 recovers $\hat{x}$ and the first B Fourier coefficients of ρ exactly, up to a global rotation.*

In practice, we do not expect T_ρ to be a circulant matrix, implying $S_B(\hat{\rho}) > 0$. Hence, in practice, the spectral algorithm only approximates the solution, and the estimation error depends on $S_B(\hat{\rho})$. This is captured by the following result, based on the Davis-Kahan Theorem [14].

Theorem II.4. *Consider $\lambda_0^T \geq \lambda_1^T \geq \ldots \geq \lambda_{2B}^T$ and $\lambda_0^C \geq \lambda_1^C \geq \ldots \geq \lambda_{2B}^C$ to be the eigenvalues of T_ρ and $C_{v_{opt}}$, respectively. We denote $P_{\max} := \max_{0 \leq k \leq B} P_{\hat{x}}[k]$ and $\delta_\kappa = \max\left(\min_{j \neq \kappa} \left|\lambda_\kappa^C - \lambda_j^T\right|, \min_{j \neq \kappa} \left|\lambda_j^C - \lambda_\kappa^T\right|\right)$. Assume the following conditions hold:*

1) *$\hat{x}$ is non-vanishing.*
2) *κ of Step 5 of Algorithm 2 has both λ_κ^T and λ_κ^C with an eigenspace of dimension 1.*
3) *The estimate calculated in Step 8 of Algorithm 2 satisfies $\tilde{x}_{\text{est}}^*[\Phi_{\frac{2\pi l}{2B+1}} \odot \tilde{x}] \geq 0$.*
4) *$S_B(\hat{\rho}) \leq \delta_\kappa^2$.*

Then, there exists $0 \leq l \leq 2B$, such that

$$\left\| \hat{x}_{\text{est}} - \Phi_{\frac{2\pi l}{2B+1}} \odot \hat{x} \right\|_F^2 \leq 2(2B+1) P_{\max} \left[1 - \sqrt{1 - \frac{S_B(\hat{\rho})}{\delta_\kappa^2}} \right], \tag{II.8}$$

where $\Phi_\varphi[k] = e^{-\iota k \varphi}$.

III. 2-D MRA over SO(2)

We now extend the algorithms from 1-D signals to 2-D images. We assume that V is the space of bandlimited images in the sense that they can be presented with finitely many Fourier-Bessel coefficients [15]. Namely, the sought image is of the form

$$x(\theta, r) = \sum_{(k,q) \in I} \hat{x}[k,q] e^{\iota k \theta} J_q(r), \quad \theta \in [0, 2\pi), \quad r \in [0,1],$$

where $u_{k,q}(\theta, r) = e^{\iota k \theta} J_q(r)$ is the general Fourier-Bessel function, with $J_q(r)$ being the cylindrical Bessel function of order q. The set $\{u_{k,q}(\theta, r)\}_{k \in \mathbb{Z}, q \in \mathbb{Z}_{\geq 0}}$ spans all the "nice enough" functions over the unit disc. The images are bandlimited in the sense that there exists an angular bandwidth B and a maximal radial frequency Q_k such that the set $I = \left\{(k,q) : -B \leq k \leq B,\, 0 \leq q \leq Q_k - 1\right\}$ is finite. We denote the total number of coefficients by $|I|$. It can be readily seen that

$$x(\theta - \varphi, r) = \sum_{(k,q) \in I} \hat{x}[k,q] e^{-\iota k \varphi} u_{k,q}(\theta, r). \tag{III.1}$$

