

The sample complexity of multi-reference alignment

AMELIA PERRY*, JONATHAN WEED†, AFONSO S. BANDEIRA‡, PHILIPPE RIGOLLET§ AND AMIT SINGER¶

Massachusetts Institute of Technology

Massachusetts Institute of Technology

Courant Institute of Mathematical Sciences, New York University

Massachusetts Institute of Technology

Princeton University

Abstract. The growing role of data-driven approaches to scientific discovery has unveiled a large class of models that involve latent transformations with a rigid algebraic constraint. Among them, multi-reference alignment (MRA) is a simple model that captures fundamental aspects of the statistical and algorithmic challenges arising from this new paradigm. In this model, an unknown signal is subject to two types of corruption: a latent cyclic shift and the more traditional additive white noise. The goal is to recover the signal at a certain precision from independent samples. While at high signal-to-noise ratio (SNR), the number of observations needed to recover a generic signal is proportional to $1/\text{SNR}$, we show that it rises to $1/\text{SNR}^3$ in the more realistic low SNR regime. We propose an algorithm that achieves this optimal dependence on the SNR. Furthermore, we extend our results to cover heterogeneous MRA model where the samples come from a mixture of signals, as is often the case in applications such as Cryo-Electron Microscopy, where molecules may have different conformations. We provide the first known procedure that provably achieves signal recovery in the low SNR regime for heterogeneous MRA.

AMS 2000 subject classifications: Primary 62F10; secondary 92C55, 16W22.

Key words and phrases: Multi-reference alignment, Method of invariants, Bispectrum, Cryo-EM.

*Supported in part by NSF grant DMS-1541100, NSF CAREER Award CCF-1453261 and a grant from the MIT NEC Corporation.

†Supported in part by NSF Graduate Research Fellowship DGE-1122374.

‡Supported in part by NSF Grant DMS-1317308, NSF DMS-1712730 and NSF DMS-1719545. Part of this work was done while the author was with the Mathematics Department at MIT.

§Supported in part by NSF CAREER DMS-1541099, NSF DMS-1541100, NSF DMS-1712596, DARPA W911NF-16-1-0551, ONR N00014-17-1-2147 and a grant from the MIT NEC Corporation.

¶Supported in part by Award Number R01GM090200 from the NIGMS, FA9550-12-1-0317 from AFOSR, Simons Investigator Award and Simons Collaboration on Algorithms and Geometry from Simons Foundation, and the Moore Foundation Data-Driven Discovery Investigator Award.

1. INTRODUCTION

More and more scientific and engineering disciplines rely on data collection and analysis to guide scientific discovery. This is the case for example in Cryo-Electron Microscopy (Cryo-EM), which was selected by the journal *Nature Methods* as Method of the Year 2015 “for its newfound ability to solve protein structures at near-atomic resolution” [Edi16, Nog16]. This promising imaging method comes with new challenges that resist known signal processing techniques. In Cryo-EM, a molecule is imaged from different, unknown viewing directions, which can be thought of as latent three-dimensional rotations acting on each molecule sample before the image is captured. As we shall see, the presence of the latent rotation has a singular effect not only on the computational but also on the statistical complexity of this problem. This feature is shared by various other examples of the same kind in areas such as robotics [RCBL16], structural biology [Dia92, TS12, PMMC11, PC14]; radar [ZvdHGG03, PZAF05]; crystalline simulations [SSK13]; and image registration problems in a number of important contexts, such as in geology, medicine, and paleontology [DM98, FZB02].

Multi-reference alignment (MRA) [BCSZ14] is one of the simplest models that is able to capture fundamental aspects of this class of problems, rendering it ideal for theoretical study. In this model one observes n independent copies y_1, \dots, y_n of $y \in \mathbb{R}^d$ given by

$$(1) \quad y = R_\ell \theta + \sigma \xi,$$

where R_ℓ is a cyclic shift by an unknown number ℓ of coordinates: the j th coordinate of $R_\ell \theta \in \mathbb{R}^d$ is given by $(R_\ell \theta)_j = \theta_{j+\ell \pmod{d}}$. We assume Gaussian noise $\xi \sim \mathcal{N}(0, I_d)$. The goal is to estimate the unknown vector $\theta \in \mathbb{R}^d$.

The MRA model is illustrated in Figure 1. The unknown vector $\theta \in \mathbb{R}^d$ can be thought of as a discretization of a continuous one-dimensional signal; in each observation y_i , it is subject to two types of corruption: the latent cyclic shift R_{ℓ_i} and the additive Gaussian noise $\sigma \xi_i$.

We refer to $\|\theta\|^2/\sigma^2$ as the signal-to-noise ratio (SNR); without loss of generality we assume in the sequel that $\|\theta\| = 1$, implying $\text{SNR} = 1/\sigma^2$. In practice, the SNR is completely determined by the experimental setup.

In this paper, we study the sample complexity of MRA, that is the number of observations needed to recover a generic signal with a given accuracy as a function of the SNR. Our results reveal a striking difference between the high and low SNR regimes. On the one hand, the picture at a high SNR is fairly standard in signal processing: the sample complexity scales proportionally to $1/\text{SNR}$. On the other hand, using information theoretic arguments, we show that the presence of the latent cyclic shifts has a profound effect on the sample complexity at low SNR, where the optimal sample complexity becomes proportional to $1/\text{SNR}^3$. We propose an algorithm based on the estimated bispectrum coupled with a tensor decomposition to achieve this sample complexity efficiently and provably. This draws the first complete picture for the sample complexity of the MRA model. We complement the results with the first algorithm for the heterogeneous case where θ in (1) is randomly drawn from a finite family of linearly independent vectors.

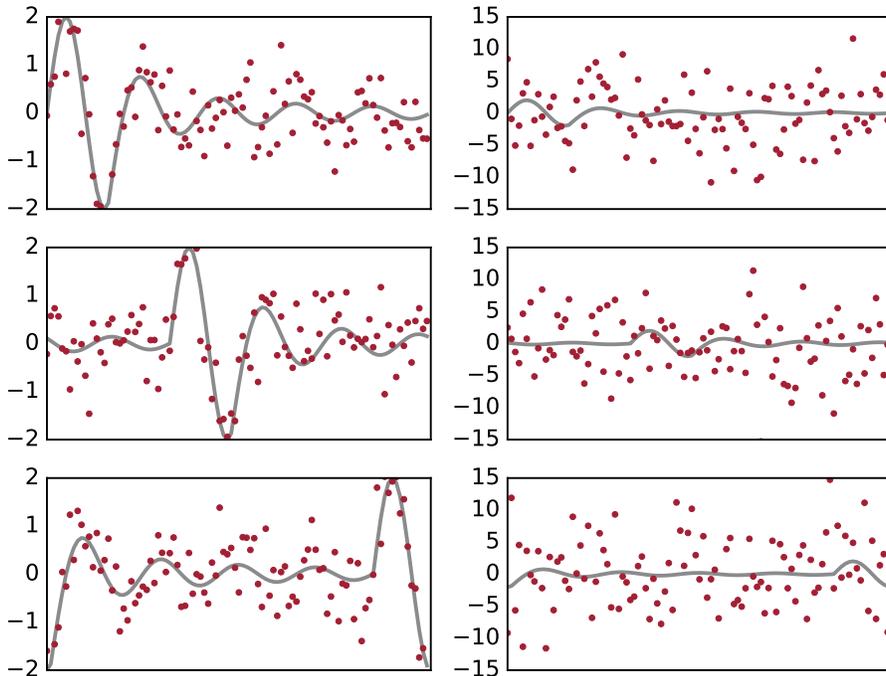


FIG 1. *Instances of the multi-reference alignment problem, at low ($\sigma = .5$, left column) and high ($\sigma = 5$, right column) noise levels. The true underlying signal appears in gray, and the noisy sample appears in red. When the noise level is low, features of the signal are still visible despite the noise; in the presence of large noise, however, the signals cannot reliably be aligned.*

2. OVERVIEW

In this section, we give an overview of our contributions and how they fit in the existing literature.

