

Bispectrum Inversion with Application to Multireference Alignment

Tamir Bendory[†], Nicolas Boumal[†], Chao Ma, Zhizhen Zhao 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 non-convex optimization over the smooth manifold of phases. Empirically, in the absence of noise, these non-convex 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. 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 optimal 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], [2], [3], [4], radar [5], [6], image registration and super-resolution [7], [8], [9], and has been the subject of recent theoretical analysis [10], [11]. The MRA model reads

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

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 shift.

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 (highly) 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 shifts 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

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

where \Re extracts the real part of a complex number. This approach requires only one pass over the data: for each observation, the best shift can be computed in $O(N \log N)$, and the aligned observations can be averaged online. This results in a total computational cost of $O(MN \log N)$, see Table I. While this approach is simple and efficient, it necessarily fails below a critical SNR—see Figure I.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], [17], [18], [19], [20], [21]. The computational complexity of aligning all pairs

[†] These two authors contributed equally to the work.

TB, NB, CM and AS are with Princeton University in PACM and the Mathematics Department. ZZ is with the University of Illinois at Urbana-Champaign in the Department of Electrical and Computer Engineering. The authors were partially supported 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.

individually is $O(M^2N \log N)$, 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], [26], [27], [28]. For example, for a continuous, 2D version of the MRA model, it was shown that the Cramér–Rao lower bound (CRLB) on the variance for translation estimation is proportional to the noise variance σ^2 [25]; crucially, it does not improve with M .

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.

In this work, we propose to estimate the sought signal x directly by 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. This approach allows us to deal with *any noise level* as long as we can acquire sufficiently many measurements. In particular, for fixed signal length N , the required number of measurements M to achieve accurate estimation scales as σ^6 . In this sense, our approach achieves the information-theoretic estimation rate for MRA—see Theorem 1 in [11] and accompanying discussion. Additionally, its computational complexity and memory requirements are relatively low. Indeed, for the proposed approach, the only operations whose computational costs grow with M are computations of averages over the data. These can be performed on-the-fly and are easily parallelizable.

In order to recover the Fourier phases from the bispectrum (i.e., bispectrum inversion), we propose two non-convex optimization algorithms on the manifold of phases. We also suggest three additional techniques that can be used to initialize these algorithms, based on frequency marching, SDP relaxation and integer programming. The first two methods provably recover the phases exactly in the absence of noise.

In addition to 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 [29], [30]. It is also used to investigate the cosmic background radiation [31], [32], seismic signal processing [33], image deblurring [34], feature extraction for radar [35], analysis of EEG signals [36] and classification [37] (see also [38], [39] and references therein). 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 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 and y denote the underlying signal and its discrete Fourier transform (DFT), respectively. The sign of a complex scalar a is defined as $\text{phase}(a) = a/|a|$ if $a \neq 0$ and zero otherwise. We further use \bar{a} and $\Re\{a\}$ for its conjugate and real part, respectively. 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 $N \times N$. All indexing is modulo 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

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

It is clear that c_1, c_2, c_3 are invariants under circular shifts of x . For higher-order invariants based on auto-correlations, see for instance [40].

The first feature is the mean of the signal which is the auto-correlation function of order one (i.e., c_1 in (II.1)). The distribution of the mean of ξ_j is then given by $\mu_{\xi_j} \sim \mathcal{N}\left(\mu_x, \frac{\sigma^2}{N}\right)$ and we can estimate μ_x as

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

Estimating the signal’s mean supplies only limited information about the signal itself. Thus, we consider also the auto-

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

Table I: Comparison of main MRA approaches. $F(N)$ 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.