Thus, in Fourier-Bessel space, the observations are given by

$$\hat{y}_i[k,q] = g_i \cdot \hat{x}[k,q] + \hat{\varepsilon}_i[k,q], \quad i = 1, \ldots, n, \tag{III.2}$$

where $\hat{y}[k,q]$ is the (k,q)-th Fourier-Bessel coefficient of the i-th observation, $g \cdot \hat{x}[k,q] = \hat{x}[k,q] e^{-\iota k \phi}$, where ϕ is the angle associated with g, $\hat{\varepsilon}_i[0,q] \sim \mathcal{N}(0, \sigma^2)$, and for $k \geq 0$ the real and imaginary part of $\hat{\varepsilon}_i[k,q]$ are i.i.d. normal variables with zero mean and variance $\sigma^2/2$ that obey the conjugation $\hat{\varepsilon}_i[k,q] = \hat{\varepsilon}_i^*[-k,q]$. In this sequel, we order the Fourier-Bessel coefficients in lexicographical order and think of them as a column vector with $|I|$ complex entries.

The following lemma gives the first two moments.

Lemma III.1. *Consider the 2-D MRA problem* (III.2)*. The first two moments are given by*

$$M_1 = 2\pi \hat{x} \odot \mathcal{R}, \tag{III.3}$$

$$M_2 = 2\pi D_{\hat{x}} \mathcal{T}_\rho D_{\hat{x}^*} + \sigma^2 \mathbb{I}_{|I|}. \tag{III.4}$$

Here, $\mathcal{R} \in \mathbb{C}^{|I|}$ and $\mathcal{T}_\rho \in \mathbb{C}^{|I| \times |I|}$ are block matrices with blocks indexed by $-B \leq k_1, k_2 \leq B$ and are defined as:

$$(\mathcal{T}_\rho)_{k_1 k_2} = \hat{\rho}[k_1 - k_2] \mathbb{1}_{Q_{|k_1|} \times Q_{|k_2|}},$$
$$(\mathcal{R})_{k_1} = \hat{\rho}[k_1] \mathbb{1}_{Q_{|k_1|} \times 1},$$

where $\mathbb{1}_{i \times j}$ denotes a matrix of ones of size $i \times j$.

A. Frequency Marching Algorithm

We consider a frequency marching algorithm, similar to Algorithm 1. The difference stems from the addition of the radial coordinate that provides more information and, thus, stability, as shown in a related problem in [18].

Proposition III.2. *Assuming $M_{1,est} = M_1$, $M_{2,est} = M_2$, there exists a frequency marching algorithm, analog to Algorithm 1, that recovers $\hat{x}$ and $\hat{\rho}$ exactly, up to a global rotation.*

B. Spectral Algorithm

As in the 1-D case, a spectral algorithm can be devised for the 2-D setting, analogous to Algorithm 2, with Toeplitz and circulant matrices replaced by their block counterparts. If the block Toeplitz matrix is block circulant, the algorithm recovers the signal (up to a global rotation) from second-moment data. Otherwise, the error scales with its distance from the nearest block circulant matrix.

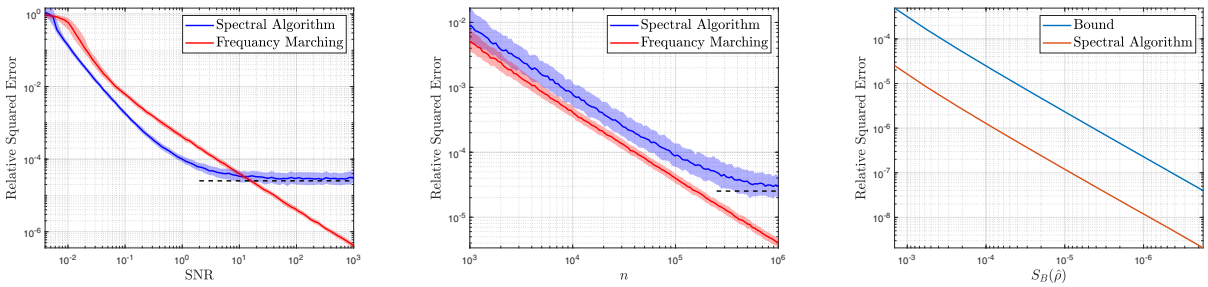

(a) (b) (c)

Fig. 1: (a) Recovery error of the frequency marching and spectral algorithms, with 20% error margins, as a function of the SNR for $n = 10^6$ observations. (b) Recovery error of the frequency marching and spectral algorithms, with 20% error margins, as a function of the number of observations for SNR $= 100$. (c) Recovery error of the spectral algorithm as a function of $S_B(\hat{\rho})$, compared to the theoretical bound from Theorem II.4.

