

Bispectrum Inversion With Application to Multireference Alignment

Tamir Bendory¹, Nicolas Boumal², Chao Ma, Zhizhen Zhao³, and Amit Singer

Abstract—We consider the problem of estimating a signal from noisy circularly translated versions of itself, called *multireference alignment* (MRA). One natural approach to MRA could be to estimate the shifts of the observations first, and infer the signal by aligning and averaging the data. In contrast, we consider a method based on estimating the signal directly, using features of the signal that are invariant under translations. Specifically, we estimate the power spectrum and the bispectrum of the signal from the observations. Under mild assumptions, these invariant features contain enough information to infer the signal. In particular, the bispectrum can be used to estimate the Fourier phases. To this end, we propose and analyze a few algorithms. Our main methods consist of nonconvex optimization over the smooth manifold of phases. Empirically, in the absence of noise, these nonconvex algorithms appear to converge to the target signal with random initialization. The algorithms are also robust to noise. We then suggest three additional methods. These methods are based on frequency marching, semidefinite relaxation, and integer programming. The first two methods provably recover the phases exactly in the absence of noise. In the high noise level regime, the invariant features approach for MRA results in stable estimation if the number of measurements scales like the cube of the noise variance, which is the information-theoretic rate. Additionally, it requires only one pass over the data, which is important at low signal-to-noise ratio when the number of observations must be large.

Index Terms—Bispectrum, multireference alignment, phase retrieval, non-convex optimization, optimization on manifolds,

semidefinite relaxation, phase synchronization, frequency marching, integer programming, cryo-EM.

I. INTRODUCTION

WE CONSIDER the problem of estimating a discrete signal from multiple noisy and translated (i.e., circularly shifted) versions of itself, called *multireference alignment* (MRA). This problem occurs in a variety of applications in biology [1]–[4], radar [5], [6], image registration and super-resolution [7]–[9], and has been the subject of recent theoretical analysis [10], [11]. The MRA model reads

$$\xi_j = R_{r_j} x + \varepsilon_j, \quad j = 1, \dots, M, \quad (1)$$

where ε_j are i.i.d. normal random vectors with variance σ^2 and the underlying signal x is in $\mathbb{R}^N$ or in $\mathbb{C}^N$. Operator R_{r_j} rotates the signal x circularly by r_j locations, namely, $(R_{r_j} x)[n] = x[n - r_j]$, where indexing is zero-based and considered modulo N (throughout the paper). While both x and the translations $\{r_j\}$ are unknown, we stress that the goal here is merely to estimate x . This estimation is possible only up to an arbitrary translation.

A chief motivation for this work arises from the imaging technique called single particle Cryo-Electron Microscopy (Cryo-EM), which allows to visualize molecules at near-atomic resolution [12], [13]. In Cryo-EM, we aim to estimate a three dimensional (3D) object from its two-dimensional (2D) noisy projections, taken at unknown viewing directions [14], [15]. While typically the recovery process involves alignment of multiple observations in a low signal-to-noise ratio (SNR) regime, the underlying goal is merely to estimate the 3D object. In this manner, with the unknown shifts corresponding to the unknown viewing directions, MRA can be understood as a simplified model for Cryo-EM.

Existing approaches for MRA can be classified into two main categories. The first class of methods aims to estimate the set of translations $\{r_j\}$ first. Given this set, estimating x can be achieved easily by aligning all observations ξ_j and then averaging to reduce the noise. The second class, which we favor in this paper, consists of methods which aim to estimate the signal directly, without estimating the shifts.

Considering the first class, one intuitive approach to estimating the translations is to fix a template observation, say ξ_1 , and to estimate the relative translations by cross-correlation. This is called template matching. Specifically, r_j is estimated as

$$\hat{r}_j = \arg \max_k \Re \left\{ \sum_{n=0}^{N-1} \xi_1[n] \overline{\xi_j[n+k]} \right\}, \quad j = 2, \dots, M,$$

Manuscript received May 16, 2017; revised October 6, 2017; accepted November 6, 2017. Date of publication November 20, 2017; date of current version January 16, 2018. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Antti Tolli. The authors were supported in part by the National Institute of General Medical Sciences under Award Number R01GM090200, in part by the Air Force Office of Scientific Research under Grant FA9550-17-1-0291, in part by Simons Investigator Award and Simons Collaboration on Algorithms and Geometry from Simons Foundation, and in part by the Moore Foundation Data-Driven Discovery Investigator Award. The work of N. Boumal was supported by the National Science Foundation under Grant DMS-1719558. The work of Z. Zhao was supported in part by the National Center for Supercomputing Applications Faculty Fellowship and in part by the University of Illinois at Urbana-Champaign College of Engineering Strategic Research Initiative. (T. Bendory and N. Boumal contributed equally to this work.) (Corresponding author: Tamir Bendory.)

T. Bendory, N. Boumal, C. Ma, and A. Singer are with the Program in Applied and Computational Mathematics and the Mathematics Department, Princeton University, Princeton, NJ 08544 USA (e-mail: tamir.bendory@princeton.edu; nboumal@math.princeton.edu; chaom@princeton.edu; amits@math.princeton.edu).

Z. Zhao is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: zhizhenz@illinois.edu).

This paper has supplementary downloadable multimedia material available at <http://ieeexplore.ieee.org> provided by the authors. This includes an appendix to the paper. This material is 208 KB in size.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2017.2775591

TABLE I
COMPARISON OF MAIN MRA APPROACHES

Method	Computational complexity	Storage requirement	Comments
Template alignment	$O(MN \log N)$	$O(N)$	Fails at moderate SNR (see Figure I.1)
Angular synchronization	$O(M^2 N \log N)$	$O(M^2)$	Fails at low SNR
Expectation maximization	$O(TMN \log N)$	$O(MN)$	Empirically accurate; #iterations T grows with noise level
Invariant features (this paper)	$O(MN^2 + F(N))$	$O(N^2)$	Under mild conditions, accurate estimation if M grows as σ^6

$F(N)$ denotes the complexity of inverting the bispectrum. For instance, for the FM algorithm, $F(N) = O(N^2)$. Storage requirements include the possibility of streaming computations where possible.




Fig. 1. Alignment of two translated versions of the same signal in the presence of i.i.d. Gaussian noise with various standard deviations σ . The true signal in $\mathbb{R}^{100}$ is a window of length 22 and height 1. Each row presents two observations and their cross-correlation. Importantly, beyond a certain threshold, noise makes pairwise alignment impossible.

where $\Re\{z\}$ and $\bar{z}$ denote the real part and the conjugate of a complex number z . This approach requires only one pass over the data: for each observation, the best shift can be computed in $O(N \log N)$ flops, and the aligned observations can be averaged online. This results in a total computational cost of $O(MN \log N)$ flops, see Table I. While this approach is simple and efficient, it necessarily fails below a critical SNR—see Fig. 1 for a representative example.

The issue with template matching is that we rely on aligning each observation to only one template: this is error prone at low SNR. Instead, to derive a more robust estimator, one can look for the most suitable alignment among all pairs of observations. The M^2 relative shifts thus computed must then be reconciled into a compatible choice of M shifts for the individual observations. This is a discrete version of the *angular synchronization* problem, see [16]–[21]. The computational complexity of aligning all pairs individually is $O(M^2 N \log N)$, while storing the results uses $O(M^2)$ memory.

Alternative algorithms for estimating the translations are based on different SDP relaxations [22], [23], iterative template alignment [24], zero phase representations [5] and neural networks [6]. The statistical limits of alignment tasks were derived for a variety of setups and noise models, see for instance [25]–[28]. For example, for a continuous, 2D version of the MRA model, it was shown that the Cramér–Rao lower bound

(CRLB) for translation estimation is proportional to the noise variance σ^2 [25]; crucially, it does not improve with M , even if the underlying signal is known. This is motivation to consider the second category of MRA methods, where shifts are not estimated.

Section VI elaborates on expectation maximization (EM) which tries to compute the maximum marginalized likelihood estimator (MMLE) of the signal—marginalization is done over the shifts. This method acknowledges the difficulty of alignment by working not with estimates of the shifts themselves, but rather with estimates of the probability distributions of the shifts. As a result, EM achieves excellent numerical performance in practice. However its computational complexity is high and its performance is not understood in theory.

It has been shown recently that the sample complexity of MRA, under assumption that shifts are distributed uniformly, is proportional to σ^6 in the low SNR regime. In other words, the number of measurements M needs to scale like σ^6 to retain a constant estimation error [11].

In this work, we propose a framework which achieves this sample complexity by estimating the sought signal x directly using features that are invariant under translations. For instance, the mean of x is invariant under translation and can be estimated easily from the mean of all observations. We further use the power spectrum and the bispectrum of the observations—

which are Fourier-transform based invariants—to estimate the magnitudes and phases of the signal’s Fourier transform, respectively.

For any fixed noise level (which may be arbitrarily large), these features can be estimated accurately provided sufficiently many measurements are available. Hence, our approach allows to deal with *any noise level*. Besides achieving the sample complexity, the computational complexity and memory requirements of the methods we describe are relatively low. Indeed, the only operations whose computational cost grows with M are computations of averages over the data. These can be performed on-the-fly and are easily parallelizable. We mention that a recent tensor decomposition algorithm also achieves this estimation rate [29].

Given estimators for the mean and power spectrum of x , estimating the DC component and Fourier magnitudes of x is straightforward. In this paper, we thus focus on the task of recovering the Fourier phases of x from an estimator of its bispectrum. We propose two non-convex optimization algorithms on the manifold of phases for this task, which we call bispectrum inversion. We also discuss three additional algorithms which do not require initialization (and hence could be used to initialize others), based on frequency marching, SDP relaxation and integer programming. The first two methods recover the phases exactly in the absence of noise.

Beyond MRA, the bispectrum plays a central role in a variety of signal processing applications. For instance, it is a key tool to separate Gaussian and non-Gaussian processes [30], [31]. It is also used to investigate the cosmic background radiation [32], [33], seismic signal processing [34], image deblurring [35], feature extraction for radar [36], analysis of EEG signals [37], MIMO systems [38] and classification [39] (see also [40]–[44] and references therein). In Section III, we review previous works on bispectrum inversion [34], [45], [46]. Reliable algorithms to invert the bispectrum, as studied here, may prove useful in some of these applications.

The paper is organized as follows. Section II discusses the invariant feature approach for MRA. Section III presents the non-convex algorithms on the manifold of phases for bispectrum inversion. Section IV is devoted to additional algorithms that can be used to initialize the non-convex algorithms. Section V analyzes one of the proposed non-convex algorithm. Section VI elaborates on the EM approach for MRA, Section VII shows numerical experiments and Section VIII offers conclusions and perspective.