2.1 Existing methods

The difficulty of the multi-reference alignment problem resides in the fact that both the signal $\theta \in \mathbb{R}^d$ and the shifts $\ell_1, \dots, \ell_n \in \mathbb{Z}_d$ are unknown. If the shifts were known, one could easily estimate θ by taking the average of $R_{\ell_i}^{-1}y_i, i = 1, \dots, n$. In fact, this simple observation is the basis of the so-called “synchronization” approach [Sin11, BCSZ14, BCS15]: first estimate the shifts by $\tilde{\ell}_1, \dots, \tilde{\ell}_n \in \mathbb{Z}_d$ and then estimate θ by averaging the $R_{\tilde{\ell}_i}^{-1}y_i$ ’s. While the synchronization approach can be employed at high SNR, it is limited by the fact that at low SNR, even alignment of observations to the true signal yields inaccurate shift estimates [ADBS16].

Alternative approaches estimate distributions over possible cyclic shifts instead. This relaxation is reminiscent of relaxation employed in replacing the hard labelling of k -means with the soft labelling of mixture modelling in clustering. This approach leads naturally to an estimate of the underlying signal via Expectation-Maximization algorithms or analogous Bayesian approaches [Sch12b, LBTG13]. Nevertheless, these procedures lack theoretical guarantees, and may get stuck in local minima.

2.2 The method of invariants

Throughout this paper, we take a radically different approach called the *method of invariants*, which bypasses the estimation of the shifts altogether. This method exploits features of the signal that are *invariant* under cyclic shifts, or more generally under the action of a group of transformations of the signal. The norm $\|\theta\|$ of θ is an example of an invariant feature: it is clearly invariant under cyclic shifts of θ and can therefore be estimated consistently in the MRA model without difficulty, assuming σ is known.

More generally, the method of invariants leverages invariant polynomials in the coordinates of θ . These are captured by moment tensors $T_r(\theta)$ defined by

$$(2) \quad T_r(\theta) := \frac{1}{d} \sum_{\ell=1}^d (R_\ell \theta)^{\otimes r}.$$

For example, the first moment tensor reduces to the entrywise mean of the signal and the second to its autocorrelation matrix, a circulant matrix containing the d autocorrelations of θ . In what follows we regard r as a constant—in the context of this paper we always have $r \leq 5$. It is worth noting that the autocorrelation matrix carries the same information as the power spectrum—the square of the absolute value of the Fourier coefficients of the signal—which is oftentimes used as an invariant feature in signal processing. As a result, it does not carry enough information to allow for estimation of θ , since it provides only the magnitudes of the Fourier coefficients, but not their phases.

In particular, the third moment tensor corresponds to the triple autocorrelation of the signal. The triple autocorrelation has a Fourier-analytic analogue called the *bispectrum*, which is used in signal processing. The bispectrum of a signal θ is given by

$$\mathcal{B}(k_1, k_2) = \hat{\theta}_{k_1} \hat{\theta}_{k_2} \hat{\theta}_{-k_1-k_2},$$

where $\hat{\theta}$ is the Fourier transform of θ , $k_1, k_2 \in [d]$, and the indices are taken modulo d . It was originally introduced in a statistical context [Bri91, Tuk84] and is known to uniquely determine generic signals up to cyclic shift [Kak09]. It is not difficult to see, for example, that it uniquely determines signals that have nonvanishing spectrum. This fact has been exploited to obtain estimates for alignment problems [SG92, Gia89, BBM⁺17].

An apparent drawback of this approach is that the sample average based estimator for T_3 has a variance of order σ^6/n when σ is large. It suggests that in the low SNR regime, any approach relying on the bispectrum requires at least order $1/\text{SNR}^3$ samples. In fact, we show that this number of samples is actually a fundamental requirement of the problem when the shifts are sampled from the uniform distribution, independent of the approach taken (following ideas developed in [BRW17]). Note that the dependency on the SNR is very different from that observed in more classical problems in signal processing, for which order $1/\text{SNR}$ samples usually suffice. This shows that the cyclic transformations fundamentally change the difficulty of the problem. A similar phenomenon has been demonstrated for a Boolean version of MRA [APS17].

To complement our lower bound, we also propose an algorithm capable of provably estimating generic signals θ with a number of samples that is optimal in the SNR, but potentially not in the dimension d . It operates by decomposing the

third moment tensor. For this reason our algorithm can be viewed as a procedure that provably inverts the bispectrum for generic signals in a stable form, an important problem in signal processing and the imaging sciences. A detailed description can be found later in the paper.

2.3 Non-generic signals

The tensor-based methods for the multi-reference alignment problem we present work only for signals that are suitably generic, for instance, signals with non-vanishing power spectrum. In fact, non-generic signals can exhibit significantly worse behavior with a sample complexity of order $1/\text{SNR}^{\sqrt{d}}$ rather than $1/\text{SNR}^3$ [BRW17] but this pessimistic scenario does not seem to be representative of signals encountered in practice. In fact, signals that cannot be estimated with $1/\text{SNR}^3$ samples exhibit power spectra satisfying algebraic relations and form a set of zero Lebesgue measure. The optimal sample complexity for non-generic signals is sensitive to specific properties of the power spectrum, and even apparently benign signal processing techniques such as low-pass filtering may produce signals that are *harder* to estimate than the original one.

2.4 The heterogeneity problem

One of the main challenges in Cryo-EM reconstruction is the heterogeneity problem, where one observes noisy projection images of multiple unknown conformations of the same molecule. The MRA model can be extended to accommodate heterogeneity by assuming that in (1), the vector $\theta \in \mathbb{R}^d$ is also a latent variable drawn from a finite set of unknown vectors $\mathcal{C} = \{\theta^{(1)}, \dots, \theta^{(K)}\}$. The goal here is to recover the set \mathcal{C} up to a cyclic shift and the distribution of θ on this set.

Most approaches to the heterogeneity problem attempt to cluster the samples in K clusters, corresponding to each conformation or different signal. For Cryo-EM, the prevailing method is marginalized maximum likelihood as implemented for example in RELION [Sch12a, Sch12b, Sch16]. Semidefinite relaxations [BCS15, LS16] can also address heterogeneity but none of these methods comes with theoretical guarantees at low SNR.

Our approach based on the method of invariants coupled with tensor decomposition techniques extends to the heterogeneous setup. It yields the first algorithm capable of provably solving the heterogeneous MRA at arbitrarily low SNR, albeit at a potentially suboptimal sample complexity of $1/\text{SNR}^5$.

2.5 Connections to Cryo-EM

One of the main motivations to study the multi-reference alignment problem is that it serves as a simpler surrogate for Cryo-EM. This paper indicates potentially fruitful directions for work on the Cryo-EM problem. We offer theoretical vindication for the use of invariant methods in Cryo-EM, a proposal which dates back to Zvi Kam [Kam80]. In addition, Cryo-EM has the potential to analyze heterogeneous mixtures and thereby determine the structures of complexes in different functional states. Our analysis of the properties of an analogous model for heterogeneous multi-reference alignment serves as a first step towards a complete statistical theory of heterogeneous Cryo-EM.

2.6 Notation

We use $[d]$ to represent the set $\{1, \dots, d\}$ and I_d to represent the $d \times d$ identity matrix. The smallest and largest singular values of a matrix are denoted σ_{\min} and σ_{\max} , respectively. The symbol $\text{poly}(\cdot)$ refers to an unspecified polynomial with constant coefficients. C_d is used to refer to a constant that may depend on d but not on other parameters, and it may refer to a different constant in different appearances throughout the text. The expression $f(n) = O(g(n))$ means that there exists a constant C such that $f(n) \leq Cg(n)$ for all n , and we write $O_d(g(n))$ when the constant may depend on d . We write $g(n) = \Omega(f(n))$ or $g(n) = \Omega_d(f(n))$ when $f(n) = O(g(n))$ or $f(n) = O_d(g(n))$, respectively.

3. FUNDAMENTAL LIMITATIONS

In this section, we establish the fundamental limits of MRA and point to shortcomings of existing strategies to achieve optimal sample complexity.