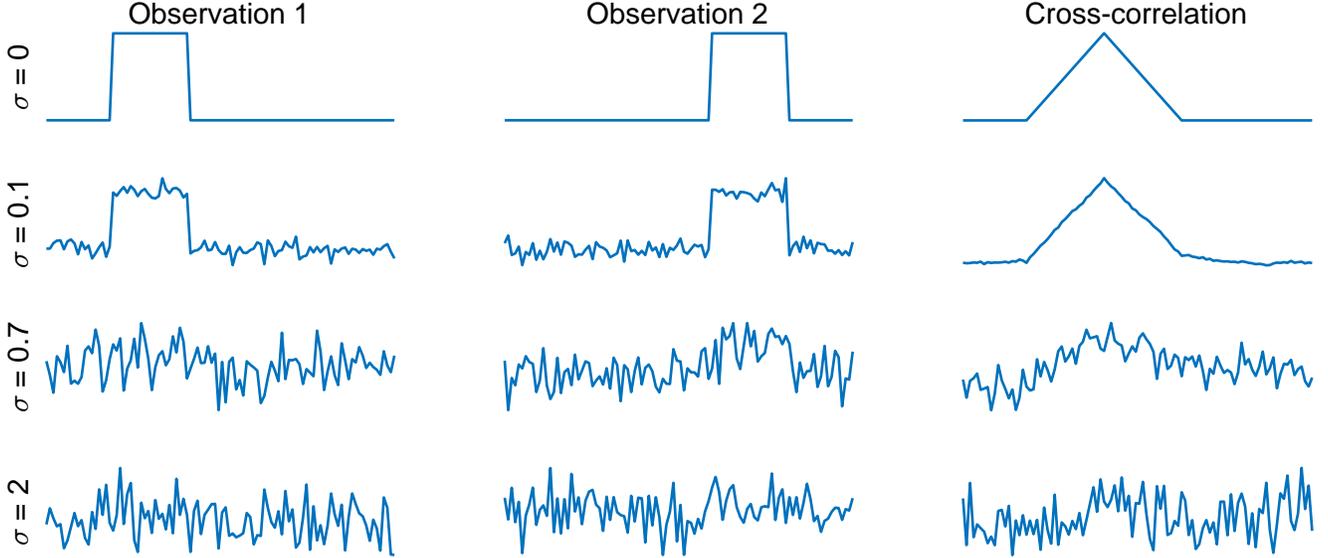


Figure I.1: Alignment of two translated versions of the same signal in the presence of i.i.d. Gaussian noise with various standard deviations σ . The true signal in \mathbb{R}^{100} is a window of length 22 and height 1. Each row presents two observations and their cross-correlation. Importantly, beyond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

correlation function of order two (i.e., c_2 in (II.1)). 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:

$$\text{DFT}(R_s x)[k] = y[k] \cdot e^{2\pi i k s / N}. \quad (\text{II.3})$$

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 ξ_j . Therefore, we estimate the power spectrum of x as:

$$\hat{P}_x[k] = \frac{1}{M} \sum_{j=1}^M (P_{\xi_j}[k] - N\sigma^2). \quad (\text{II.4})$$

It can be shown that \hat{P}_x is unbiased and its variance is dominated by $\frac{\sigma^4}{M}$. 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 [42], [43], [44], [45], [46]. 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.

We stress that in several special cases there exists a one-to-one mapping between the signal and its power spectrum.² One example is a minimum phase signal in which the signal can be recovered from its power spectrum by solving a convex program [47]. In [48], [44], it was shown that a one-dimensional sparse signal is uniquely determined from its power spectrum if the signal's support satisfies some regularity conditions. The authors of [44] also suggested a tractable algorithm with theoretical guarantees. In contrast to the one-

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

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

²In the phase retrieval literature, invertibility is commonly defined up to rotation, translation and conjugate-reflection.

dimensional case, almost all multidimensional signals are determined uniquely from their power spectrum. Nevertheless, theoretical guarantees for existing algorithms are yet to be derived.

Since phase retrieval is in general ill posed, we use the autocorrelation function of order three (that is, c_3 in (II.1)) 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:

$$B_x[k_1, k_2] = y[k_1]\overline{y[k_2]}y[k_2 - k_1]. \quad (\text{II.5})$$

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 (II.3). Indeed, for any shift R_s ,

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

In matrix notation, the bispectrum relation can be expressed as

$$B_x = Y \circ T(y), \quad (\text{II.6})$$

where $Y := yy^*$ and $T(y)$ is a circulant matrix whose first row is y , that is, $T(y)[k_1, k_2] = y[k_2 - k_1]$. Observe that if x is real, then $y[k] = y[-k]$ so that $T(y)$ and B are Hermitian matrices. In Appendix B, we show that, if the signal x has zero mean, then the bispectrum of any observation ξ_j is an unbiased estimator for the bispectrum of x :

$$\mathbb{E} \{ B_{\xi_j}[k_1, k_2] \} = B_x[k_1, k_2]. \quad (\text{II.7})$$

Consequently, we want 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 $\frac{\sigma^6}{M}$. 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

$$\hat{B}_{x-\mu_x} = \frac{1}{M} \sum_{j=1}^M B_{\xi_j-\hat{\mu}_x}, \quad (\text{II.8})$$

which is asymptotically unbiased. For finite M , the bias induced by the approximation $\hat{\mu}_x \approx \mu_x$ is significantly smaller than the standard deviation of (II.8).