Throughout the paper we use the following notation. Vectors x in $\mathbb{R}^N$ or $\mathbb{C}^N$ and $y \in \mathbb{C}^N$ denote the underlying signal and its discrete Fourier transform (DFT), respectively. In the sequel, all indices are understood modulo N , namely, in the range $0, \dots, N-1$. The phase of a complex scalar a , defined as $a/|a|$ if $a \neq 0$ and zero otherwise, is denoted by $\text{phase}(a)$ or $\tilde{a}$. The conjugate-transpose of a vector z is denoted by z^* . We use $\circ$ to denote the Hadamard (entry-wise) product, $\mathbb{E}$ for expectation, $\text{Tr}(Z)$ for the trace and $\|Z\|_F$ for the Frobenius norm of a matrix Z . We reserve $T(z)$ for circulant matrices determined by their first row z , i.e., $T(z)[k_1, k_2] = z[k_2 - k_1]$, and $\mathcal{H}^N$ for the set of Hermitian matrices of size $N \times N$.

II. MULTIREFERENCE ALIGNMENT VIA INVARIANT FEATURES

We propose to solve the MRA problem directly using features that are invariant under translations. Unlike pairwise alignment, this approach fuses information from all M observations together—not just of pairs—and it only aims to recover the signal itself—not the translations. The essence of this idea was discussed as a possible extension in [22, Appendix A]. The invariant features can be understood either as auto-correlation functions or as their Fourier transform. In this work, we make use of the first three invariants defined as

$$\begin{aligned} c_1 &= \mu_x = \frac{1}{N} \sum_{n=0}^{N-1} x[n], \\ c_2[n_1] &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] \overline{x[n - n_1]}, \\ c_3[n_1, n_2] &= \frac{1}{N} \sum_{n=0}^{N-1} x[n] \overline{x[n - n_1]} x[n + n_2], \end{aligned} \quad (2)$$

for $n_1, n_2 = 0, \dots, N-1$. It is clear that c_1, c_2, c_3 are invariant under circular shifts of x . For higher-order invariants based on auto-correlations, see for instance [47].

The first feature is the mean of the signal which is the auto-correlation function of order one (i.e., c_1 in (2)). The distribution of the mean of ξ_j is then given by $\mu_{\xi_j} \sim \mathcal{N}(\mu_x, \frac{\sigma^2}{N})$ and we can estimate μ_x as

$$\hat{\mu}_x = \frac{1}{M} \sum_{j=1}^M \left(\frac{1}{N} \sum_{n=0}^{N-1} \xi_j[n] \right) \sim \mathcal{N} \left(\mu_x, \frac{\sigma^2}{NM} \right). \quad (3)$$

Estimating the signal’s mean supplies only limited information about the signal itself. Thus, we consider also the auto-correlation function of order two (i.e., c_2 in (2)). Its Fourier transform, the power spectrum, is explicitly defined as

$$P_x[k] = |y[k]|^2,$$

for all k , where y is the DFT of x . An alternative way to understand the invariance of the power spectrum under shifts is through the effect of shifts on the DFT of a signal:

$$\text{DFT}(R_s x)[k] = y[k] \cdot e^{-2\pi i k s / N}. \quad (4)$$

Thus, shifts only affect the phases of the DFT, so that $P_{R_s x} = P_x$ for any shift R_s . Furthermore, owing to independence of the noise with respect to the signal itself and to the shift,

$$\mathbb{E} \{ P_{\xi_j}[k] \} = P_x[k] + N\sigma^2,$$

where the second term is the power spectrum of the noise ε_j . Therefore, we estimate the power spectrum of x as:

$$\hat{P}_x[k] = \frac{1}{M} \sum_{j=1}^M (P_{\xi_j}[k] - N\sigma^2). \quad (5)$$

It can be shown that $\hat{P}_x$ is unbiased and its variance is dominated by $\frac{\sigma^4}{M}$ for large σ . Hence, $\hat{P}_x \rightarrow P_x$ as $M \rightarrow \infty$. In particular,

accurate estimation of the power spectrum requires M to scale like σ^4 . In the sequel, we assume that σ is known.¹

Recovering a signal from its power spectrum is commonly referred to as *phase retrieval*. This problem received considerable attention in recent years, see for instance [48]–[54]. It is well known that almost no one-dimensional signal can be determined uniquely from its power spectrum. Therefore, we use the power spectrum merely to estimate the signal's Fourier magnitudes. As explained next, we use the auto-correlation of third order and its Fourier transform, the bispectrum, to estimate the Fourier phases.

Since phase retrieval is in general ill posed, we use the auto-correlation function of order three (that is, c_3 in (2)) through its Fourier transform, the bispectrum, to estimate the Fourier phases of the sought signal. The bispectrum is a function of two frequencies $k_1, k_2 = 0, \dots, N-1$ and is defined as [55]:

$$B_x[k_1, k_2] = y[k_1] \overline{y[k_2]} y[k_2 - k_1]. \quad (6)$$

Note that, if $y[0] \neq 0$, the power spectrum is explicitly included in the bispectrum since $P_x[k] = B_x[k, k]/y[0]$. The fact that the bispectrum is invariant under shifts can also be deduced from (4). Indeed, for any shift R_s ,

$$\begin{aligned} B_{R_s x}[k_1, k_2] &= \left(y[k_1] e^{-2\pi i k_1 s/N} \right) \left(\overline{y[k_2]} e^{2\pi i k_2 s/N} \right) \\ &\quad \cdot \left(y[k_2 - k_1] e^{2\pi i (k_1 - k_2) s/N} \right) \\ &= B_x[k_1, k_2]. \end{aligned}$$

In matrix notation, we express this as

$$B_x = y y^* \circ T(y), \quad (7)$$

where $T(y)$ is a circulant matrix whose first row is y , that is, $\overline{T(y)[k_1, k_2]} = y[k_2 - k_1]$. Observe that if x is real, then $y[k] = \overline{y[-k]}$ so that $T(y)$ and B_x are Hermitian matrices. Simple expectation calculations lead to the conclusion that

$$\mathbb{E} \{ B_{\xi_j} \} = B_x + \sigma^2 N^2 \mu_x A, \quad (8)$$

where $A = A_{\mathbb{R}}$ or $A = A_{\mathbb{C}}$ depending on $x \in \mathbb{R}^N$ or $x \in \mathbb{C}^N$ and

$$A_{\mathbb{R}} = \begin{bmatrix} 3 & 1 & 1 & 1 & \dots & 1 \\ 1 & 1 & 0 & 0 & \dots & 0 \\ 1 & 0 & 1 & 0 & \dots & 0 \\ 1 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & \dots & 1 \end{bmatrix}, \quad A_{\mathbb{C}} = \begin{bmatrix} 2 & 1 & 1 & 1 & \dots & 1 \\ 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \end{bmatrix}.$$

Since the bias term is proportional to μ_x , we propose to estimate $\hat{B}_{x-\mu_x}$ by averaging over $B_{\xi_j - \mu_x}$ for all j . This estimator is unbiased and its variance is controlled by $\frac{\sigma^6}{M}$ for large σ . Therefore, M is required to scale like σ^6 to ensure accurate estimation. In practice, μ_x is not known exactly. Thus, we estimate

the bispectrum by

$$\hat{B}_{x-\mu_x} = \frac{1}{M} \sum_{j=1}^M B_{\xi_j - \hat{\mu}_x}, \quad (9)$$

which is asymptotically unbiased. For finite M and large σ , bias induced by the approximation $\hat{\mu}_x \approx \mu_x$ is significantly smaller than the standard deviation of (9).

The bispectrum contains information about the Fourier phases of x because, defining $\tilde{y}[k] = \text{phase}(y[k])$ and $\tilde{B}_x[k_1, k_2] = \text{phase}(B_x[k_1, k_2])$ where phase extracts the phase of a complex number (and returns 0 if that number is 0), we have

$$\tilde{B}_x[k_1, k_2] = \tilde{y}[k_1] \overline{\tilde{y}[k_2]} \tilde{y}[k_2 - k_1]. \quad (10)$$

In matrix notation, the normalized bispectrum takes the form $\tilde{B}_x = \tilde{y} \tilde{y}^* \circ T(\tilde{y})$.

Contrary to the power spectrum, the bispectrum is usually invertible. Indeed, in the absence of noise, the bispectrum determines the sought signal uniquely under moderate conditions:

Proposition II.1: For $N \geq 5$, let $x \in \mathbb{C}^N$ be a signal whose DFT y obeys $y[k] \neq 0$ for $k = 1, \dots, K$, possibly also for $k = 0$, and zero otherwise. Up to integer time shifts, x is determined exactly by its bispectrum provided $K \geq \frac{N+1}{2}$.

For $N \geq 5$, let $x \in \mathbb{R}^N$ be a real signal whose DFT y obeys $y[k] \neq 0$ for $k = 1, \dots, K$ and $k = N-1, \dots, N-K$, possibly also for $k = 0$, and zero otherwise. Up to integer time shifts, x is determined exactly by its bispectrum provided $\frac{N}{3} \leq K \leq \frac{N-1}{2}$.

Proof: This is a direct corollary of Lemmas V.1 and B.1. ■

We stress that the bispectrum estimator in (9) is not a bispectrum itself, since the set of bispectra is not a linear space: $\hat{B}_{x-\mu_x}$ is not invertible as such [42]. Algorithms we propose aim to find a *stable* inverse, in the sense that the recovered signal will have a bispectrum which is close to the estimated bispectrum in $\mathbb{C}^{N \times N}$. The following propositions combined argue formally that this can be done in the MRA model. The proofs in Appendix A are constructive.

Proposition II.2 (Stable bispectrum inversion): There exists an estimator $\hat{x}$ with the following property. For any signal x in $\mathbb{R}^N$ or $\mathbb{C}^N$ whose DFT is non-vanishing, there exist a precision $\delta = \delta(x) > 0$ and a sensitivity $L = L(x) < \infty$ such that if an estimator $\hat{B}_x$ of B_x satisfies $\|\hat{B}_x - B_x\|_F \leq \delta$, then $\hat{x} = \hat{x}(\hat{B}_x)$ satisfies $\min_{r=0 \dots N-1} \|x - R_r \hat{x}\|_2 \leq L \|\hat{B}_x - B_x\|_F$.

Proposition II.3 (Bispectrum estimation): For any signal x in $\mathbb{R}^N$ or $\mathbb{C}^N$ whose DFT is non-vanishing, for any required precision $\delta > 0$ and for any probability $p < 1$, there exists a constant $C = C(x, p, \delta) < \infty$ such that, for any noise level $\sigma > 0$, if the number of observations M exceeds $C \cdot (\sigma^2 + \sigma^6)$, the estimator

$$\hat{B}_x = \frac{1}{M} \sum_{j=1}^M B_{\xi_j} - \sigma^2 N^2 \hat{\mu}_x A$$

satisfies $\|\hat{B}_x - B_x\|_F \leq \delta$ with probability at least p .

We mention that uniqueness in the continuous setup was considered in [56]. The more general setting of bispectrum over compact groups was considered in [57]–[60].