3.1 Lower bounds for sample complexity

Since observations in the MRA model (1) are invariant under a global cyclic shift, one may only identify θ up to such a global shift. To account for this fact, it is natural to employ the following shift-invariant distance between vectors $\theta, \tau \in \mathbb{R}^d$:

$$\rho(\theta, \tau) = \min_{\ell \in \mathbb{Z}_d} \|\theta - R_\ell \tau\|_2.$$

By applying an independent and uniform random cyclic shift to each observation, we can always reduce the MRA model to the case where ℓ_1, \dots, ℓ_n are drawn i.i.d. uniformly from $[d]$. In this case, the distribution of y in (1) is a uniform mixture of the d Gaussian distributions $\mathcal{N}(\theta, \sigma^2 I), \dots, \mathcal{N}(R_{d-1} \theta, \sigma^2 I)$. If y is generated according to this distribution, we call it a “sample from MRA with signal θ .” The statistical properties of this Gaussian mixture are analyzed in [BRW17].

If σ is small—that is, if the SNR is sufficiently large—then the signals can be aligned (for example, via the synchronization approach [Sin11]), and therefore θ can be estimated accurately on the basis of n samples from MRA with signal θ as long as $n \geq C_d/\text{SNR}$ for some constant C_d . This is the same dependence that would be expected in the absence of shifts. Strikingly, the situation in the high-noise regime (when the SNR is low) is very different: estimation is impossible unless $n \geq C_d/\text{SNR}^3$ for some constant C_d that depends on d .

THEOREM 1. *Fix $d > 2$ and $\varepsilon \in (0, 1)$. There exists a constant C_d such that the following holds with constant probability: for any estimator $\hat{\theta}$ based on n samples from (1) there exists $\theta \in \mathbb{R}^d$ with $\|\theta\|_2 = 1$, such that $\rho(\hat{\theta}, \theta) \geq \varepsilon$ whenever $n \leq C_d \sigma^6 \varepsilon^{-2}$.*

This result is established using a standard two-point testing reduction [LeC73] by exhibiting two signals θ and τ such that $\rho(\theta, \tau) \geq 2\varepsilon$ but the distributions of n samples arising from MRA with signal θ and from MRA with signal τ are statistically indistinguishable, a fact which we prove using the tight information-theoretic bounds developed in [BRW17].

Note that this lower bound is limited to the case of uniform shifts. In general, it may be possible to leverage departure from uniformity to obtain a better sample complexity.

3.2 The importance of high frequencies

As we will see in the next section, the lower bound C_d/SNR^3 is in fact tight for generic signals. In general, the optimal number of samples depend on specific properties of the support of the Fourier transform of the original signal. This dependence is often counter-intuitive, as the following example shows.

Many approaches to the alignment problem implicitly rely on the strategy of first estimating low frequencies of a signal, and then using this initial estimate to estimate higher frequencies (see [BGPS16]). In other words, these strategies assume that estimating a low-pass version of a signal is no harder than estimating the original signal. Surprisingly, this is *not* the case in general: high-frequency information is sometimes necessary to estimate information at lower frequencies.

Let us take $d \geq 14$ congruent to 2 (mod 4) and $\theta \in \mathbb{R}^d$ a signal whose Fourier transform $\hat{\theta}$ satisfies $\hat{\theta}_1 = \hat{\theta}_{-1} = 0$ but otherwise has full support. We show in the appendix that we can estimate θ with $O_d(1/\text{SNR}^3)$ samples. Surprisingly, if we low-pass θ by setting $\hat{\theta}_j = 0$ for all $|j| > 4$, then $\Omega_d(1/\text{SNR}^4)$ samples are needed to estimate it. An illustration of this phenomenon appears in Figure 2.

The difficulty in recovering the low-pass version arises from the following simple observation: if $\hat{\theta}_j = 0$ for all $j \notin \{\pm 2, \pm 3, \pm 4\}$, then the only nonzero entry of the bispectrum is $\mathcal{B}(2, 2)$. This implies that the bispectrum carries no information about the phase of $\hat{\theta}_3$, and as a result, [BRW17, Theorem 3] implies that any two such signals which agree on their second and fourth Fourier coefficient are statistically indistinguishable unless $n \geq C_d \sigma^8$.

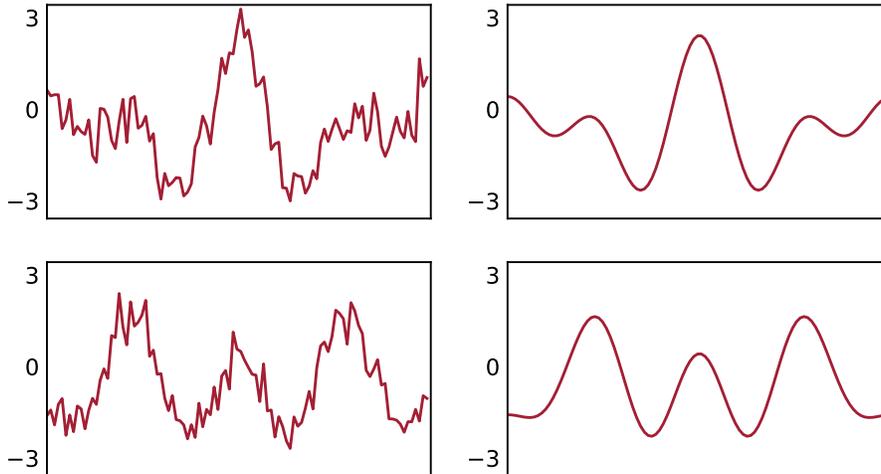


FIG 2. *Two signals whose Fourier transforms have almost full support (left column) and their corresponding low-pass versions (right column). Distinguishing between the original signals is possible with $O(1/\text{SNR}^3)$ samples, but distinguishing between the low-pass versions requires $\Omega(1/\text{SNR}^4)$ samples. This illustrates the importance of high frequencies in the MRA model.*

4. EFFICIENT RECOVERY VIA TENSOR METHODS

In this section, we describe an efficient algorithm that matches the sample complexity of Theorem 1. Namely it outputs an estimator $\tilde{\theta}$ of θ such that $\rho(\tilde{\theta}, \theta) \leq \varepsilon$ whenever $n \geq C_d \sigma^6 \varepsilon^{-2}$, for some constant C_d not necessarily matching the constant appearing in Theorem 1.

Our approach uses the method of invariants by estimating invariant features in the third moment tensor T_3 defined in (2). To estimate this tensor, we use the following empirical version that removes spurious symmetries:

$$(3) \quad \tilde{T}_3 = \frac{1}{dn} \sum_{i=1}^n \sum_{j=1}^d ((R_j y_i)^{\otimes 3} - 3 \text{sym}(R_j y_i \otimes I_d))$$

where

$$(4) \quad \text{sym}(A)_{a_1 \dots a_k} = \frac{1}{k!} \sum_{\pi} A_{\pi(a_1) \dots \pi(a_k)}.$$

The following lemma characterizes the bias and variance of this estimator.

LEMMA 1. *The estimator \tilde{T}_3 is an unbiased estimator of T_3 . Moreover, each entry of \tilde{T}_3 has a variance of order σ^6/n so long as σ is bounded below by a positive constant.*

4.1 Jennrich's Algorithm for Tensor Decomposition

In this section, we discuss an efficient algorithm that provably solves the multi-reference alignment problem for a generic signal $\theta \in \mathbb{R}^d$. It involves the spectral decomposition of the tensor of empirical third moments. While using tensor decomposition in the context of multi-reference alignment is one of the contributions of this paper, such decompositions have been long studied and a sophisticated machinery has been developed over the years; for a survey see, for example, [Moi14, Chapter 3].

The specific algorithm that we will use to estimate θ from T_3 (see (2)) is a standard tensor decomposition algorithm known as *Jennrich's algorithm* (proposed in [Har70] and credited to Robert Jennrich). The version described below allows the recovery of vectors u_1, \dots, u_r (up to simple transformations) from a noisy version of the tensor

$$(5) \quad T = \sum_{j=1}^r u_j \otimes u_j \otimes v_j \in \mathbb{R}^{m \times m \times p}$$

We call a tensor T satisfying (5) for some $r \leq m$ a *low rank tensor*.