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

$$\tilde{B}_x[k_1, k_2] = \tilde{y}[k_1]\overline{\tilde{y}[k_2]}\tilde{y}[k_2 - k_1]. \quad (\text{II.9})$$

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

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 $N \geq 3$, 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 A.1. \square

Furthermore, for a signal x with non-vanishing DFT, it can be shown that the Jacobian of the map $x \mapsto (\mu_x, P_x, B_x)$ is invertible so that, by the inverse function theorem, the sensitivity of the inverse map is finite. In other terms: for such signals, if the invariants (μ_x, P_x, B_x) are known within some bounded tolerance, they define the signal within some bounded tolerance as well. Since the invariant features can be estimated with variance dominated by $\frac{\sigma^6}{M}$ as described above, this indicates that methods which estimate a fixed signal x based on these invariants can withstand any noise level σ , provided M grows like σ^6 (with a constant possibly depending on x).

We mention that the uniqueness in the continuous setup was considered in [49]. The more general setting of bispectrum over compact groups was considered in [50], [51], [52], [53].

Recall that the MRA model 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. 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 [54].

Figure II.1 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 ($\|\cdot\|_F$ denotes the Frobenius norm)

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

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

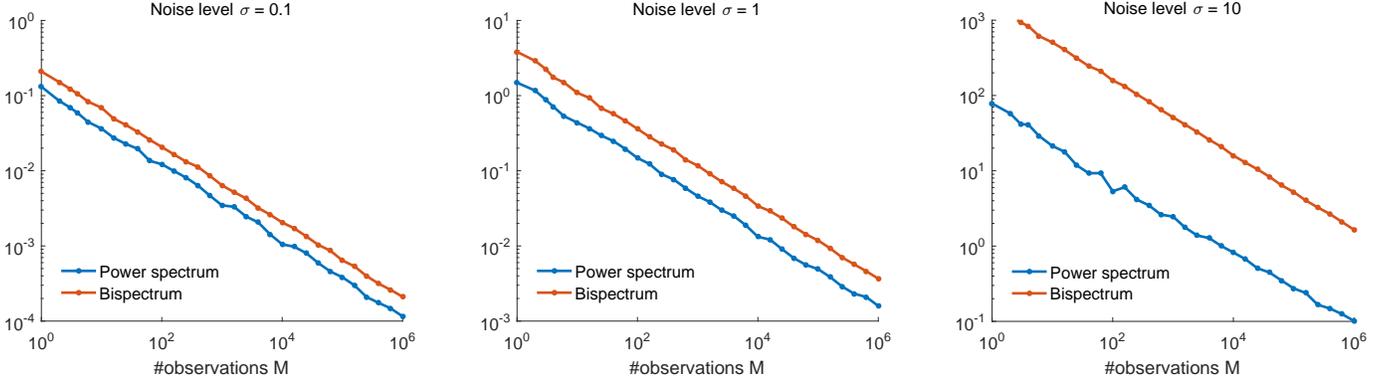


Figure II.1: Relative error of estimating the power spectrum and bispectrum for different noise levels as a function of the number of observations M . Results are averaged over 10 repetitions for each value of M on a fixed real signal of length $N = 41$ with i.i.d. normal random entries. Importantly, 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 $\xi_j, j = 1, \dots, M$ according to (I.1) and noise level σ

Output: \hat{x} : estimation of x

Estimate invariant features:

- 1) Compute $\hat{\mu}_x$ according to (II.2)
- 2) Compute \hat{P}_x according to (II.4)
- 3) Compute $\hat{B}_{x-\mu_x}$ according to (II.8)

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

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 [55], [56]. The second approach suggests to estimate the signal by least-squares solution and

phase unwrapping [56], [33]. 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 that can be used to initialize the non-convex algorithm.

A. Local non-convex algorithm over the manifold of phases

In this section, similarly to (II.9), 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

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

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} . One possibility is to choose $W = \sqrt{|\hat{B}_{x-\mu_x}|}$, where the absolute value and the square root are taken entry-wise. Expanding the squared Frobenius norm yields

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

where

$$\langle U, V \rangle = \Re\{\text{Tr}(U^*V)\} \quad (\text{III.2})$$

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

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

with

$$M(z) := W^{(2)} \circ \tilde{B} \circ \overline{T(z)}, \quad (\text{III.3})$$

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 [57], [58], try to solve:

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

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

where we use the notation $W^{(2)} := W \circ W$. Hence, optimization problem (III.1) is equivalent to

$$\max_{z \in \mathbb{C}^N} f(z) = \langle z, M(z) z \rangle \quad \text{subject to} \quad |z[k]| = 1, \forall k. \quad \text{(III.4)}$$

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 .