¹If σ is not known, it can be estimated from the data as

$$\hat{\sigma}^2 = \frac{1}{N} \cdot \text{variance} \left(\sum_{n=0}^{N-1} \xi_j[n] \right)_{j=1, \dots, M}.$$




Fig. 2. Relative error of estimating the power spectrum and bispectrum for different noise levels as a function of the number of observations M . Results are averaged over 10 repetitions for each value of M on a fixed real signal of length $N = 41$ with i.i.d. normal random entries. The signal-to-noise ratio is then $1/\sigma^2$. Importantly, the relative error decreases as $1/\sqrt{M}$ regardless of noise level.

Algorithm 1: Outline of the invariant approach for MRA.

Input: Set of observations ξ_j , $j = 1, \dots, M$ according to (1) and noise level σ

Output: $\hat{x}$: estimation of x

Estimate invariant features:

- 1) Compute $\hat{\mu}_x$ according to (3)
- 2) Compute $\hat{P}_x$ according to (5)
- 3) Compute $\hat{B}_{x-\mu_x}$ according to (9)

Estimate the signal's DFT:

- 1) Estimate $y[0]$ from $\hat{\mu}_x$. For other frequencies:
- 2) Estimate the magnitudes of y from $\hat{P}_x$
- 3) Estimate the phases of y from $\hat{B}_{x-\mu_x}$ (e.g., Algorithm 2)

Return: $\hat{x}$: inverse DFT of the estimated y

The MRA model here assumes i.i.d. Gaussian noise. However, the estimation is performed by averaging in the bispectrum domain, where noise affecting individual entries is correlated. Consequently, one may want to use a more robust estimator, such as the median. Yet, computing the median of complex matrices is computationally expensive, while computing the average can be performed efficiently and on-the-fly, that is, without requiring to store all observations. For Gaussian noise, we have noticed numerically that using the mean or the median for bispectrum estimation leads to comparable estimation errors (experiments not shown). In other noise models, e.g., with outliers, it might be useful to consider the median or the median of means method, see for instance [61].

Fig. 2 presents the relative estimation error of the power spectrum and bispectrum as a function of the number of observations M . For the bispectrum, the relative error is computed as

$$\text{relative error} := \frac{\|B_{x-\mu_x} - \frac{1}{M} \sum_{j=1}^M B_{\xi_j - \hat{\mu}_x}\|_F}{\|B_{x-\mu_x}\|_F},$$

and similarly for the power spectrum. As expected, the slope of all curves is approximately $1/2$ in logarithmic scale, implying that the estimation error decreases as $O(1/\sqrt{M})$. The invariant features approach for MRA is summarized in Algorithm 1.

Consider the case in which the number of samples M may be very large whereas the size of the object is fixed, namely $N \ll M$. This case is of interest in many applications, such as

cryo-EM [14], [15]. In this regime, the invariant features approach has two important advantages over methods that rely on estimating the translations. First, in the invariant features approach, we average over the M observations (which is computationally cheap), and then apply a more complex algorithm (say, to recover a signal from its bispectrum) whose input size is a function of N but is independent of M . Hence, the overall complexity of this approach can be relatively low. Second, the alignment-based method requires storing all M observations, namely, MN samples, which is unnecessary in the invariant features approach. There, for each observation, we just need to compute its invariants, to be averaged over all observations: this can be done online (in streaming mode) and in parallel.

III. NON-CONVEX ALGORITHMS FOR BISPECTRUM INVERSION

After estimating the first Fourier coefficient $y[0]$ as $N\hat{\mu}_x$, our approach for MRA by invariant features consists of two parts. We use the power spectrum to estimate the signal's Fourier magnitudes and the bispectrum for the phases. The first part is straightforward: $|y[k]|$ can be estimated as $\sqrt{\hat{P}_x[k]}$ if $\hat{P}_x[k] \geq 0$, and as 0 otherwise. Hereafter, we focus on estimation of the phases of the DFT, $\tilde{y}$.

In the literature, two main approaches were suggested to invert the discrete bispectrum. The first is based on estimating the frequencies one after the other by exploiting simple algebraic relations [45], [46]. The second approach suggests to estimate the signal by least-squares solution and phase unwrapping [34], [46]. We improve these methods and suggest a few new algorithms. The algorithms are split into two sections. This section is devoted to two new non-convex algorithms based on optimization on the manifold of phases. Both of these algorithms require initialization. While experimentally it appears that random initialization works well, for completeness, in the next section we propose three additional algorithms which do not need initialization and hence could be used to initialize the non-convex algorithms.

A. Local Non-convex Algorithm Over the Manifold of Phases

In this section, similarly to (10), we let $\tilde{B}$ denote our estimate of the phases of B_x . Since $\tilde{B} \approx \tilde{y}\tilde{y}^* \circ T(\tilde{y})$, one way to model

recovery of the Fourier phases $\tilde{y}$ is by means of the non-convex least-squares optimization problem

$$\min_{z \in \mathbb{C}^N} \left\| W \circ (\tilde{B} - zz^* \circ T(z)) \right\|_F^2 \quad \text{subject to } |z[k]| = 1, \forall k. \quad (11)$$

The matrix $W \in \mathbb{R}^{N \times N}$ is a weight matrix with nonnegative entries. These weights can be used to indicate our confidence in each entry of $\tilde{B}$. Expanding the squared Frobenius norm yields

$$\begin{aligned} \left\| W \circ (\tilde{B} - zz^* \circ T(z)) \right\|_F^2 &= \|W \circ \tilde{B}\|_F^2 \\ &+ \|W \circ zz^* \circ T(z)\|_F^2 - 2 \langle W \circ \tilde{B}, W \circ zz^* \circ T(z) \rangle, \end{aligned}$$

where

$$\langle U, V \rangle = \Re\{\text{Tr}(U^* V)\}, \quad (12)$$

is the real inner product associated to the Frobenius norm. Under the constraints on z , the first two terms are constant and the inner product term is equivalent to

$$\langle W \circ \tilde{B}, W \circ zz^* \circ T(z) \rangle = \langle z, M(z)z \rangle,$$

with

$$M(z) := W^{(2)} \circ \tilde{B} \circ \overline{T(z)}, \quad (13)$$

where we use the notation $W^{(2)} := W \circ W$. One possibility is to choose $W = \sqrt{|\hat{B}_{x-\mu_x}|}$, where the absolute value and the square root are taken entry-wise, so that $M(z) = \hat{B}_{x-\mu_x} \circ \overline{T(z)}$. Hence, optimization problem (11) is equivalent to

$$\max_{z \in \mathbb{C}^N} f(z) = \langle z, M(z)z \rangle \quad \text{subject to } |z[k]| = 1, \forall k. \quad (14)$$

We can also impose $z[0] = \text{phase}(\hat{\mu}_x)$. If x is real, we have the additional symmetry constraints $z[k] = \overline{z[-k]}$.

Since the cost function f is continuous and the search space is compact, a solution exists. Of course, the solution is not unique, in accordance with the invariance of the bispectrum under integer time-shifts of the underlying discrete signal. This is apparent through the fact that the cost function f is invariant under the corresponding (discrete) transformations of z . This is true independently of the data $\tilde{B}$ and W . The proof is in Appendix C.

Lemma III.1: The cost function f is invariant under transformations of z that correspond to integer time-shifts of the underlying signal.

To solve this non-convex program, we use the Riemannian trust-region method (RTR) [62], whose usage is simplified by the toolbox Manopt [63]. RTR enjoys global convergence to second-order critical points, that is, points which satisfy first- and second-order necessary optimality conditions [64] and a quadratic local convergence rate. Empirically, in the noiseless case it appears that the algorithm recovers the target signal with *random initialization*, all local minima are global (with minor technicality for even N in the real case) and all second-order critical points have an escape direction, that is, saddles are “strict”. Numerical experiments demonstrate reasonable robustness in the face of noise. This algorithm is summarized in Algorithm 2 and studied in detail in Section V.

Algorithm 2: Non-convex optimization on phase manifold.

Input: The normalized bispectrum $\tilde{B}[k_1, k_2]$ and a weight matrix $W \in \mathbb{R}^{N \times N}$

Output: $\hat{y}$: an estimation of $\tilde{y}$

Compute: Using RTR [62], [63], approximately solve:

$$\hat{y} = \arg \max_{z \in \mathbb{C}^N} \Re\{z^* M(z)z\} \quad \text{subject to } |z[k]| = 1, \forall k, \\ \text{(if } x \text{ is real) } z[k] = \overline{z[-k]}, \forall k,$$

$$\text{where } M(z) := (W \circ W) \circ \tilde{B} \circ \overline{T(z)}.$$

B. Iterative Phase Synchronization Algorithm

In this section we present an alternative heuristic to the non-convex algorithm on the manifold of phases. This algorithm is based on iteratively solving the phase synchronization problem. Suppose we get an estimation of $\tilde{y}$, say $\hat{y}_{k-1}$. If $\hat{y}_{k-1} \approx \tilde{y}$ is non-vanishing, then this estimation should approximately satisfy the bispectrum relation:

$$\tilde{B} \circ \overline{T(\hat{y}_{k-1})} \approx \hat{y}_{k-1} \hat{y}_{k-1}^*.$$

The underlying idea is now to push the current estimation towards $\tilde{y}$ by finding a rank-one approximation of $\tilde{B} \circ \overline{T(\hat{y}_{k-1})}$ with unit modulus entries. This problem can be formulated as:

$$\arg \max_{z \in \mathbb{C}^N} \Re\left\{z^* \left(\tilde{B} \circ \overline{T(\hat{y}_{k-1})}\right) z\right\} \quad \text{subject to } |z[\ell]| = 1, \forall \ell, \quad (15)$$

where we treat the matrix $\tilde{B} \circ \overline{T(\hat{y}_{k-1})}$ as a constant. This problem is called phase synchronization. Many algorithms have been suggested to solve the phase synchronization problem. Among them are the eigenvector method, SDP relaxation, projected power method, Riemannian optimization and approximate message passing [16]–[20]. Notice that the solution of (15) is only defined up to a global phase, namely, if z is optimal, then so is $ze^{i\phi}$ for any angle ϕ . To resolve this ambiguity, we require knowledge of the phase of the mean, $\tilde{y}[0]$ (which is easy to estimate from the data) and we pick the global phase of $\hat{y}_k$ such that $\hat{y}_k[0] = \tilde{y}[0]$.

The k th iteration of our algorithm thus (tries to) solve the phase synchronization problem with respect to the matrix $M_{k-1} := \tilde{B} \circ \overline{T(\hat{y}_{k-1})}$, where $\hat{y}_{k-1}$ is the solution of the previous estimation. Assuming the signal is real, we also impose at each iteration the conjugate-reflection property of $\hat{y}_k[\ell] = \hat{y}_k[-\ell]$ for all ℓ so that M_k is Hermitian. In the numerical experiments in Section VII, we solve (15) by the Riemannian trust-region method described in [17]. Empirically, the performance of this algorithm and Algorithm 2 is indistinguishable. The algorithm is summarized in Algorithm 3.