While Jennrich's algorithm provably solves the problem in polynomial time for generic signals, we note that this algorithm may fail for worst-case choices of signals. Indeed, the statistical rates of estimation for worst-case signals are provably very different and can require exponentially many samples (on the dimensional of the signal) [BRW17]. Our results do not simply hold with high probability on an input distribution of the signal, but instead hold for almost every signal, akin to the smooth analysis framework of Spielman and Teng [ST04].

Jennrich’s Algorithm ([Har70, LRA93]).
Input: Low rank tensor $T = \sum_{j=1}^r u_j \otimes u_j \otimes v_j \in \mathbb{R}^{m \times m \times p}$.
Output: Matrix $U = [u_j, \dots, u_r] \in \mathbb{R}^{m \times r}$ (up to permutation and scaling of columns).

- ▶ Choose random unit vectors $a, b \in \mathbb{R}^p$, and form matrices $A, B \in \mathbb{R}^{m \times m}$ with entries:

$$A_{ij} = \sum_k T_{ijk} a_k, \quad A = \sum_{j=1}^r \langle v_j, a \rangle u_j \otimes u_j,$$

$$B_{ij} = \sum_k T_{ijk} b_k, \quad B = \sum_{j=1}^r \langle v_j, b \rangle u_j \otimes u_j$$
- ▶ Let W be the matrix whose columns are the first r left singular vectors of A .
- ▶ Compute $M = W^\top A W (W^\top B W)^{-1}$.
- ▶ Output $U = W P$, where $M = P D P^{-1}$ is the eigendecomposition of M .

Jennrich’s algorithm requires two eigendecompositions and can therefore be implemented very efficiently even on large scale problems. It also enjoys robustness guarantees as illustrated by the following theorem.

THEOREM 2 ([GVX14], Theorem 5.2). *Let T be a tensor of the form (5) with all u_j linearly independent, and define $\kappa = \sigma_{\max}(U)/\sigma_{\min}(U)$. Moreover, fix $\varepsilon > 0$ and let \tilde{T} satisfy $\|\tilde{T} - T\|_F \leq \varepsilon$. Then Jennrich’s algorithm applied to \tilde{T} returns unit vectors $\tilde{u}_j, j = 1, \dots, r$ such that there exists a permutation π and scalars β_j such that*

$$(6) \quad \max_{j \in [r]} \|\tilde{u}_j - \beta_j u_{\pi(j)}\|_\infty \leq \varepsilon \text{ poly}(m, \kappa)$$

with high probability.

We show in the appendix that T_3 is indeed a low rank tensor of the form (5), with $m = p = d$, $u_j = v_j = R_{j-1} \theta$ (for $j = 1, \dots, r$), and $U = [\theta, R_1 \theta, \dots, R_{d-1} \theta]$. Moreover, it can be shown that $\kappa(U) = \max_{j,k} \{|\hat{\theta}_j|/|\hat{\theta}_k|\} =: \kappa(\theta)$. This quantity is generically finite, and is bounded by $\text{poly}(d)$ with high probability if, for example $\theta \sim \mathcal{N}(0, d^{-1} I_d)$. Under these assumptions Jennrich’s algorithm applied to \tilde{T}_3 outputs in particular \tilde{u}_1 , such that

$$\|\tilde{u}_1 - \beta_1 R_j \theta\|_\infty \leq \frac{\sigma^3 d^c}{\sqrt{n}}, \quad c > 0,$$

with high probability for some $j \in [d]$ and some $\beta_j \in \mathbb{R}$. The scaling factor β_1 can be easily removed as follows. Estimate β_1 by $\tilde{\beta}_1$:

$$\tilde{\beta}_1 = \tilde{u}_1^\top \mathbf{1} / \tilde{\mu}, \quad \tilde{\mu} = \frac{1}{n} \sum_{i=1}^n y_i^\top \mathbf{1}.$$

We call this procedure **homoJen** for “homogeneous MRA via Jennrich”. It outputs $\tilde{\theta} = \tilde{u}_1 / \tilde{\beta}_1$. It enjoys the following theoretical guarantees.

THEOREM 3. *Fix $\sigma > .1$ and $\delta \in (0, 1)$ and assume $.1 \leq \|\theta\|_2 \leq 10$. Then, for any $\varepsilon > 0$, the `homoJen` algorithm based on n samples from (1) outputs $\tilde{\theta}$ such that $\rho(\tilde{\theta}, \theta) \leq \varepsilon$ with probability at least $1 - \delta$ whenever*

$$n \geq \sigma^6 \varepsilon^{-2} \text{poly}(d, \kappa(\theta), 1/\delta).$$

Note that the constants .1 and 10 are arbitrary and may be replaced by any other constants.

Note that in view of the lower bound appearing in Theorem 1, the sample complexity of the modified Jennrich algorithm is optimal apart possibly for factors depending polynomially in d . In fact, the left half of Figure 3 illustrates the practical effectiveness of the class of methods proposed here: Given samples in the multi-reference alignment problem (1) for $d = 5$, this plot shows the performance of our estimator based on Jennrich’s algorithm for various values of number of samples n and signal to noise ratio SNR (inversely proportional to σ^2). For each value of n and SNR, 500 trials were performed and the mean squared error plotted in color, with blue representing low estimation error (good performance) and red representing high estimation error (bad performance, recovering only noise). The black lines are lines of constant n/SNR^3 , representing the theoretical prediction of Theorem 3.

Several other bispectrum-based algorithms appear in the literature; see [BBM⁺17] for a recent empirical study. These may perform better in practice, but they largely do not come with the theoretical guarantees of the algorithm proposed here, and they do not extend to the heterogeneous case discussed below.

5. HETEROGENEITY

This section describes one of the main contributions of our paper, an algorithm provably able to solve the heterogeneous multi-reference alignment problem. We recall this model here for completeness. Recall that in this model, we observe

$$(7) \quad y_i = R_{\ell_i} \theta^{(Z_i)} + \sigma \xi_i, \quad i = 1, \dots, n,$$

where $Z_1, \dots, Z_n \in \{1, \dots, K\}$ are i.i.d. latent variables such that $\Pr(Z_i = k) = \pi_k, k \in [K]$ that are independent of all other variables and $\theta^{(k)} \in \mathbb{R}^d, k \in [K]$ are unknown vectors. The other variables are specified as in the homogeneous model (1). The goal here is to recover the set of vectors $\theta^{(k)} \in \mathbb{R}^d, k = 1, \dots, K$ up to a cyclic shift and the probability mass function $\{\pi_k\}_{k \in [K]}$.

The method of invariants described above can be extended to handle the heterogeneous model (7). In this case our method proceeds by estimating the mixtures of signals from an unbiased estimator \tilde{T}_5 for the 5-tensor

$$T_5 = \sum_{k=1}^K \sum_{\ell=1}^d \frac{\pi_k}{d} (R_{\ell} \theta^{(k)})^{\otimes 5}.$$

Using manipulations similar to the ones arising in the proof of Lemma 1, it is not hard to show that T_5 can be rewritten as

$$T_5 = \mathbb{E}[y_i^{\otimes 5}] - 10\sigma^2 \text{sym}(T_3 \otimes I_d) - 15\sigma^4 \text{sym}(T_1 \otimes I_d^{\otimes 2})$$