IV. Numerical Results

We present the numerical results of both algorithms for the 2-D case. The performance for the 1-D case is similar, but due to space constraints, it is omitted. In all experiments, the ground truth image was generated with $B = 10$ and $Q = 2$. The magnitudes of all Fourier–Bessel coefficients were set to one, and their phases were drawn from a uniform distribution, subject to the constraint that the image remains real. The real and imaginary components of the distribution were initially drawn from a uniform distribution over the interval $[0, 1]$ and subsequently corrected according to Equation (II.6) to ensure $S_B(\hat{\rho}) = 0$ while preserving both the positivity and normalization of the distribution in the signal domain (as it is a distribution). Next, we perturbed the distribution as

$$\hat{\rho}[k] \to e^{\iota\eta\sqrt{k}}\hat{\rho}[k], \tag{IV.1}$$

with $\eta = 0.1$. In the experiments presented in Figure 1a and 1b, this perturbation resulted in $S_B(\hat{\rho}) = 0.0014$. The observations $\hat{y}_i$ are then generated, with Gaussian noise added according to Equation (III.2). We define the recovery error as

$$\frac{\min_{\varphi\in[0,2\pi)}\left\|\hat{x}_{\text{est}} - e^{-\iota k\varphi}\odot\hat{x}\right\|_F^2}{\|\hat{x}\|_F^2}. \tag{IV.2}$$

In the experiment whose results are shown in Figure 1a, we present the recovery error as a function of the SNR, defined by

$$\text{SNR} = \frac{\sum_{k,q} P_{\hat{x}}[k,q]}{(2B+1)Q\sigma^2},$$

with $n = 10^6$ observations. We report the median error with 20% error margins, computed over 400 trials per SNR value. As expected, the error decreases as the SNR increases. For higher SNR values, the error of the spectral algorithm stagnates because $S_B(\hat{\rho}) \neq 0$ (see Theorem II.4). In contrast, the error of the frequency marching algorithm approaches zero at high SNR, as predicted by Proposition III.2. The error rate of the frequency marching is proportional to $1/\text{SNR}$. However, at lower SNR values, the spectral algorithm outperforms the frequency marching algorithm. This is because the spectral algorithm processes all second-moment information simultaneously, while the frequency marching algorithm, due to its sequential nature, accumulates errors over frequencies.

The experiment shown in Figure 1b presents the recovery error as a function of the number of observations for SNR $=$ 100. We report the median error over 800 trials with 20% error margins. In this high-SNR regime, the frequency marching algorithm outperforms the spectral algorithm, achieving an error rate proportional to $1/n$.

In the experiment presented in Figure 1c, we assume access to the exact moments and compare the error of the spectral algorithm with the theoretical bound from Theorem II.4. This comparison is performed by varying the parameter η in (IV.1) from 10^{-1} to 10^{-3} and optimizing over all possible rotations of ρ to achieve the minimal bound. The theoretical bound closely follows the numerical results of the spectral algorithm as a function of $S_B(\hat{\rho})$. However, a noticeable gap remains, indicating that Theorem II.4 is not tight.

V. Vision: Provable Algorithms for Cryo-EM

This paper introduces provable frequency marching and spectral algorithms for the MRA problem with the group of in-plane rotations, $\mathrm{SO}(2)$. This represents a significant advancement toward developing rigorous algorithms for more complex MRA models.