IV. INITIALIZATION-FREE ALGORITHMS

The previous section was devoted to non-convex algorithms to invert the bispectrum. In this section we present three additional algorithms based on frequency marching (FM), SDP relaxation and phase unwrapping. These algorithms do not require initialization and therefore could be used to initialize the non-convex algorithms.

Algorithm 3: Iterative phase synchronization algorithm.

Input: The normalized bispectrum $\tilde{B}$, initial estimation $\hat{y}_0$, phase of the mean $\tilde{y}[0]$

Output: $\hat{y}$: estimation of $\tilde{y}$

Set $k = 0$

while stopping criterion does not trigger **do**:

- $k \leftarrow k + 1$

- Compute $\hat{y}_k$ as a solution of (15)

- Fix the global phase: $\hat{y}_k \leftarrow \hat{y}_k \cdot \frac{\tilde{y}[0]}{\hat{y}_k[0]}$

- If x is real, symmetrize: $\hat{y}_k \leftarrow \text{phase}((\hat{y}_k)_{\downarrow\uparrow})$, see (B.1)

end while

Return: $\hat{y} \leftarrow \hat{y}_k$

We prove that FM and the SDP recover the Fourier phases exactly in the absence of noise under the assumption that we can fix $\tilde{y}[1]$. If the signal has non-vanishing DFT, $\tilde{y}[1]$ can be estimated from the bispectrum using the fact that $\tilde{y}[1]^N$ equals

$$\text{phase}(B_x[N-1, 1]B_x[1, 2] \cdot B_x[1, 2] \cdots B_x[1, N-1]).$$

(Any N th root can be used for $\tilde{y}[1]$, corresponding to the N possible shifts of x .) In all cases, we argue that forcing $\tilde{y}[1] = 1$ is acceptable if N is large. Indeed, recall that a shift by ℓ entries in the space domain is equivalent to modulating the k th Fourier coefficient by $e^{-2\pi i \ell k / N}$. In particular, it means that the phase $\tilde{y}[1]$ can be shifted by $e^{-2\pi i \ell / N}$ for an arbitrary $\ell \in \mathbb{Z}$. Thus, for signals of length $N \gg 1$, the phase $\tilde{y}[1]$ can be set arbitrarily with only small error. In the numerical experiments of Section VII, we give the correct value of $\tilde{y}[1]$ to the algorithms in order to assess their best possible behavior.

We begin by discussing the FM algorithm, which is a simple propagation method: it is exact in the absence of noise. Notwithstanding, its estimation for the low-frequency coefficients is sensitive to noise. Because of its recursive nature, error in the low frequencies propagates to the high frequencies, resulting in unreliable estimation. The other two algorithms are more computationally demanding but appear more robust.

A. Frequency Marching Algorithm

The FM algorithm is a simple propagation algorithm in the spirit of [45], [46] that aims to estimate $\tilde{y}$ one frequency at a time. This algorithm has computational complexity $O(N^2)$ and it recovers $\tilde{y}$ exactly for both real and complex signals in the absence of noise, assuming $\tilde{y}[1]$ is known.

Let us denote $\tilde{B}[k_1, k_2] = e^{i\psi[k_1, k_2]}$ and $\tilde{y}[k] = e^{i\psi[k]}$. Accordingly, we can reformulate (6) as

$$\Psi[k_1, k_2] = \psi[k_1] - \psi[k_2] + \psi[k_2 - k_1] \bmod 2\pi,$$

where the modulo is taken over the sum of all three terms. Using this relation, we can start to estimate the missing phases. The first unknown phase, $\psi[2]$, can be estimated by:

$$\begin{aligned} \Psi[1, 2] &= \psi[1] - \psi[2] + \psi[1] \bmod 2\pi \\ \Rightarrow \hat{\psi}[2] &= 2\psi[1] - \Psi[1, 2] \bmod 2\pi, \end{aligned}$$

where $\hat{\psi}[2]$ refers to the estimator of $\psi[2]$ (defined modulo 2π). We can estimate the next phase in the same manner:

$$\begin{aligned} \Psi[1, 3] &= \psi[1] - \psi[3] + \psi[2] \bmod 2\pi \\ \Rightarrow \hat{\psi}[3] &= \psi[1] + \hat{\psi}[2] - \Psi[1, 3] \bmod 2\pi. \end{aligned}$$

For higher frequencies, we have more measurements to rely on. For the fourth entry, we now can derive two estimators as follows:

$$\begin{aligned} \Psi[1, 4] &= \psi[1] - \psi[4] + \psi[3] \bmod 2\pi \\ \Rightarrow \hat{\psi}^{(1)}[4] &= \psi[1] + \hat{\psi}[3] - \Psi[1, 4] \bmod 2\pi, \end{aligned}$$

and

$$\begin{aligned} \Psi[2, 4] &= \psi[2] - \psi[4] + \psi[2] \bmod 2\pi \\ \Rightarrow \hat{\psi}^{(2)}[4] &= 2\hat{\psi}[2] - \Psi[2, 4] \bmod 2\pi. \end{aligned}$$

In the noiseless case, it is clear that $\psi[4] = \hat{\psi}^{(1)}[4] = \hat{\psi}^{(2)}[4]$. In a noisy environment, we can reduce the noise by averaging the two estimators, where averaging is done over the set of phases (namely, over the rotation group $\text{SO}(2)$) as explained in Appendix D. Specifically,

$$e^{i\hat{\psi}[4]} = \text{phase}\left(e^{i\hat{\psi}^{(1)}[4]} + e^{i\hat{\psi}^{(2)}[4]}\right).$$

We can iterate this procedure. To estimate phase q , we want to consider all entries of $\tilde{B}[k, \ell] = \tilde{y}[k]\tilde{y}[\ell]\tilde{y}[\ell - k]$ such that exactly one of the indices k, ℓ , or $\ell - k$ is equal to q and all other indices are in $1, \dots, q-1$, so that all other phases involved have already been estimated. A simple verification shows that only entries $\tilde{B}[p, q]$, $p = 1, \dots, q-1$, have that property. Furthermore, because of symmetry in the bispectrum (25), half of these entries are redundant so that only entries $\tilde{B}[p, q]$, $p = 1, \dots, \lfloor \frac{q}{2} \rfloor$ remain. As a result, estimation of the k th phase relies on averaging over $\lfloor \frac{k}{2} \rfloor$ equations, as summarized in Algorithm 4, with the following simple guarantee. The above construction yields the following proposition.

Proposition IV.1: Let $\tilde{B} = \tilde{B}_x$ be the normalized bispectrum as defined in (10) and assume that $\tilde{y}[1]$ is known. If $y[k] \neq 0$ for $k = 1, \dots, K$, then Algorithm 4 recovers the Fourier phases $\tilde{y}[k]$, $k = 1, \dots, K$ exactly.

We note in closing that, if the signal x is real, symmetries in the phases $\tilde{y}$ and $\tilde{B}$ can be exploited easily in FM.

B. Semidefinite Programming Relaxation

In this section we assume that the DFT y is non-vanishing so that the bispectrum relation can be manipulated as

$$\tilde{B} = \tilde{y}\tilde{y}^* \circ T(\tilde{y}) \equiv \tilde{B} \circ \overline{T(\tilde{y})} = \tilde{y}\tilde{y}^*,$$

where $\overline{T(\tilde{y})}$ is its entry-wise conjugate. The developments are easily adapted if the signal has zero mean. Similarly to the FM algorithm, we assume that $\tilde{y}[0]$ and $\tilde{y}[1]$ are available. We aim to estimate $\tilde{y}$ by a convex program. As a first step, we decouple the bispectrum equation and write the problem of estimating $\tilde{y}$

Algorithm 4: Frequency marching algorithm.

Input: Normalized bispectrum $\tilde{B}[k_1, k_2] = e^{i\Psi[k_1, k_2]}$, $\tilde{y}[0]$ and $\tilde{y}[1] \neq 0$

Output: $\hat{y}$: estimation of $\tilde{y}$

1) Set $\hat{y}[0] = \tilde{y}[0]$ and $e^{i\hat{\psi}[1]} = \tilde{y}[1]$

2) **For** $k = 2, \dots, N$ **do**:

a) Average the phase measurements:

$$u = \text{phase} \left(\sum_{\ell=1}^{\lfloor \frac{k}{2} \rfloor} e^{i(\hat{\psi}[\ell] + \hat{\psi}[k-\ell] - \Psi[\ell, k])} \right)$$

b) Estimate $\hat{\psi}[k]$ through:

$$e^{i\hat{\psi}[k]} = \begin{cases} u, & u \neq 0, \\ 1, & u = 0. \end{cases}$$

Return: $\hat{y} \leftarrow e^{i\hat{\psi}}$

as the following non-convex optimization problem:

$$\begin{aligned} \min_{Z \in \mathcal{H}^N, z \in \mathbb{C}^N} & \left\| W \circ (\tilde{B} \circ \overline{T(z)} - Z) \right\|_F^2 \\ \text{subject to } & Z = zz^*, \\ & \text{diag}(Z) = 1, \\ & z[0] = \tilde{y}[0], z[1] = \tilde{y}[1], \\ & (\text{if } x \text{ is real}) z[k] = \overline{z[-k]}, \forall k, \end{aligned} \quad (16)$$

where $\mathcal{H}^N$ is the set of Hermitian matrices of size N and $W \in \mathbb{R}^{N \times N}$ is a real weight matrix with positive entries. In particular, in the numerical experiments we set $W = |B|$.

In the absence of noise, the minimizers of (16) satisfy the bispectrum equation. However, in general these cannot be computed in polynomial time. In order to make the problem tractable, we relax the non-convex coupling constraint $Z = zz^*$ to the convex constraint $Z \succeq zz^*$ (that is, $Z - zz^*$ is positive semidefinite). The convex relaxation is then given by

$$\begin{aligned} \min_{Z \in \mathcal{H}^N, z \in \mathbb{C}^N} & \left\| W \circ (\tilde{B} \circ \overline{T(z)} - Z) \right\|_F^2 \\ \text{subject to } & Z \succeq zz^*, \\ & \text{diag}(Z) = 1, \\ & z[0] = \tilde{y}[0], z[1] = \tilde{y}[1], \\ & (\text{if } x \text{ is real}) z[k] = \overline{z[-k]}, \forall k. \end{aligned} \quad (17)$$

Upon solving (17), which can be done in polynomial time with interior point methods, the phases $\tilde{y}$ are estimated from $\text{phase}(z)$. In practice, we use CVX to solve this problem [65]. The algorithm is summarized in Algorithm 5. We note that problem (17) is not a standard SDP, in that its cost function is nonlinear.

In the noiseless case, the SDP relaxation (17) recovers the missing phases exactly. Interestingly, the proof is not so much based on optimality conditions as it is on an algebraic property

Algorithm 5: Semidefinite relaxation algorithm.

Input: The normalized bispectrum $\tilde{B}$, $\tilde{y}[0]$ and $\tilde{y}[1]$

Output: $\hat{y}$: estimation of $\tilde{y}$

Solve the SDP with nonlinear cost function (17), for example using CVX [65]

Return: $\hat{y} \leftarrow \text{phase}(z)$

of circulant matrices. The proof of the following property is given in Appendix E.