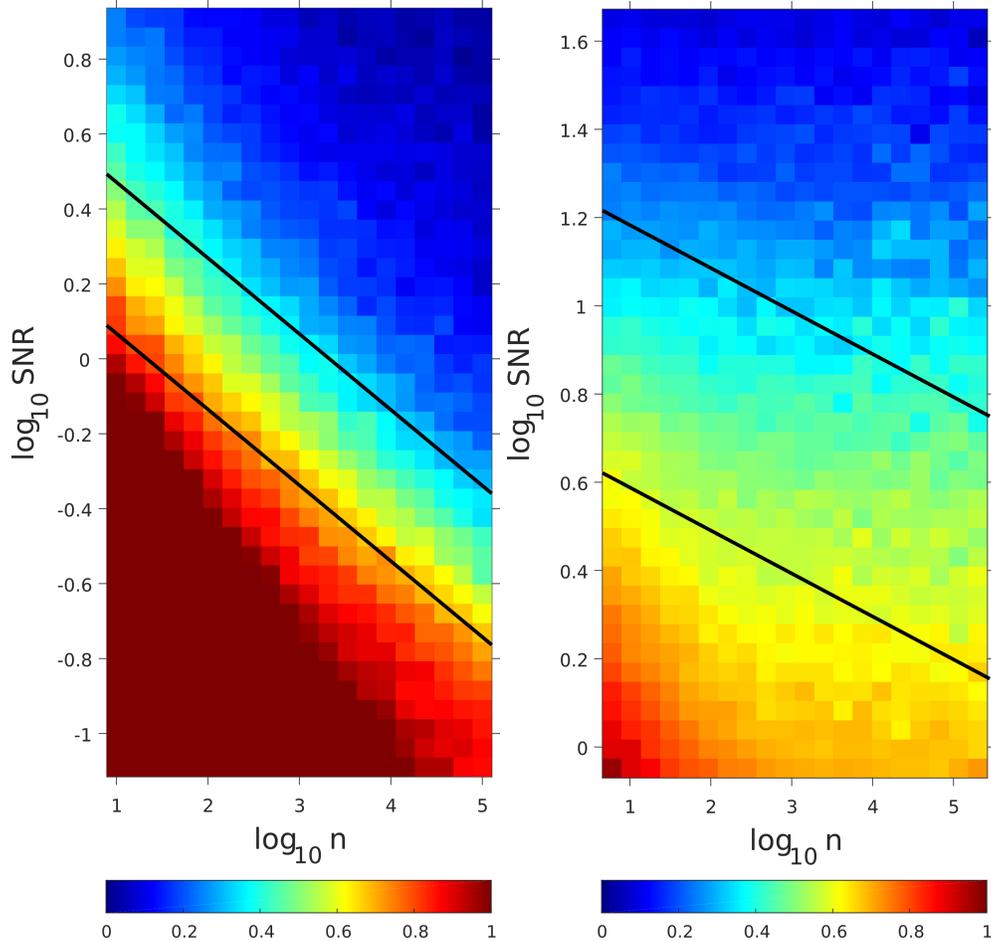


FIG 3. **Left:** Empirical performance of `homoJen` to solve the homogeneous multi-reference alignment problem, with $d = 5$ and 500 trials per cell. Color indicates the root mean square of the error $\rho(\tilde{\theta}, \theta)$, where $\tilde{\theta}$ is the output of `homoJen`; blue indicates less error. The lines have constant n/SNR^3 , representing the theoretical guarantee of Theorem 3. **Right:** Empirical performance of `heteroJen` to solve the heterogeneous multi-reference alignment problem, with $d = 6$, $K = 3$, and 500 trials per cell. Color indicates the sample average of $\sum_k \min_j \rho(\tilde{\theta}_j, \theta_k)$, where $\{\tilde{\theta}_1, \dots, \tilde{\theta}_K\}$ denotes the output of `heteroJen`; blue indicates less error. The lines have constant n/SNR^5 , representing the theoretical guarantee of Theorem 4.

where $T_1 = \mathbb{E}[y_i]$ and T_3 is defined in (2). Therefore an unbiased estimator of T_5 is given by

$$\tilde{T}_5 = \frac{1}{n} \sum_{i=1}^n y_i^{\otimes 5} - 10\sigma^2 \text{sym}(\tilde{T}_3 \otimes I_d) - 15\sigma^4 \text{sym}\left(\frac{1}{n} \sum_{i=1}^n y_i \otimes I_d^{\otimes 2}\right),$$

where \tilde{T}_3 is given in (3). Moreover, each entry of \tilde{T}_5 has variance of order σ^{10}/n so long as σ is bounded below by a positive constant.

We propose a method that consists in applying Jennrich’s algorithm to an appropriate flattening of \tilde{T}_5 . We call it `heteroJen`. This algorithm hinges on the following observation: the 5-tensor T_5 can be flattened into a 3-tensor of shape

$d^2 \times d^2 \times d$ that admits the following low-rank decomposition:

$$\sum_{k=1}^K \sum_{\ell=1}^d \frac{\pi_k}{d} (R_\ell \theta^{(k)})^{\otimes 2} \otimes (R_\ell \theta^{(k)})^{\otimes 2} \otimes (R_\ell \theta^{(k)}).$$

The algorithm `heteroJen` then proceeds by plugging \tilde{T}_5 into the above flattening operation and then applying Jennrich’s algorithm to the resulting 3-tensor of shape $d^2 \times d^2 \times d$. Theorem 2 implies that this procedure outputs vectors u_i , $1 \leq i \leq dK$ with the following guarantees: there exist scalars β_i and a bijection $a \times b : [dK] \rightarrow [d] \times [K]$ satisfying

$$\|\tilde{u}_i - \beta_i (R_{a(i)} \theta^{(b(i))})^{\otimes 2}\|_\infty \leq C_K \frac{\sigma^5}{\sqrt{n}} \text{poly}(d),$$

with high probability.

We compute \tilde{v}_i as the leading eigenvector of the $d \times d$ matrix \tilde{u}_i . Letting \tilde{V}_3 be the $d^3 \times dK$ matrix with columns $\tilde{v}_i^{\otimes 3}$, we estimate $\tilde{\alpha} \in \mathbb{R}^{dK}$ as the least-squares solution to $\tilde{V}_3 \tilde{\alpha} = \text{vec}(T_3)$. The vectors $\tilde{w}_i := \tilde{\alpha}^{1/3} \tilde{v}_i$ (with entrywise exponentiation) now comprise dK redundant estimates to the K original signals θ_k ; we remove this redundancy via a clustering procedure. These steps are rigorously described and justified in the appendix.

This method succeeds when $K \leq d/2$ as illustrated by the following theorem. It involves the condition number of the matrix $[\text{vec}((R_1 \theta^{(1)})^{\otimes 2}), \dots, \text{vec}((R_d \theta^{(K)})^{\otimes 2})]$, which we denote $\kappa(\{\theta^{(1)}, \dots, \theta^{(K)}\})$.

THEOREM 4. *Fix $\sigma > .1$ and $\delta \in (0, 1)$. Assume that $.1 \leq \|\theta^{(k)}\|_2 \leq 10$ for all $k \in [K]$ and that $K \leq \lceil d/2 \rceil$. Then, for any $\varepsilon > 0$, the `heteroJen` algorithm based on n samples from (7) outputs $\{\tilde{\theta}_1, \dots, \tilde{\theta}_K\}$ such that*

$$\sum_k \min_j \rho(\tilde{\theta}_j, \theta_k) \leq \varepsilon,$$

with probability at least $1 - \delta$ whenever

$$n \geq C_K \sigma^{10} \varepsilon^{-2} \text{poly}(d, \kappa, 1/\delta),$$

where $\kappa = \kappa(\{\theta^{(1)}, \dots, \theta^{(K)}\})$.

We prove this result in the appendix, along with the following lemma.

LEMMA 2. *The condition number $\kappa = \kappa(\{\theta^{(1)}, \dots, \theta^{(K)}\})$ is generically finite.*

To the best of our knowledge, `heteroJen` is the first efficient method shown to solve the heterogeneous MRA problem with arbitrarily large noise variance (and enough samples). We conjecture, however, that the σ^{10} dependence is not optimal, and that generic heterogeneous mixtures (perhaps with only a constant number of components) should have sample complexity $C_K \sigma^6 \varepsilon^{-2} \text{poly}(d)$ in well conditioned cases, akin to the homogeneous case.

Figure 3 illustrates the effectiveness of the class of methods proposed here: Given samples in the heterogeneous multi-reference alignment problem (7), this plot shows the performance of our estimator based on Jennrich’s algorithm for various values of n and SNR. The black lines represent the theoretical prediction of Theorem 4.

Acknowledgments

We thank Alex Wein for many insightful discussions on the topic of this paper.