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

Proof. See Appendix C. \square

To solve this non-convex program, we use the Riemannian trust-region method (RTR) [57], whose usage is simplified by the toolbox Manopt [58]. RTR enjoys global convergence to second-order critical points, that is, points which satisfy first- and second-order necessary optimality conditions [59] 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.

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:

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

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 (III.5)
- 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 (A.1)

end while

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

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:

$$\arg \max_{z \in \mathbb{C}^N} \Re \left\{ z^* \left(\tilde{B} \circ \overline{T(\hat{y}_{k-1})} \right) z \right\} \quad \text{subject to} \quad |z[\ell]| = 1, \forall \ell, \quad \text{(III.5)}$$

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], [17], [18], [19]. Notice that the solution of (III.5) 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 (III.5) 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 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.

We prove that FM and the SDP recover the Fourier phases exactly in the absence of noise under the assumption that we can fix $\tilde{y}[1]$. Since the global shift is unidentifiable, the phase $\tilde{y}[1]$ can be set to any of N evenly spaced points on the unit circle without loss of generality. Thus, for signals of length $N \gg 1$, the phase $\tilde{y}[1]$ can be set arbitrarily with only small error. Alternatively, one could run the algorithms a few times with a different value for $\tilde{y}[1]$ each time in search of the best value. 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 [55], [56] 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 (II.5) 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

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

property. Furthermore, because of symmetry in the bispectrum (V.4), 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.

Proposition IV.1. *Let $\tilde{B} = \tilde{B}_x$ be the normalized bispectrum as defined in (II.9) 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, this leads to symmetries in the phases \tilde{y} and \tilde{B} that 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

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

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} as the following non-convex optimization problem:

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

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 $W = |\tilde{B}|$.

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 (IV.2), for example using CVX [60]

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

In the absence of noise, the minimizers of (IV.1) 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

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

Upon solving (IV.2), 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 [60]. The algorithm is summarized in Algorithm 5. We note that problem (IV.2) is not a standard SDP, in that its cost function is nonlinear.

In the noiseless case, the SDP relaxation (IV.2) recovers the missing phases exactly. Interestingly, the proof is not so much based on optimality conditions as it is on an algebraic property of circulant matrices.

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 .*

Proof. See Appendix E. \square

The following theorem is a direct corollary of Lemma IV.2:

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 (IV.2) attains 0 (which is the case in the absence of noise), then the SDP has a unique solution given by $z = \tilde{y}$ and $Z = zz^*$.*

Proof. See Appendix F. \square

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

$$\Psi[k_1, k_2] + 2\pi\chi[k_1, k_2] = \psi[k_1] - \psi[k_2] + \psi[k_2 - k_1], \quad (\text{IV.3})$$

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

$$\Psi_{\text{vec}} + 2\pi\chi_{\text{vec}} = A\psi, \quad (\text{IV.4})$$

where the sparse matrix $A \in \mathbb{R}^{N^2 \times N}$ encodes the right hand side of (IV.3). It can be verified that A is of rank $N - 1$ (see for instance [61]), with null space corresponding to the time-shift-induced ambiguity on the phases (II.3). 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

$$\min_{\psi \in \mathbb{R}^N} \|\Psi_{\text{vec}} + 2\pi\chi_{\text{vec}} - A\psi\|_p, \quad (\text{IV.5})$$

for some ℓ_p norm. Observe that any error in estimating χ may cause a big estimation error of ψ in (IV.5). 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 (IV.4) 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 [62], 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 (IV.4) 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

$$\min_{\chi_{\text{vec}} \in \mathbb{Z}^{N^2}} \left\| \frac{1}{2\pi} C\Psi_{\text{vec}} + C\chi_{\text{vec}} \right\|_2, \quad (\text{IV.6})$$

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 (IV.6) with the LLL (Lenstra–Lenstra–Lovasz) algorithm, which can be run in polynomial time [63]. 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 [64, Ch. 17].

We note that (IV.6) 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

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 (IV.4), $CA = 0$ and $\Psi_{\text{vec}} \in \mathbb{R}^{N^2}$ is a column-stacked version of Ψ , e.g., using code from [65].