Lemma IV.2: Let $\hat{u}$ be the DFT of a vector $u \in \mathbb{C}^N$ obeying $u[k] = \overline{u[-k]}$, so that $\hat{u}$ is real. If $u[0] = u[1] = 1$ and $\hat{u}$ is non-negative, then $u[k] = 1$ for all k .

The following theorem is a direct corollary of Lemma IV.2. The main proof idea is as follows. Consider $u = \overline{\tilde{y}} \circ z$ where (Z, z) is optimal for the SDP; then, the constraints ensure $u[0] = u[1] = 1$. Furthermore, one can see via the Schur complement that the constraints force $T(u)$ to be positive semidefinite. Since the eigenvalues of $T(u)$ are the DFT of u , it follows that $\hat{u}$ is non-negative, so that the lemma above applies and $u \equiv 1$, or, equivalently, $z = \tilde{y}$. Details of the proof are in Appendix F.

Theorem IV.3: For a real signal with non-vanishing DFT y , if all weights in W are positive, $\tilde{y}[0]$ and $\tilde{y}[1]$ are known and the objective value of (17) attains 0 (which is the case in the absence of noise), then the SDP has a unique solution given by $z = \tilde{y}$ and $Z = zz^*$.

We close with an important remark about the symmetry breaking purpose of constraint $z[1] = \tilde{y}[1]$ in the SDP. Because the signal x can be recovered only up to integer time shifts, even in the noiseless case, without this constraint there are at least N distinct solutions (z, Z) to the SDP. Because SDP is a convex program, any point in the convex hull of these N points is also a solution. Thus, if the symmetry is not broken, the set of solutions contains many irrelevant points. Furthermore, interior point methods tend to converge to a center of the set of solutions, which in this case is never one of the desired solutions.

C. Phase Unwrapping by Integer Programming Algorithm

The next algorithm is based on solving an over-determined system of equations involving integers. Let us denote $\tilde{y}[k] = e^{i\psi[k]}$ and $\tilde{B}[k_1, k_2] = e^{i\Psi[k_1, k_2]}$ so the normalized bispectrum model is given by

$$e^{i\Psi[k_1, k_2]} = e^{i(\psi[k_1] - \psi[k_2] + \psi[k_2 - k_1])}.$$

By taking the logarithm, we get the algebraic relation

$$\Psi[k_1, k_2] + 2\pi\chi[k_1, k_2] = \psi[k_1] - \psi[k_2] + \psi[k_2 - k_1], \quad (18)$$

where, as a result of phase wrapping, χ takes on integer values. Let Ψ_{vec} and χ_{vec} be the column-stacked versions of Ψ and χ , respectively. Then, the model reads

$$\Psi_{\text{vec}} + 2\pi\chi_{\text{vec}} = A\psi, \quad (19)$$

where the sparse matrix $A \in \mathbb{R}^{N^2 \times N}$ encodes the right hand side of (18). It can be verified that A is of rank $N - 1$ (see for instance [66]), with null space corresponding to the time-shift-induced ambiguity on the phases (4). Note that both the

integer vector χ_{vec} and the phases ψ are unknown. Given χ_{vec} , the phases ψ can be obtained easily by solving

$$\min_{\psi \in \mathbb{R}^N} \|\Psi_{\text{vec}} + 2\pi\chi_{\text{vec}} - A\psi\|_p, \quad (20)$$

for some ℓ_p norm. Observe that any error in estimating χ may cause a big estimation error of ψ in (20). These errors can be thought of as outliers. Hence, we choose to use least unsquared deviations (LUD), $p = 1$, which is more robust to outliers. The more challenging task is to estimate the integer vector $\chi_{\text{vec}} \in \mathbb{Z}^{N^2}$. To this end, we first eliminate ψ from (19) as follows. Let $C \in \mathbb{R}^{(N^2 - (N-1)) \times N^2}$ be a full rank matrix such that $CA = 0$, that is, the columns of C^T are in the null space of A^T . Matrix C can be designed by at least two methods. One, suggested in [67], exploits the special structure of A to design a sparse matrix composed of integer values. Another, which we use here, is to take C to have orthonormal rows which form a basis of the kernel of A^T . Numerical experiments (not shown) indicate that the latter approach is more stable. Next, we multiply both sides of (19) from the left by C to get

$$C(\Psi_{\text{vec}} + 2\pi\chi_{\text{vec}}) = CA\psi = 0.$$

Therefore, the integer recovery problem can be formulated as

$$\min_{\chi_{\text{vec}} \in \mathbb{Z}^{N^2}} \left\| \frac{1}{2\pi} C\Psi_{\text{vec}} + C\chi_{\text{vec}} \right\|_2, \quad (21)$$

where we minimize over all integers. Note that $C\Psi_{\text{vec}}$ is a known vector. The problem is then equivalent to finding a lattice vector with the basis C which is as close as possible to the vector $-C\Psi_{\text{vec}}/(2\pi)$. While the problem is known to be NP-hard, we approximate the solution of (21) with the LLL (Lenstra–Lenstra–Lovasz) algorithm, which can be run in polynomial time [68]. The LLL algorithm computes a lattice basis, called a *reduced basis*, which is approximately orthogonal. It uses the Gram–Schmidt process to determine the quality of the basis. For more details, see [69, Ch. 17].

We note that (21) is under-determined as the matrix C is of rank $N^2 - \text{rank}(A) = N^2 - (N - 1)$. While the LLL algorithm works with under-determined systems, in our case we can solve it for a determined system since we can fix the first $N - 1$ entries of χ_{vec} to be zero.² Once we have estimated χ_{vec} , we solve (20) with $p = 1$. This approach is summarized in Algorithm 6.

V. ANALYSIS OF OPTIMIZATION OVER PHASES

In this section, we study the non-convex optimization problem (14) and give more implementation details to solve it, since numerical experiments identify this as the method of choice for MRA from invariant features among all methods compared. We start by considering the general case of a complex signal $x \in \mathbb{C}^N$ and consider the real case in Appendix B. Recall that we aim to maximize

$$f(z) = \langle z, M(z)z \rangle, \quad M(z) := W^{(2)} \circ \tilde{B} \circ \overline{T(z)},$$

²We omit the proof of this property here and only mention that it is based on the derivation in [67].

Algorithm 6: Phase unwrapping by integer programming.

Input: The normalized bispectrum $\tilde{B}[k_1, k_2] = e^{i\Psi[k_1, k_2]}$

Output: $\hat{y}$: estimation of $\tilde{y}$

1) (integer programming) Apply the LLL algorithm to estimate the integer vector χ_{vec} from

$$\min_{\chi_{\text{vec}} \in \mathbb{Z}^{N^2}} \|C\Psi_{\text{vec}}/(2\pi) + C\chi_{\text{vec}}\|_2,$$

where A is given in (19), $CA = 0$ and $\Psi_{\text{vec}} \in \mathbb{R}^{N^2}$ is a column-stacked version of Ψ , e.g., using code from [70].

2) (least-unsquared minimization) Let $\hat{\chi}_{\text{vec}}$ be the solution of stage 1. Then, solve

$$\hat{\psi} = \arg \min_{\psi \in \mathbb{R}^N} \|\Psi_{\text{vec}} + 2\pi\hat{\chi}_{\text{vec}} - A\psi\|_1.$$

Return: $\hat{y} \leftarrow e^{i\hat{\psi}}$

where the inner product is defined by (12), W is a real weighting matrix and $W^{(2)} := W \circ W$. The optimization problem lives on a *manifold*, that is, a smooth nonlinear space. Indeed, the smooth cost function $f(z)$ is to be maximized over the set

$$\mathcal{M} = \{z \in \mathbb{C}^N : |z[0]| = \dots = |z[N-1]| = 1\},$$

which is a Cartesian product of N unit circles in the complex plane (a torus). Theory and algorithms for optimization on manifolds can be found in the monograph [71]. We follow this formalism here. Details can also be found in [17], which deals with the similar problem of phase synchronization, using similar techniques. For the numerical experiments below, we use the toolbox Manopt which provides implementations of various optimization algorithms on manifolds [63].

Under mild conditions, the global optima of (14) correspond exactly to $\tilde{y}$ up to integer time shifts. This fact is proven in Appendix G.

Lemma V.1: For $N \geq 3$, let $x \in \mathbb{C}^N$ be a signal whose DFT y is nonzero for frequencies k in $\{1, \dots, K\}$, possibly also for $k = 0$, and zero otherwise. Up to integer time shifts, x is determined exactly by its bispectrum B provided $K \geq \frac{N+1}{2}$. Furthermore, the global optima of (14) correspond exactly to the relevant phases of y —up to the effects of integer time shifts—provided $W[k, \ell]$ is positive when $B[k, \ell] \neq 0$.

The problem at hand is

$$\max_{z \in \mathcal{M}} f(z). \quad (22)$$

This is smooth but non-convex, so that in general it is hard to compute the global optimum. We derive first- and second-order necessary optimality conditions. Points which satisfy these conditions are called *critical* and *second-order critical* points, respectively. Known algorithms converge to critical points (e.g., Riemannian gradient descent) and even to second-order critical points (e.g., Riemannian trust-regions) regardless of initialization [62], [64], [71]. Empirically, despite non-convexity, the global optimum appears to be computable reliably in favorable noise regimes.

As we proceed to consider optimization algorithms for (14), the gradient of f will come into play:

$$\nabla f(z) = M(z)z + M(z)^*z + M^{\text{adj}}(zz^*),$$

where $M^{\text{adj}} : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^N$ is the adjoint of M with respect to the inner product $\langle \cdot, \cdot \rangle$. Formally, the adjoint is defined such that, for any $z \in \mathbb{C}^N$, $X \in \mathbb{C}^{N \times N}$,

$$\langle z, M^{\text{adj}}(X) \rangle = \langle M(z), X \rangle.$$

Specifically, in Appendix H we show that

$$M^{\text{adj}}(X)[k] = \text{Tr} \left(T_k^\top \left(W^{(2)} \circ \tilde{B} \circ \overline{X} \right) \right), \quad (23)$$

where T_k is a circulant matrix with ones in its k th (circular) diagonal and zero otherwise, namely,

$$(T_k)[\ell', \ell] = \begin{cases} 1 & \text{if } \ell' = \ell - k, \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

As it turns out, under the symmetries of the problem at hand, there is no need to evaluate M^{adj} explicitly. Indeed, $\tilde{B}$ obeys

$$\tilde{B}[k_2 - k_1, k_2] = \tilde{y}[k_2 - k_1] \overline{\tilde{y}[k_2]} \tilde{y}[k_1] = \tilde{B}[k_1, k_2]. \quad (25)$$

This property is preserved when $\tilde{B}$ is obtained by averaging bispectra of multiple observations, as in (9). Assuming the same symmetry for the real weights W , we find below that $M^{\text{adj}}(zz^*) = M(z)z$. See Appendix I.