APPENDIX A: TECHNICAL PROOFS

A.1 Proof of Theorem 1

In what follows, let c_d and C_d be constants depending on d whose value may change from line to line. Let θ be any zero-mean signal such that $\|\theta\|_2 = 1$, and let $\hat{\theta}$ be its Fourier transform. There exists a coefficient $\hat{\theta}_j$ with $j \neq 0$ such that $|\hat{\theta}_j| \geq c_d$. Define τ by setting

$$\hat{\tau}_k = \begin{cases} e^{i\delta}\hat{\theta}_j & \text{if } k = j \\ e^{-i\delta}\hat{\theta}_{-j} & \text{if } k = -j \\ \theta_k & \text{otherwise,} \end{cases}$$

where $\delta = c_d\varepsilon$. It is clear that $\rho(\theta, \tau) \geq c_d\varepsilon$ for some constant c_d , since for any shift R we have $\|\hat{\theta}_j - \widehat{R\tau}_j\|_2 \geq c_d|1 - e^{i\delta}|$. On the other hand, since θ and τ have the same mean and power spectrum, [BRW17, Theorem 3] implies that if P_θ and P_τ are the Gaussian mixtures corresponding to MRA with signals θ and τ respectively, then the Kullback–Leibler divergence $D(P_\theta \| P_\tau)$ is at most $C_d\sigma^{-6}\varepsilon^2$. The chain rule for divergence implies $D(P_\theta^n \| P_\tau^n) \leq C_d\sigma^{-6}\varepsilon^2n$, so Pinsker’s inequality implies that if $C_d\sigma^{-6}\varepsilon^2n \leq .01$, then no statistical procedure can distinguish between n samples from MRA with signals θ and τ with probability better than $2/3$. \square

A.2 Analysis of Low-pass example

To show that θ can be recovered with $O_d(1/\text{SNR}^3)$ samples, it suffices to show that the phases of the Fourier coefficients of θ can be reconstructed uniquely from its bispectrum. Given a complex number z , denote by $\arg(z)$ its phase. By applying a circular shift, we can assume without loss of generality that $\arg(\hat{\theta}_2) \in [0, 4\pi/d)$ and that $\arg(\hat{\theta}_3) \in [0, \pi)$. It is easy to check that the identity $2 \sum_{k=2}^{(d-6)/4} \arg(\mathcal{B}(2, 2k)) + \arg(\mathcal{B}((d-2)/2, (d-2)/2)) = \frac{d}{2} \arg(\hat{\theta}_2)$ holds modulo 2π , and the assumption that $\arg(\hat{\theta}_2) \in [0, 4\pi/d)$ implies that the choice of $\arg(\hat{\theta}_2)$ is unique. This implies that all even-indexed phases can be recovered. We also have the simple identity $\arg(\hat{\theta}_6) + \arg(\mathcal{B}(3, 3)) = 2 \arg(\hat{\theta}_3)$ modulo 2π , and the assumption that $\arg(\hat{\theta}_3) \in [0, \pi)$ implies that the choice is unique. Combined with the knowledge of $\arg(\hat{\theta}_2)$, this implies recoverability of all odd-indexed phases.

A.3 Proof of Lemma 1

If $\xi_i \sim \mathcal{N}(0, I_d)$, then both $\mathbb{E}[\xi_i]$ and $\mathbb{E}[\xi_i^{\otimes 3}]$ are zero. This implies that

$$\mathbb{E}_{ij}[(R_j y_i)^{\otimes 3}] = \mathbb{E}_{ij}[(R_j \theta + \sigma \xi)^{\otimes 3}] = \mathbb{E}_j(R_j \theta)^{\otimes 3} + 3 \text{sym}((\mathbb{E}_j R_j \theta) \otimes I_d),$$

so \tilde{T}_3 is an unbiased estimator of T_3 .

Each entry of y_i is a Gaussian with variance σ^2 , so the entries of $\text{sym}(y_i \otimes I_d)$ have variance of order σ^2 , and the entries of $y_i^{\otimes 3}$ have variance of order σ^6 ; the latter dominates for σ bounded away from 0. The claim follows. \square

A.4 Proof of Theorem 3

We construct an algorithm as follows. Following Lemma 1, we use some number n of samples to construct an estimate \tilde{T}_3 of T_3 . We apply Jennrich's algorithm to produce unit vectors \tilde{u}_i such that, for some permutation π and some scalars β_i , $\|\tilde{u}_{\pi(i)} - \beta_i R_i \theta\|_\infty \leq \alpha$, following the guarantee of Theorem 2. To recover this unknown scaling, we compute the sample mean $\tilde{\mu} = \frac{1}{n} \sum_i \langle y_i, \frac{1}{d} \mathbf{1} \rangle$, estimate a scale factor as $\tilde{\beta} = \langle \tilde{u}_1, \frac{1}{d} \mathbf{1} \rangle / \tilde{\mu}$, and output $\tilde{\theta} = \tilde{\beta}^{-1} u_1$. Here $\mathbf{1}$ denotes the all-ones vector in \mathbb{R}^d .

As we are only concerned with polynomial dependence and not detailed bounds, we write $A \approx B$ if we can bound

$$|A - B| \leq \alpha \text{poly}(d, \kappa(\theta), \delta^{-1}) + \sigma/\sqrt{n} \text{poly}(d, \kappa(\theta), \delta^{-1})$$

so long as¹ $\alpha \leq 1$ and $\sigma/\sqrt{n} \leq 1$; we apply this also to vectors in 2-norm or (equivalently) most other common norms.

Theorem 2 guarantees us that $\tilde{u} \approx \beta_i R_i \theta$. Taking norms, we have $1 \approx |\beta_i| \|\theta\|_2$; as $\|\theta\|_2$ is bounded above and below by constants, we have that $|\beta_i| \approx 1/\|\theta\|_2$ is also of constant order. Note also that by Chebyshev we have that $|\tilde{\mu} - \mu| \leq \sigma\sqrt{\delta}/\sqrt{2nd}$ with probability $1 - \delta/2$, so that $\mu \approx \tilde{\mu}$. From $\tilde{u} \approx \beta_i R_i \theta$ we also derive that

$$(8) \quad \langle \tilde{u}, \mathbf{1} \rangle / d \approx \beta_i \mu \approx \beta_i \tilde{\mu}.$$

We know that $\|\theta\|_2 \leq \sigma_{\max}(U)$ and that $\sigma_{\min}(U) \leq \|U \cdot \frac{1}{\sqrt{d}} \mathbf{1}\|_2 = d|\mu|$, so that $|\tilde{\mu}| \approx |\mu| \geq d\kappa(\theta)/\|\theta\|_2$; we are thus justified in dividing (8) by $\tilde{\mu}$ to obtain $\tilde{\beta} \approx \beta_i$, and $\beta_i/\tilde{\beta} \approx 1$. We now bound the total estimation error as follows:

$$\begin{aligned} \|\tilde{\beta}^{-1} \tilde{u} - R_i \theta\|_2 &\leq \|\tilde{\beta}^{-1} \tilde{u} - \beta_i^{-1} \tilde{u}\|_2 + \|\beta_i^{-1} \tilde{u} - R_i \theta\|_2 \\ &\leq |\tilde{\beta}^{-1} - \beta_i^{-1}| + |\beta_i|^{-1} \alpha \\ &= |\beta_i|^{-1} \left(\alpha + \left| \frac{\beta}{\tilde{\beta}} - 1 \right| \right) \approx 0. \end{aligned}$$

Thus in order to bound this estimation error to within ε , it suffices to require bounds of the form $\sigma/\sqrt{n} \leq \varepsilon/\text{poly}(d, \kappa(\theta), \delta^{-1})$ and $\alpha \leq \varepsilon/\text{poly}(d, \kappa(\theta), \delta^{-1})$. By Theorem 2, we achieve this bound on α from Jennrich's algorithm so long as $\|\tilde{T}_3 - T_3\|_F \leq \varepsilon/\text{poly}(d, \kappa, \delta^{-1})$. This estimation error is achieved with probability $1 - \delta/2$ so long as $n \geq \sigma^6 \text{poly}(d, \kappa, \delta^{-1})$, which also subsumes the explicit bound on σ/\sqrt{n} . By a union bound over the two probabilistic steps in this argument, the desired accuracy guarantee holds with probability $1 - \delta$. \square

A.5 Proof of Theorem 4

As in the proof of Theorem 3, we write $A \approx B$ if we can bound $|A - B| \leq \alpha \text{poly}(d, \kappa(U), \delta^{-1}) + \sigma^3/\sqrt{n} \text{poly}(d, \kappa(U), \delta^{-1})$ given that² $\alpha \leq 1$ and $\sigma^3/\sqrt{n} \leq 1$; we apply this also to vectors in 2-norm or (equivalently) most other common norms.