Our primary motivation is cryo-EM, which can be modeled as an MRA model with the group of 3-D rotations $\mathsf{SO}(3)$ and an additional linear operator (which is not modeled by (I.1)) [8]. Despite being widely used and occasionally yielding excellent results, existing cryo-EM algorithms lack theoretical guarantees and robust validation measures. Moreover, cryo-EM algorithms are susceptible to well-documented pitfalls and model biases, such as the "Einstein from noise" phenomenon [16], [17]. To address these challenges, our vision is to design provable algorithms for cryo-EM—whether through frequency marching, spectral methods, or alternative techniques—to enhance confidence in cryo-EM reconstructions and ultimately drive new biological discoveries.

References

[1] G. Drozatz, T. Bendory, N. Sharon, "Provable algorithms for multi-reference alignment over $\mathsf{SO}(2)$," *arXiv preprint arXiv:2504.19140*, 2025.

[2] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu, "Multireference alignment using semidefinite programming," in *Proceedings of the 5th conference on Innovations in theoretical computer science*, 2014, pp. 459–470.
[3] 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.
[4] A. S. Bandeira, B. Blum-Smith, J. Kileel, J. Niles-Weed, A. Perry, and A. S. Wein, "Estimation under group actions: recovering orbits from invariants," *Applied and Computational Harmonic Analysis*, vol. 66, pp. 236–319, 2023.
[5] T. Bendory, N. Dym, D. Edidin, and A. 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.
[6] 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.
[7] Y. Cheng, "Single-particle cryo-EM—how did it get here and where will it go," *Science*, vol. 361, no. 6405, pp. 876–880, 2018.
[8] T. Bendory, A. Bartesaghi, and A. Singer, "Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities," *IEEE signal processing magazine*, vol. 37, no. 2, pp. 58–76, 2020.
[9] E. Abbe, T. Bendory, W. Leeb, J. M. Pereira, N. Sharon, and A. Singer, "Multireference alignment is easier with an aperiodic translation distribution," *IEEE Transactions on Information Theory*, vol. 65, no. 6, pp. 3565–3584, 2018.
[10] A. Perry, J. Weed, A. S. Bandeira, P. Rigollet, and A. Singer, "The sample complexity of multireference alignment," *SIAM Journal on Mathematics of Data Science*, vol. 1, no. 3, pp. 497–517, 2019.
[11] T. Bendory, D. Edidin, W. Leeb, and N. Sharon, "Dihedral multi-reference alignment," *IEEE Transactions on Information Theory*, vol. 68, no. 5, pp. 3489–3499, 2022.
[12] N. Sharon, J. Kileel, Y. Khoo, B. Landa, and A. Singer, "Method of moments for 3d single particle ab initio modeling with non-uniform distribution of viewing angles," *Inverse Problems*, vol. 36, no. 4, p. 044003, 2020.
[13] T. F. Chan, "An optimal circulant preconditioner for toeplitz systems," *SIAM Journal on Scientific and Statistical Computing*, vol. 9, no. 4, pp. 766–771, 1988. [Online]. Available: https://doi.org/10.1137/0909051
[14] C. Davis and W. M. Kahan, "The rotation of eigenvectors by a perturbation. iii," *SIAM Journal on Numerical Analysis*, vol. 7, no. 1, pp. 1–46, 1970.
[15] Z. Zhao, Y. Shkolnisky, and A. Singer, "Fast steerable principal component analysis," *IEEE transactions on computational imaging*, vol. 2, no. 1, pp. 1–12, 2016.
[16] R. Henderson, "Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise," *Proceedings of the National Academy of Sciences*, vol. 110, no. 45, pp. 18 037–18 041, 2013.
[17] A. Balanov, W. Huleihel, and T. Bendory, "Einstein from noise: Statistical analysis," *bioRxiv*, pp. 2024–07, 2024.
[18] T. Amir, T. Bendory, N. Dym, and D. Edidin, "The stability of generalized phase retrieval problem over compact groups," *arXiv preprint arXiv:2505.04190*, 2025.