Lemma V.2: If $\tilde{B}[k_2 - k_1, k_2] = \tilde{B}[k_1, k_2]$ and $W[k_2 - k_1, k_2] = W[k_1, k_2]$ for all k_1, k_2 , then $M^{\text{adj}}(zz^*) = M(z)z$ for all $z \in \mathbb{C}^N$.

Thus, under the symmetries assumed in Lemma V.2, the gradient of f simplifies and we get a simple expression for the Hessian as well:

$$\nabla f(z) = 2M(z)z + M(z)^*z,$$

$$\nabla^2 f(z)[\dot{z}] = 2M(\dot{z})z + 2M(z)\dot{z} + M(\dot{z})^*z + M(z)^*\dot{z}. \quad (26)$$

For unconstrained optimization, the first-order necessary optimality conditions are $\nabla f(z) = 0$. In the presence of the constraint $z \in \mathcal{M}$, the conditions are different. Namely, following [71, eq. (3.37)], since $\mathcal{M}$ is a submanifold of $\mathbb{C}^N$, first-order necessary optimality conditions state that the orthogonal projection of the gradient $\nabla f(z)$ to the tangent space to $\mathcal{M}$ at z must vanish. The result of this projection is called the *Riemannian gradient*. Formally, the tangent space is obtained by linearizing (differentiating) the constraints $|z[k]|^2 = \langle z[k], z[k] \rangle = 1$ for all k , yielding

$$\text{T}_z \mathcal{M} = \{ \dot{z} \in \mathbb{C}^N : \langle z[k], \dot{z}[k] \rangle = 0, \forall k \}.$$

Orthogonal projection of $u \in \mathbb{C}^N$ to the tangent space $\text{T}_z \mathcal{M}$ can be computed entry-wise by subtracting from each $u[k]$ its component aligned with $z[k]$. Let $\text{Proj}_z : \mathbb{C}^N \rightarrow \text{T}_z \mathcal{M}$ denote this projection. This operation admits a compact matrix notation as

$$\begin{aligned} u &\mapsto \text{Proj}_z(u) = u - \Re\{u \circ \bar{z}\} \circ z \\ &= u - \Re\{\text{ddiag}(uz^*)\}z, \end{aligned}$$

where $\text{ddiag} : \mathbb{C}^{N \times N} \rightarrow \mathbb{C}^{N \times N}$ sets all non-diagonal entries of a matrix to zero. Equipped with this notion and the expression

for $\nabla f(z)$ (26), it follows that the Riemannian gradient of f at z on $\mathcal{M}$ is

$$\text{grad}f(z) := \text{Proj}_z(\nabla f(z)) = \nabla f(z) - D(z)z,$$

with

$$D(z) := \Re\{\text{ddiag}(\nabla f(z)z^*)\} = \Re\{\text{diag}(\nabla f(z) \circ \bar{z})\}.$$

Lemma V.3: If $z \in \mathcal{M}$ is optimal for (22), then $\text{grad}f(z) = 0$; equivalently, $\text{diag}(\nabla f(z)z^*) = \nabla f(z) \circ \bar{z}$ is real.

Proof: See [72, Rem. 4.2 and Cor. 4.2]. For the equivalence, notice that $\text{Proj}_z(u) = 0$ if and only if $u[k] = \Re\{u[k]z[k]\}z[k]$ for all k , and multiply by $\bar{z}[k]$ on both sides using $|z[k]| = 1$. ■

A point z which satisfies these conditions is called a critical point. Likewise, we can define a notion of *Riemannian Hessian* as the linear, self-adjoint operator on $\text{T}_z \mathcal{M}$ which captures infinitesimal changes in the Riemannian gradient around z . Without getting into technical details, we follow [71, eq. (5.15)] and define (with D the directional derivative operator):

$$\begin{aligned} \text{Hess}f(z)[\dot{z}] &:= \text{Proj}_z(D(z \mapsto \text{grad}f(z))(z)[\dot{z}]) \\ &= \text{Proj}_z(\nabla^2 f(z)[\dot{z}] - D(z)\dot{z} - (DD(z)[\dot{z}]z)), \end{aligned}$$

where $DD(z)[\dot{z}]$ is a real, diagonal matrix. Its contribution to the Hessian is zero, since $(DD(z)[\dot{z}]z)$ vanishes under the projection Proj_z . Hence,

$$\text{Hess}f(z)[\dot{z}] = \text{Proj}_z(\nabla^2 f(z)[\dot{z}] - D(z)\dot{z}).$$

The Riemannian Hessian intervenes in the second-order necessary optimality conditions as follows.

Lemma V.4: If $z \in \mathcal{M}$ is optimal for (22), then $\text{grad}f(z) = 0$ and $\text{Hess}f(z) \preceq 0$, that is, for all $\dot{z} \in \text{T}_z \mathcal{M}$ we have

$$\langle \dot{z}, \text{Hess}f(z)[\dot{z}] \rangle = \langle \dot{z}, \nabla^2 f(z)[\dot{z}] \rangle - \langle \dot{z}, D(z)\dot{z} \rangle \leq 0.$$

Proof: See [72, Rem. 4.2 and Cor. 4.2]. In the equality, we used the fact that Proj_z is self-adjoint and $\dot{z} \in \text{T}_z \mathcal{M}$. ■

A point z which satisfies these conditions is called a *second-order critical point*. With unit weights, the following lemma shows that second-order critical points z , in the noiseless case, cannot have an arbitrarily bad objective value $f(z)$. This result is weak, however, since empirically it is observed that in the noiseless case local optimization methods consistently converge to global optima whose value are N^2 , suggesting that all second-order critical points are global optima in this simplified scenario. While we do not have a proof for this stronger conjecture, we provide the lemma below because it is analogous to [17, Lemma 14] which, in that reference, is a key step toward proving global optimality of second-order critical points.

Lemma V.5: In the absence of noise and with unit weights, a second-order critical point z of (22) satisfies

$$\nabla f(z) \circ \bar{z} \geq 2(\sqrt{3} - 1) > 0.$$

In particular, this implies

$$f(z) = \frac{1}{3} \langle z, \nabla f(z) \rangle \geq \frac{2(\sqrt{3} - 1)}{3} N.$$

Proof: See Appendix J for the proof of the inequality. It follows from two key considerations. First, because z is a critical

point, Lemma V.3 indicates that $\nabla f(z) \circ \bar{z}$ is real. Second, because z is second-order critical, the Riemannian Hessian at z must be negative semidefinite by Lemma V.4. Applied to all tangent directions at z which perturb only one phase at a time implies the desired inequality. The fact that $3f(z) = \langle z, \nabla f(z) \rangle$ follows from (26). ■

One final ingredient that is necessary to optimize f over $\mathcal{M}$ is a means of moving away from a current iterate $z \in \mathcal{M}$ to the next by following a tangent vector $\dot{z}$. A simple means of achieving this is through a *retraction* [71, Def. 4.1.1]. For $\mathcal{M}$, an obvious retraction is the following:

$$\text{Retr}_z(\dot{z}) = \text{phase}(z + \dot{z}) \in \mathcal{M}. \quad (27)$$

With the formalism of (22) and the above derivations, we can now run a local Riemannian optimization algorithm. As an example, the Riemannian gradient ascent algorithm would iterate the following:

$$\begin{aligned} z^{(t+1)} &= \text{Retr}_{z^{(t)}} \left(\eta^{(t)} \text{grad} f(z^{(t)}) \right) \\ &= \text{phase} \left(z^{(t)} + \eta^{(t)} \text{grad} f(z^{(t)}) \right), \end{aligned}$$

where $\eta^{(t)} > 0$ is an appropriately chosen step size and $z^{(0)} \in \mathcal{M}$ is an initial guess. It is relatively easy to choose the step sizes such that the sequence $z^{(t)}$ converges to critical points regardless of $z^{(0)}$, with a linear local convergence rate [71, §4]. In practice, we prefer to use the Riemannian trust-region method (RTR) [62], whose usage is simplified by the toolbox Manopt [63]. RTR enjoys global convergence to second-order critical points [64] and a quadratic local convergence rate.

In this section, the analysis focused on complex signals. For real signals, we can follow the same methodology while taking the symmetry in the Fourier domain into account. This analysis is given in Appendix B.

VI. EXPECTATION MAXIMIZATION

In this section, we detail the expectation maximization algorithm (EM) [73] applied to MRA. As the numerical experiments in Section VII demonstrate, EM achieves excellent accuracy in estimating the signal. However, compared to the invariant features approach proposed in this paper, it is significantly slower and requires many passes over the data (thus excluding online processing).

Let $X = [\xi_1, \dots, \xi_M]$ be the data matrix of size $N \times M$, following the MRA model (1). The maximum marginalized likelihood estimator (MMLE) for the signal x given X is the maximizer of the likelihood function $L(x; X) = p(X|x)$ (the probability density of X given x). This density could in principle be evaluated by marginalizing the joint distribution $p(X, r|x)$ over the unknown shifts $r \in \{0, \dots, N-1\}^M$. This, however, is intractable as it involves summing over N^M terms.

Alternatively, EM tries to estimate the MMLE as follows. Given a current estimate for the signal x_k , consider the expected value of the log-likelihood function, with respect to the conditional distribution of r given X and x_k :

$$Q(x|x_k) = \mathbb{E}_{r|X, x_k} \{ \log p(X, r|x_k) \}. \quad (28)$$

This step is called the E-step. Then, iterate by computing the M-step:

$$x_{k+1} = \arg \max_x Q(x|x_k). \quad (29)$$

For the MRA model, this can be done in closed form. Indeed, the log-likelihood function follows from the i.i.d. Gaussian noise model:

$$\log p(X, r|x) = -\frac{1}{2\sigma^2} \sum_{j=1}^M \|R_{r_j} x - \xi_j\|_2^2 + \text{constant}. \quad (30)$$

To take the expectation with respect to r , we need to compute $w_k^{\ell, j}$: for each observation j , this is the probability that the shift r_j is equal to ℓ , given X and assuming $x = x_k$. This also follows easily from the i.i.d. Gaussian noise model:

$$w_k^{\ell, j} \propto \exp \left(-\frac{1}{2\sigma^2} \|R_\ell x_k - \xi_j\|_2^2 \right), \quad (31)$$

(with appropriate scale so that $\sum_{\ell=0}^{N-1} w_k^{\ell, j} = 1$). This allows to write Q down explicitly:

$$Q(x|x_k) = -\frac{1}{2\sigma^2} \sum_{j=1}^M \sum_{\ell=0}^{N-1} w_k^{\ell, j} \|R_\ell x - \xi_j\|_2^2 + \text{constant}.$$

This is a convex quadratic expression in x with maximizer

$$x_{k+1} = \frac{1}{M} \sum_{j=1}^M \sum_{\ell=0}^{N-1} w_k^{\ell, j} R_\ell^{-1} \xi_j. \quad (32)$$

In words: given an estimator x_k , the next estimator is obtained by averaging all shifted versions of all observations, weighted by the empirical probabilities of the shifts. Considering all shifts of all observations would, in principle, induce an iteration complexity of $O(MN^2)$, but fortunately, for each observation, the matrix of its shifted versions is circulant, which makes it possible to use FFT to reduce the overall computational cost to $O(MN \log N)$. See the available code for details. We note that Matlab naturally parallelizes the computations over M .