From Theorem 2, we are guaranteed that $\tilde{u}_i \approx \beta_i (R_{a(i)} \theta_{b(i)})^{\otimes 2}$; taking norms, we have $1 \approx |\beta_i| \|\theta_{b(i)}\|_2$, so that $|\beta_i| \approx \|\theta_{b(i)}\|_2^{-1}$ is of constant order. Now by the

¹The precise bound of 1 here is arbitrary.

²The precise bound of 1 here is arbitrary.

Davis–Kahan theorem, if $v_{\max}(M)$ denotes either choice of unit-length eigenvector of M corresponding to the eigenvalue of largest magnitude, we have

$$\tilde{v}_i := v_{\max}(\tilde{u}_i) \approx \varepsilon_i R_{a(i)} \theta_{b(i)} / \|\theta_{b(i)}\|_2,$$

for some sign $\varepsilon_i = \pm 1$. Then we have $\tilde{V}_3 \approx V_3$, where as above, \tilde{V}_3 is the $d^3 \times dK$ matrix whose columns are $\tilde{v}_i^{\otimes 3}$, and V_3 has columns $\varepsilon_i (R_{a(i)} \theta_{b(i)})^{\otimes 3} / \|\theta_{b(i)}\|_2^3$. Estimating T_3 by \tilde{T}_3 according to Lemma 1, we have $\tilde{T}_3 \approx T_3$ by Chebyshev, with probability $1 - \delta/2$. Note then that $V_3 \alpha = \text{vec}(T_3)$, where $\alpha_i = \varepsilon_i \|\theta_{b(i)}\|_2^3 / dK$. By the perturbation theory of linear systems, we are now guaranteed that, letting $\tilde{\alpha}$ be the least squares solution to $\tilde{V}_3 \tilde{\alpha} = \text{vec}(\tilde{T}_3)$, we have $\tilde{\alpha} \approx \alpha$, so long as the system is well-conditioned, which we defer to the following lemma:

LEMMA 3. $\kappa(V_3) \leq \kappa(U) \text{poly}(d)$.

As α_i is of constant order, it follows that $\tilde{w}_i := \tilde{\alpha}_i^{1/3} \tilde{v}_i \approx R_{a(i)} \theta_{b(i)}$, so that $\rho(\tilde{\alpha}_i^{1/3} \tilde{v}_i, \theta_{b(i)}) \approx 0$. We are thus guaranteed dK good estimates to the original K signals. We next discuss how to remove this redundancy by clustering.

Define the pseudometric on \mathbb{R}^d defined by $\rho_2(x, y) = \min_{1 \leq \ell \leq d} \|x^{\otimes 2} - (R_\ell y)^{\otimes 2}\|_2$. Note that

$$\rho_2(w_i, w_{i'}) \approx \rho_2(R_{a(i)} \theta_{b(i)}, R_{a(i')} \theta_{b(i')}) = \rho_2(\theta_{b(i)}, \theta_{b(i')}).$$

If $b(i) = b(i')$, so that the two estimates w_i and $w_{i'}$ should represent the same signal, we thus have $\rho(w_i, w_{i'}) \approx 0$. If $b(i) \neq b(i')$, we have

$$\begin{aligned} \rho_2(w_i, w_{i'}) &\approx \rho_2(\theta_{b(i)}, \theta_{b(i')}) = \min_{\ell} \|U(e_{b(i),0} - e_{b(i'),\ell})\|_2 \\ &\geq \sqrt{2} \sigma_{\min}(U) \\ &= 1/\text{poly}(d, \kappa(U)), \end{aligned}$$

where $e_{b,\ell} \in \mathbb{R}^{dK}$ is the standard basis vector corresponding to signal b and rotation ℓ . It follows that, provided α and σ^6/n are inverse-polynomially small in $d, \kappa(U), \delta^{-1}$, we exactly recover the clusters of estimates \tilde{w}_i corresponding to the same signal θ_k , simply by comparing on the metric ρ_2 and thresholding. Drawing one estimate \tilde{w}_i from each cluster, we obtain one estimate of each signal.

To conclude, in order to bound this estimation error to within ε , it suffices to require bounds of the form $\sigma^3/\sqrt{n} \leq \varepsilon/\text{poly}(d, \kappa(\theta), \delta^{-1})$ and $\alpha \leq \varepsilon/\text{poly}(d, \kappa(\theta), \delta^{-1})$. By Theorem 2, we achieve this bound on α from Jennrich’s algorithm so long as $\|\hat{T}_5 - T_5\|_F \leq \varepsilon/\text{poly}(d, \kappa, \delta^{-1})$. This estimation error is achieved with probability $1 - \delta/2$ so long as $n \geq \sigma^{10} \text{poly}(d, \kappa, \delta^{-1})$, which also subsumes the explicit bound on σ^3/\sqrt{n} . By a union bound over the two probabilistic steps in this argument, the desired accuracy guarantee holds with probability $1 - \delta$. \square

A.6 Proof of Lemma 3

We apply the following transformations which do not alter the condition number: we transform the rows by the third tensor power of a DFT, we permute the columns to sort by signal and rotation, and we negate columns according to the signs ε_i . It thus suffices to control the condition number of the $d^3 \times dK$ matrix

V_3' whose columns are $(\hat{R}_j \hat{\theta}_k)^{\otimes 3} / \|\theta_k\|_2^3$, where $\hat{R}_i = \text{diag}(\{\omega^{ij}\}_j)$ is the Fourier representation of a rotation action ($\omega = e^{2\pi/d}$), and θ_k is the Fourier transform of θ_k . Meanwhile, let V_2' be the $d^2 \times dK$ matrix with columns $(\hat{R}_j \hat{\theta}_k)^{\otimes 2}$, the Fourier transform of U , so that $\kappa(V_2') = \kappa(U)$.

Let $v \in dK$; then we have

$$\begin{aligned}
\|V_3'v\|_2^2 &= \sum_{\ell=1}^d \left\| V_2' \text{diag} \left(\{\omega^{j\ell} (\hat{\theta}_k)_\ell \|\theta_k\|_2^{-3}\}_{jk} \right) v \right\|_2^2 \\
&\geq \sum_{\ell=1}^d \sigma_{\min}(U)^2 \left\| \text{diag} \left(\{\omega^{j\ell} (\hat{\theta}_k)_\ell \|\theta_k\|_2^{-3}\}_{jk} \right) v \right\|_2^2 \\
&= \sum_{\ell} \sigma_{\min}(U)^2 \sum_{jk} |(\hat{\theta}_k)_\ell|^2 \|\theta_k\|_2^{-6} |v|_{jk}^2 \\
&= \sigma_{\min}(U)^2 \sum_k \|\theta_k\|_2^{-6} \left(\sum_{\ell} |(\hat{\theta}_k)_\ell|^2 \right) \left(\sum_j |v|_{jk}^2 \right) \\
&= \sigma_{\min}(U)^2 \sum_k \|\theta_k\|_2^{-4} \left(\sum_j |v|_{jk}^2 \right) \\
&\geq \sigma_{\min}(U)^2 10^{-4} \|v\|_2^2,
\end{aligned}$$

so that $\sigma_{\min}(V_3') \geq \sigma_{\min}(U) 10^{-2}$. Observing the norms of columns, it is clear that $\sigma_{\max}(U)$ and $\sigma_{\max}(V_3')$ are bounded above by $\text{poly}(d)$, so we conclude that $\kappa(V_3) = \kappa(V_3') \leq \kappa(U) \text{poly}(d)$, as desired. \square