In practice, we set $x_0 \sim \mathcal{N}(0, I_N)$ to be a random guess. Furthermore, for $M \geq 3000$, we first execute 3000 batch iterations, where the EM update is computed based on a random sample of 1000 observations (fresh sample at each iteration). This inexpensively transforms the random initialization into a ballpark estimate of the signal. The algorithm then proceeds with full-data iterations until the relative change between two consecutive estimates drops below 10^{-5} (in ℓ_2 -norm, up to shifts).

VII. NUMERICAL EXPERIMENTS

This section is devoted to numerical experiments, examining all proposed algorithms. Code for all algorithms and to reproduce the experiments is available online.³ The experiments were conducted as follows. The true signal x of length $N = 41$ is a fixed window of height 1 and width 21. With this signal, the signal-to-noise ratio is $\frac{\|x\|_2^2}{\|\varepsilon\|_2^2} \approx \frac{1}{2\sigma^2}$. We generated a set of M shifted noisy versions of x as

$$\xi_j = R_{r_j} x + \varepsilon_j,$$

³<https://github.com/NicolasBoumal/MRA>




Fig. 3. Relative recovery error for the signal x as a function of the number of observations M for fixed noise level $\sigma = 1$. The curves corresponding to the optim. phase manifold (Algorithm 2) and the iter. phase synch. (Algorithm 3) overlap.




Fig. 4. Average computation times corresponding to Fig. 3.

where each shift was randomly drawn from a uniform distribution over $\{0, \dots, N-1\}$ and $\varepsilon_j \sim \mathcal{N}(0, \sigma^2 I)$ for all j . The relative recovery error for a single experiment is defined as

$$\text{relative error}(x, \hat{x}) = \min_{s \in \{0, \dots, N-1\}} \frac{\|R_s \hat{x} - x\|_2}{\|x\|_2},$$

where $\hat{x}$ is the estimation of the signal. All results are averaged over 20 repetitions. While we present here results for a specific signal, alternative signal models (e.g., random signals) showed similar numerical behavior.

The following figures compare the recovery errors for all proposed algorithms, with random initialization for those that need initialization. The non-convex algorithm on the manifold of phases, Algorithm 2, runs the Riemannian trust-region method (RTR) [62] using the toolbox Manopt [63]. Algorithm 3 runs 15 iterations with warm-start using the same toolbox. For the phase unwrapping algorithm, Algorithm 6, we use an implementation of LLL available in the MILES package [70]. The SDP is solved with CVX [65]. The EM algorithm is implemented as explained in Section VI. We compared the algorithms with an oracle who knows the random shifts r_j and therefore simply averages out the Gaussian noise. Experiments are run on a computer with 30 CPUs available. These CPUs are used to compute the invariants in parallel (with Matlab's `parfor`), while the EM algorithm benefits from parallelism to run the many thousands of FFTs it




Fig. 5. Relative recovery error for the signal x as a function of the noise level σ with $M = 10,000$ observations. The curves corresponding to the optim. phase manifold (Algorithm 2) and the iter. phase synch. (Algorithm 3) overlap.




Fig. 6. Average computation times corresponding to Fig. 5.

requires efficiently (built-in Matlab). The algorithms that need $\tilde{y}[0]$ and $\tilde{y}[1]$ are given the correct values.

Figs. 3 and 4 present the recovery error and computation time of all algorithms as a function of the number of observations M for fixed noise level $\sigma = 1$. Of course, the oracle who knows the shifts of the observations is unbeatable. Algorithms 2 and 3 outperform all invariant approach methods. The inferior performance of the SDP might be explained by the fact that we are minimizing a smooth non-linear objective. This is in contrast to SDPs with linear or piecewise linear objectives which tend to promote “simple” (i.e., low rank) solutions [74], [75, Remark 6.2]. Additionally, while this is not depicted on the figure, we note that for $\sigma = 0$ all methods get exact recovery up to machine precision. EM outperforms the best invariant features approaches by a factor of 3, at the cost of being significantly slower. For large M , the best invariant features approaches are faster than EM by a factor of 25. Note, however, that for M up to about 300, EM is faster than the other algorithms. For invariant features approaches (aside from the SDP), almost all of the time is spent computing the bispectrum estimator, while inverting the bispectrum is relatively cheap.

Figs. 5 and 6 show the recovery error and computation time as a function of the noise level σ with $M = 10,000$ observations. Surprisingly, for high noise level $\sigma \gtrsim 3$, the invariant features algorithms outperform EM.

VIII. CONCLUSIONS AND PERSPECTIVE

The goal of this paper is twofold. First, we have suggested a new approach for the MRA problem based on features that are invariant under translations. This technique enables us to deal with any noise level as long as we have access to enough measurements and particularly it achieves the sample complexity of MRA. Extensions have been proposed shortly after a first appearance of this manuscript, focusing on non-uniform distribution of translations and the heterogeneous MRA model [76], [77]. The invariant features approach has low computational complexity and it requires less memory with respect to alternative methods, such as EM.

A main ingredient of the invariant features approach is estimating the signal's Fourier phases by inverting the bispectrum. Hence, the second goal of this paper was to study algorithms for bispectrum inversion. We have proposed a few algorithms for this task. In the presence of noise, the non-convex algorithms on the manifold of phases, namely, Algorithms 2 and 3, perform the best. Empirically, these algorithms have a remarkable property: despite their non-convex landscape, they appear to converge to the target signal from random initialization. We provide some analysis for Algorithm 2 but this phenomenon is not well understood.

Our chief motivation for this work comes from the more involved problem of cryo-EM. In cryo-EM, a 3D object is estimated from its 2D projections at unknown rotations in a low SNR environment. One line of research for the object recovery is based on first estimating the unknown rotations [78]–[80]. However, the rotation estimation is performed in a very noisy environment and therefore might be inaccurate. An interesting question is to examine whether the 3D object can be estimated directly from the acquired data using features that are invariant under the unknown viewing directions [81], [82].

ACKNOWLEDGMENT

The authors are grateful to Afonso Bandeira, Roy Lederman, William Leeb, Nir Sharon and Susannah Shoemaker for many insightful discussions. The authors also thank the reviewers for their useful comments, and particularly the anonymous reviewer who proposed a simpler proof for Lemma IV.2.

REFERENCES

- [1] R. Diamond, "On the multiple simultaneous superposition of molecular structures by rigid body transformations," *Protein Sci.*, vol. 1, no. 10, pp. 1279–1287, 1992.
- [2] D. L. Theobald and P. A. Steindel, "Optimal simultaneous superpositioning of multiple structures with missing data," *Bioinformatics*, vol. 28, no. 15, pp. 1972–1979, 2012.
- [3] 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. Robot. Res.*, vol. 30, no. 6, pp. 730–754, 2011.
- [4] W. Park and G. S. Chirikjian, "An assembly automation approach to alignment of noncircular projections in electron microscopy," *IEEE Trans. Autom. Sci. Eng.*, vol. 11, no. 3, pp. 668–679, Jul. 2014.
- [5] 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," *IEE Proc., Radar, Sonar Navig.*, vol. 150, no. 6, pp. 411–418, 2003.
- [6] R. Gil-Pita, M. Rosa-Zurera, P. Jarabo-Amores, and F. López-Ferreras, "Using multilayer perceptrons to align high range resolution radar signals," in *Proc. Int. Conf. Artif. Neural Netw.*, 2005, pp. 911–916.
- [7] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*, vol. 4. Chichester, U.K.: Wiley, 1998.
- [8] H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE Trans. Image Process.*, vol. 11, no. 3, pp. 188–200, Mar. 2002.
- [9] D. Robinson, S. Farsiu, and P. Milanfar, "Optimal registration of aliased images using variable projection with applications to super-resolution," *Comput. J.*, vol. 52, no. 1, pp. 31–42, 2009.
- [10] E. Abbe, J. Pereira, and A. Singer, "Sample complexity of the boolean multireference alignment problem," *Proc. IEEE Int. Symp. Inf. Theory*, 2017, pp. 1316–1320.
- [11] A. Bandeira, P. Rigollet, and J. Weed, "Optimal rates of estimation for multi-reference alignment," arXiv preprint arXiv:1702.08546, 2017.
- [12] A. Bartesaghi *et al.*, "2.2 Å resolution cryo-EM structure of β -galactosidase in complex with a cell-permeant inhibitor," *Science*, vol. 348, no. 6239, pp. 1147–1151, 2015.
- [13] D. Sirohi *et al.*, "The 3.8 Å resolution cryo-EM structure of Zika virus," *Science*, vol. 352, no. 6284, pp. 467–470, 2016.
- [14] J. Frank, *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. London, U.K.: Oxford Univ. Press, 2006.
- [15] M. van Heel *et al.*, "Single-particle electron cryo-microscopy: Towards atomic resolution," *Quart. Rev. Biophys.*, vol. 33, no. 04, pp. 307–369, 2000.
- [16] A. Singer, "Angular synchronization by eigenvectors and semidefinite programming," *Appl. Comput. Harmon. Anal.*, vol. 30, no. 1, pp. 20–36, 2011.
- [17] N. Boumal, "Nonconvex phase synchronization," *SIAM J. Optim.*, vol. 26, no. 4, pp. 2355–2377, 2016.
- [18] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra, "Message-passing algorithms for synchronization problems over compact groups," arXiv preprint arXiv:1610.04583, 2016.
- [19] Y. Chen and E. Candes, "The projected power method: An efficient algorithm for joint alignment from pairwise differences," arXiv preprint arXiv:1609.05820, 2016.
- [20] A. S. Bandeira, N. Boumal, and A. Singer, "Tightness of the maximum likelihood semidefinite relaxation for angular synchronization," *Math. Program.*, vol. 163, no. 1, pp. 145–167, 2017.
- [21] Y. Zhong and N. Boumal, "Near-optimal bounds for phase synchronization," *SIAM J. Control Optim.*, to be published.
- [22] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu, "Multireference alignment using semidefinite programming," in *Proc. 5th Conf. Innov. Theor. Comput. Sci.*, 2014, pp. 459–470.
- [23] Y. Chen, L. Guibas, and Q. Huang, "Near-optimal joint object matching via convex relaxation," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 100–108.
- [24] P. Kosir, R. DeWall, and R. A. Mitchell, "A multiple measurement approach for feature alignment," in *Proc. IEEE Nat. Aerosp. Electron. Conf.*, vol. 1, 1995, pp. 94–101.
- [25] C. Aguerrebere, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5707–5722, Nov. 2016.
- [26] D. Robinson and P. Milanfar, "Fundamental performance limits in image registration," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1185–1199, Sep. 2004.
- [27] A. Weiss and E. Weinstein, "Fundamental limitations in passive time delay estimation—Part I: Narrow-band systems," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-31, no. 2, pp. 472–486, Apr. 1983.
- [28] E. Weinstein and A. Weiss, "Fundamental limitations in passive time-delay estimation—Part II: Wide-band systems," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-32, no. 5, pp. 1064–1078, Oct. 1984.
- [29] A. Perry, J. Weed, A. Bandeira, P. Rigollet, and A. Singer, "The sample complexity of multi-reference alignment," arXiv preprint arXiv:1707.00943, 2017.
- [30] P. L. Brockett, M. J. Hinich, and D. Patterson, "Bispectral-based tests for the detection of Gaussianity and linearity in time series," *J. Amer. Statist. Assoc.*, vol. 83, no. 403, pp. 657–664, 1988.
- [31] N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto, "Non-Gaussianity from inflation: Theory and observations," *Phys. Rep.*, vol. 402, no. 3, pp. 103–266, 2004.
- [32] X. Luo, "The angular bispectrum of the cosmic microwave background," *Astrophys. J.*, vol. 427, no. 2, pp. L71–L71, 1994.
- [33] L. Wang and M. Kamionkowski, "Cosmic microwave background bispectrum and inflation," *Phys. Rev. D*, vol. 61, no. 6, 2000, Art. no. 063504.
- [34] T. Matsuoka and T. J. Ulrych, "Phase estimation using the bispectrum," *Proc. IEEE*, vol. 72, no. 10, pp. 1403–1411, Oct. 1984.

- [35] M. M. Chang, A. M. Tekalp, and A. T. Erdem, "Blur identification using the bispectrum," *IEEE Trans. Signal Process.*, vol. 39, no. 10, pp. 2323–2325, Oct. 1991.
 - [36] T.-W. Chen, W.-D. Jin, and J. Li, "Feature extraction using surrounding-line integral bispectrum for radar emitter signal," in *Proc. 2008 IEEE Int. Joint Conf. Neural Netw. (IEEE World Congr. Comput. Intell.)*, 2008, pp. 294–298.
 - [37] T. Ning and J. D. Bronzino, "Bispectral analysis of the rat EEG during various vigilance states," *IEEE Trans. Biomed. Eng.*, vol. 36, no. 4, pp. 497–499, Apr. 1989.
 - [38] B. Chen and A. P. Petropulu, "Frequency domain blind MIMO system identification based on second-and higher order statistics," *IEEE Trans. Signal Process.*, vol. 49, no. 8, pp. 1677–1688, Aug. 2001.
 - [39] Z. Zhao and A. Singer, "Rotationally invariant image representation for viewing direction classification in cryo-EM," *J. Struct. Biol.*, vol. 186, no. 1, pp. 153–166, 2014.
 - [40] J. M. Mendel, "Tutorial on higher-order statistics (spectra) in signal processing and system theory: Theoretical results and some applications," *Proc. IEEE*, vol. 79, no. 3, pp. 278–305, Mar. 1991.
 - [41] C. L. Nikias and M. R. Raghuveer, "Bispectrum estimation: A digital signal processing framework," *Proc. IEEE*, vol. 75, no. 7, pp. 869–891, Jul. 1987.
 - [42] R. Marabini and J. M. Carazo, "Practical issues on invariant image averaging using the bispectrum," *Signal Process.*, vol. 40, nos. 2–3, pp. 119–128, 1994.
 - [43] R. Marabini and J. M. Carazo, "On a new computationally fast image invariant based on bispectral projections," *Pattern Recog. Lett.*, vol. 17, no. 9, pp. 959–967, 1996.
 - [44] A. P. Petropulu and H. Pozidis, "Phase reconstruction from bispectrum slices," *IEEE Trans. Signal Process.*, vol. 46, no. 2, pp. 527–530, Feb. 1998.
 - [45] G. B. Giannakis, "Signal reconstruction from multiple correlations: Frequency-and time-domain approaches," *J. Opt. Soc. Amer. A*, vol. 6, no. 5, pp. 682–697, 1989.
 - [46] B. M. Sadler and G. B. Giannakis, "Shift-and rotation-invariant object reconstruction using the bispectrum," *J. Opt. Soc. Amer. A*, vol. 9, no. 1, pp. 57–69, 1992.
 - [47] A. Swami, G. Giannakis, and J. Mendel, "Linear modeling of multidimensional non-gaussian processes using cumulants," *Multidimens. Syst. Signal Process.*, vol. 1, no. 1, pp. 11–37, 1990.
 - [48] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.
 - [49] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: A contemporary overview," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, May 2015.
 - [50] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Convex algorithms and limitations," in *Proc. 2013 IEEE Int. Symp. Inf. Theory*, 2013, pp. 1022–1026.
 - [51] T. Bendory, Y. Eldar, and N. Boumal, "Non-convex phase retrieval from STFT measurements," *IEEE Trans. Inf. Theory*, vol. 63, no. 12, pp. 1–18, Dec. 2017.
 - [52] R. Beinert and G. Plonka, "Ambiguities in one-dimensional discrete phase retrieval from fourier magnitudes," *J. Fourier Anal. Appl.*, vol. 21, no. 6, pp. 1169–1198, 2015.
 - [53] T. Bendory, R. Beinert, and Y. C. Eldar, "Fourier phase retrieval: Uniqueness and algorithms," in *Proc. Compressed Sensing and Its Applications: MATHEON Workshop 2015 (Applied and Numerical Harmonic Analysis)*, 2017.
 - [54] T. Bendory, D. Edidin, and Y. C. Eldar, "On signal reconstruction from FROG measurements," arXiv preprint arXiv:1706.08494, 2017.
 - [55] J. W. Tukey, "The spectral representation and transformation properties of the higher moments of stationary time series," in *The Collected Works of John W. Tukey*, vol. 1, D. R. Brillinger, Ed. Belmont, CA, USA: Wadsworth, 1984, ch. 4, pp. 165–184.
 - [56] J. I. Yellott and G. J. Iverson, "Uniqueness properties of higher-order autocorrelation functions," *J. Opt. Soc. Amer. A*, vol. 9, no. 3, pp. 388–404, 1992.
 - [57] R. Kondor, "A novel set of rotationally and translationally invariant features for images based on the non-commutative bispectrum," arXiv preprint arXiv:cs/0701127, 2007.
 - [58] R. Kakarala, "Completeness of bispectrum on compact groups," arXiv preprint arXiv:0902.0196, 2009.
 - [59] R. Kakarala, "The bispectrum as a source of phase-sensitive invariants for fourier descriptors: A group-theoretic approach," *J. Math. Imag. Vision*, vol. 44, no. 3, pp. 341–353, 2012.
 - [60] R. Kakarala, "Bispectrum on finite groups," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2009, pp. 3293–3296.
 - [61] L. Devroye et al., "Sub-gaussian mean estimators," *Ann. Statist.*, vol. 44, no. 6, pp. 2695–2725, 2016.
 - [62] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on Riemannian manifolds," *Found. Comput. Math.*, vol. 7, no. 3, pp. 303–330, 2007.
 - [63] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a Matlab toolbox for optimization on manifolds," *J. Mach. Learn. Res.*, vol. 15, pp. 1455–1459, 2014. [Online]. Available: <http://www.manopt.org>.
 - [64] N. Boumal, P.-A. Absil, and C. Cartis, "Global rates of convergence for nonconvex optimization on manifolds," to appear in *IMA Journal of Numerical Analysis*.
 - [65] M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab software for disciplined convex programming," 2008.
 - [66] T. Bendory, P. Sidorenko, and Y. C. Eldar, "On the uniqueness of frog methods," *IEEE Signal Process. Lett.*, vol. 24, no. 5, pp. 722–726, May 2017.
 - [67] J. Marron, P. Sanchez, and R. Sullivan, "Unwrapping algorithm for least-squares phase recovery from the modulo 2π bispectrum phase," *J. Opt. Soc. Amer. A*, vol. 7, no. 1, pp. 14–20, 1990.
 - [68] A. K. Lenstra, H. W. Lenstra, and L. Lovász, "Factoring polynomials with rational coefficients," *Mathematische Annalen*, vol. 261, no. 4, pp. 515–534, 1982.
 - [69] S. D. Galbraith, *Mathematics of Public Key Cryptography*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
 - [70] X.-W. Chang and T. Zhou, "Miles: Matlab package for solving mixed integer least squares problems," *GPS Solution*, vol. 11, no. 4, pp. 289–294, 2007.
 - [71] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ, USA: Princeton Univ. Press, 2009.
 - [72] W. H. Yang, L.-H. Zhang, and R. Song, "Optimality conditions for the non-linear programming problems on Riemannian manifolds," *Pac. J. Optim.*, vol. 10, no. 2, pp. 415–434, 2014.
 - [73] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. Ser. B, Methodol.*, vol. 39, no. 1, pp. 1–38, 1977.
 - [74] G. Pataki, "On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues," *Math. Oper. Res.*, vol. 23, no. 2, pp. 339–358, 1998.
 - [75] N. Boumal, "A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints," arXiv preprint arXiv:1506.00575, 2015.
 - [76] E. Abbe, T. Bendory, W. Lee, J. Pereira, N. Sharon, and A. Singer, "Multireference alignment is easier with an aperiodic translation distribution," arXiv preprint arXiv:1710.02793, 2017.
 - [77] N. Boumal, T. Bendory, R. R. Lederman, and A. Singer, "Heterogeneous multireference alignment: A single pass approach," arXiv preprint arXiv:1710.02590, 2017.
 - [78] A. Singer and Y. Shkolnisky, "Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming," *SIAM J. Imag. Sci.*, vol. 4, no. 2, pp. 543–572, 2011.
 - [79] Y. Shkolnisky and A. Singer, "Viewing direction estimation in cryo-EM using synchronization," *SIAM J. Imag. Sci.*, vol. 5, no. 3, pp. 1088–1110, 2012.
 - [80] L. Wang, A. Singer, and Z. Wen, "Orientation determination of cryo-EM images using least unsquared deviations," *SIAM J. Imag. Sci.*, vol. 6, no. 4, pp. 2450–2483, 2013.
 - [81] Z. Kam, "The reconstruction of structure from electron micrographs of randomly oriented particles," *J. Theor. Biol.*, vol. 82, no. 1, pp. 15–39, 1980.
 - [82] E. Levin, T. Bendory, N. Boumal, J. Kileel, and A. Singer, "3D ab initio modeling in cryo-EM by autocorrelation analysis," arXiv preprint arXiv:1710.08076, 2017.
- Tamir Bendory**, photograph and biography not available at the time of publication.
- Nicolas Boumal**, photograph and biography not available at the time of publication.
- Chao Ma**, photograph and biography not available at the time of publication.
- Zhizhen Zhao**, photograph and biography not available at the time of publication.
- Amit Singer**, photograph and biography not available at the time of publication.