A.7 Proof of Lemma 2

Suppose we have some nonzero linear relation $0 = \sum_{k=1}^K \sum_{\ell=1}^d c_{k,\ell} (R_\ell \theta_k) (R_\ell \theta_k)^\top$. Multiplying by a DFT on the left and its adjoint on the right, and examining the a, b entry, we have $0 = \sum_{k=1}^K (\hat{c}_k)_{a-b} \hat{\theta}_{ka} \hat{\theta}_{kb}$. Some $(\hat{c}_k)_\alpha$ is nonzero, yielding a nontrivial linear relation among the autocorrelation vectors v_k , $1 \leq k \leq K$, with $v_{k,j} = \hat{\theta}_{k,j} \hat{\theta}_{k,\alpha-j}$. These vectors satisfy the symmetry $v_{k,j} = v_{k,\alpha-j}$, but are generic on this subspace, which has dimension at least $\lceil d/2 \rceil$. Hence generically no such relation exists, and the matrix U has finite condition number. \square

REFERENCES

- [ADBS16] C. Aguerrebere, M. Delbraccio, A. Bartesaghi, and G. Sapiro. Fundamental limits in multi-image alignment. *IEEE Trans. Signal Process.*, 64(21):5707–5722, 2016.
- [APS17] E. Abbe, J. Pereira, and A. Singer. Sample complexity of the boolean multireference alignment problem. In *2017 IEEE International Symposium on Information Theory (ISIT)*, July 2017.
- [BBM⁺17] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer. Bispectrum inversion with application to multireference alignment. *Available online at arXiv:1705.00641 [cs.IT]*, 2017.
- [BCS15] A. S. Bandeira, Y. Chen, and A. Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *Available online at arXiv:1505.03840 [cs.CV]*, 2015.

- [BCSZ14] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu. Multireference alignment using semidefinite programming. In *ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science*, pages 459–470. ACM, New York, 2014.
- [BGPS16] A. Barnett, L. Greengard, A. Pataki, and M. Spivak. Rapid solution of the cryo-EM reconstruction problem by frequency marching. *SIAM J. Imaging Sci.*, 2016. To appear.
- [Bri91] D. R. Brillinger. Some history of the study of higher-order moments and spectra. *Statist. Sinica*, 1(2):465–476, 1991.
- [BRW17] A. S. Bandeira, P. Rigollet, and J. Weed. Optimal rates of estimation for multi-reference alignment. Available online at *arXiv:1702.08546 [math.ST]*, 2017.
- [Dia92] R. Diamond. On the multiple simultaneous superposition of molecular structures by rigid body transformations. *Protein Science*, 1(10):1279–1287, October 1992.
- [DM98] I. L. Dryden and K. V. Mardia. *Statistical shape analysis*. Wiley series in probability and statistics. Wiley, Chichester, 1998.
- [Edi16] Editorial. Method of the year 2015. *Nat Methods*, 13(1):1–1, 01 2016.
- [FZB02] H. Foroosh, J. Zerubia, and M. Berthod. Extension of phase correlation to subpixel registration. *IEEE Trans. Image Processing*, 11(3):188–200, 2002.
- [Gia89] G. B. Giannakis. Signal reconstruction from multiple correlations: Frequency-and time-domain approaches. *JOSA A*, 6(5):682–697, 1989.
- [GVX14] N. Goyal, S. Vempala, and Y. Xiao. Fourier PCA and robust tensor decomposition. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 584–593. ACM, 2014.
- [Har70] R. Harshman. Foundations of the PARAFAC procedure: Model and conditions for an explanatory multimodal factor analysis. Technical report, Tech. Rep. UCLA Working Papers in Phonetics 16, University of California, Los Angeles, Los Angeles, CA, December. 13, 27, 1970.
- [Kak09] R. Kakarala. Completeness of bispectrum on compact groups. *arXiv preprint arXiv:0902.0196*, 2009.
- [Kam80] Z. Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [LBTG13] D. Lyumkis, A. F. Brilot, D. L. Theobald, and N. Grigorieff. Likelihood-based classification of cryo-EM images using FREALIGN. *Journal of structural biology*, 183(3):377–388, 2013.
- [LeC73] L. LeCam. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, pages 38–53, 1973.
- [LRA93] S. Leurgans, R. Ross, and R. Abel. A decomposition for three-way arrays. *SIAM Journal on Matrix Analysis and Applications*, 14(4):1064–1083, 1993.
- [LS16] R. R. Lederman and A. Singer. A representation theory perspective on simultaneous alignment and classification. Available online at *arXiv:1607.03464 [cs.CV]*, 2016.
- [Moi14] A. Moitra. Algorithmic aspects of machine learning. *Lecture notes (MIT)*, 2014.

- [Nog16] E. Nogales. The development of cryo-EM into a mainstream structural biology technique. *Nat Methods*, 13(1):24–27, 01 2016.
- [PC14] W. Park and G. S. Chirikjian. An assembly automation approach to alignment of noncircular projections in electron microscopy. *IEEE Transactions on Automation Science and Engineering*, 11(3):668–679, 2014.
- [PMMC11] W. Park, C. R. Midgett, D. R. Madden, and G. S. Chirikjian. A stochastic kinematic model of class averaging in single-particle electron microscopy. *Int. J. Rob. Res.*, 30(6):730–754, 2011.
- [PZAF05] R. G. Pita, M. R. Zurera, P. J. Amores, and F. L. Ferreras. Using multi-layer perceptrons to align high range resolution radar signals. In W. Duch, J. Kacprzyk, E. Oja, and S. Zadrozny, editors, *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005*, volume 3697 of *Lecture Notes in Computer Science*, pages 911–916. Springer Berlin Heidelberg, 2005.
- [RCBL16] D. Rosen, L. Carlone, A. Bandeira, and J. Leonard. A certifiably correct algorithm for synchronization over the special Euclidean group. In *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, San Francisco, CA, December 2016.
- [Sch12a] S. H. W. Scheres. A bayesian view on cryo-EM structure determination. *Journal of Structural Biology*, 415(2):406–418, 2012.
- [Sch12b] S. H. W. Scheres. RELION: Implementation of a bayesian approach to cryo-EM structure determination. *Journal of Structural Biology*, 180(3):519–530, 2012.
- [Sch16] S. H. W. Scheres. Processing of structurally heterogeneous cryo-EM data in RELION. *Methods in Enzymology*, 579:125–157, 2016.
- [SG92] B. M. Sadler and G. B. Giannakis. Shift- and rotation-invariant object reconstruction using the bispectrum. *Oct. Soc. Am. A*, 9:57–69, 1992.
- [Sin11] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30(1):20 – 36, 2011.
- [SSK13] B. Sontag, A. Singer, and I. G. Kevrekidis. Noisy dynamic simulations in the presence of symmetry: Data alignment and model reduction. *Computers & Mathematics with Applications*, 65(10):1535 – 1557, 2013.
- [ST04] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- [TS12] D. L. Theobald and P. A. Steindel. Optimal simultaneous superpositioning of multiple structures with missing data. *Bioinformatics*, 28(15):1972–1979, 2012.
- [Tuk84] J. W. Tukey. The spectral representation and transformation properties of the higher moments of stationary time series. In D. R. Brillinger, editor, *The Collected Works of John W. Tukey*, volume 1, chapter 4, pages 165–184. Wadsworth, 1984.
- [ZvdHGG03] J. P. Zwart, R. van der Heiden, S. Gelsema, and F. Groen. Fast translation invariant classification of HRR range profiles in a zero phase representation. *Radar, Sonar and Navigation, IEE Proceedings*, 150(6):411–418, 2003.

AMELIA PERRY
DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE,
CAMBRIDGE, MA 02139-4307, USA
(ameliaperry@mit.edu)

AFONSO S. BANDEIRA
DEPARTMENT OF MATHEMATICS
COURANT INSTITUTE OF MATHEMATICAL SCIENCES
CENTER FOR DATA SCIENCE
NEW YORK UNIVERSITY,
NEW YORK, NY 10012, USA
(bandeira@cims.nyu.edu)

AMIT SINGER
DEPARTMENT OF MATHEMATICS
PROGRAM IN APPLIED AND COMPUTATIONAL MATHEMATICS
PRINCETON UNIVERSITY,
PRINCETON NJ, 08544, USA
(amits@princeton.edu)

JONATHAN WEED
DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE,
CAMBRIDGE, MA 02139-4307, USA
(jweed@mit.edu)

PHILIPPE RIGOLLET
DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE,
CAMBRIDGE, MA 02139-4307, USA
(rigollet@math.mit.edu)