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

we can solve it for a determined system since we can fix the first $N - 1$ entries of χ_{vec} to be zero.³ Once we have estimated χ_{vec} , we solve (IV.5) 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 (III.4) 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 A. Recall that we aim to maximize

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

where the inner product is defined by (III.2), 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 [66]. 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 [58].

Under mild conditions, the global optima of (III.4) correspond exactly to \tilde{y} up to integer time shifts:

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*

³We omit the proof of this property here and only mention that it is based on the derivation in [62].

determined exactly by its bispectrum B provided $K \geq \frac{N+1}{2}$. Furthermore, the global optima of (III.4) correspond exactly to the relevant phases of y —up to the effects of integer time shifts—provided $W[k, \ell]$ is positive when $B[k, \ell] \neq 0$.

Proof. See Appendix G. \square

The problem at hand is

$$\max_{z \in \mathcal{M}} f(z). \quad (\text{V.1})$$

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 [66], [57], [59]. Empirically, despite non-convexity, the global optimum appears to be computable reliably in favorable noise regimes.

As we proceed to consider optimization algorithms for (III.4), 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

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

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

$$(T_k)[\ell', \ell] = \begin{cases} 1 & \text{if } \ell' = \ell - k, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{V.3})$$

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

$$\tilde{B}[k_2 - k_1, k_2] = \tilde{y}[k_2 - k_1] \overline{\tilde{y}[k_2]} \tilde{y}[k_1] = \tilde{B}[k_1, k_2]. \quad (\text{V.4})$$

This property is preserved when \tilde{B} is obtained by averaging bispectra of multiple observations, as in (II.8). Assuming the same symmetry for the real weights W , we find below that $M^{\text{adj}}(zz^*) = M(z)z$.

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 $z \in \mathbb{C}^N$.*

Proof. See Appendix I. \square

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:

$$\begin{aligned} \nabla f(z) &= 2M(z)z + M(z)^*z, & (\text{V.5}) \\ \nabla^2 f(z)[\dot{z}] &= 2M(\dot{z})z + 2M(z)\dot{z} + M(\dot{z})^*z + M(z)^*\dot{z}. \end{aligned}$$

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 [66, 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

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

Orthogonal projection of $u \in \mathbb{C}^N$ to the tangent space $\mathbb{T}_z \mathcal{M}$ can be computed entry-wise by subtracting from each $u[k]$ its component aligned with $z[k]$. Let $\text{Proj}_z: \mathbb{C}^N \rightarrow \mathbb{T}_z \mathcal{M}$ 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 $\nabla f(z)$ (V.5), 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 (V.1), then $\text{grad}f(z) = 0$; equivalently, $\text{diag}(\nabla f(z)z^*) = \nabla f(z) \circ \bar{z}$ is real.*

Proof. See [67, 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 $|z[k]| = 1$. \square

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 $\mathbb{T}_z \mathcal{M}$ which captures infinitesimal changes in the Riemannian gradient around z . Without getting into technical details, we follow [66, eq. (5.15)] and define (with D be 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 (V.1), then $\text{grad}f(z) = 0$ and $\text{Hess}f(z) \preceq 0$, that is, for all $\dot{z} \in \mathbb{T}_z \mathcal{M}$ 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 [67, Rem. 4.2 and Cor. 4.2]. In the equality, we used the fact that Proj_z is self-adjoint and $\dot{z} \in \mathbb{T}_z \mathcal{M}$. \square

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 (V.1) 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. For the equality, see (V.5). \square

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* [66, Def. 4.1.1]. For \mathcal{M} , an obvious retraction is the following:

$$\text{Retr}_z(\dot{z}) = \text{phase}(z + \dot{z}) \in \mathcal{M}. \quad (\text{V.6})$$

With the formalism of (V.1) 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 [66, §4]. In practice, we prefer to use the Riemannian trust-region method (RTR) [57], whose usage is simplified by the toolbox Manopt [58]. RTR enjoys global convergence to second-order critical points [59] 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 A.

VI. EXPECTATION MAXIMIZATION

In this section, we detail the expectation maximization algorithm (EM) [68] 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 (I.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 :

$$Q(x|x_k) = \mathbb{E}_{r|X, x_k} \{\log p(X, r|x)\}. \quad (\text{VI.1})$$

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

$$x_{k+1} = \arg \max_x Q(x|x_k). \quad (\text{VI.2})$$

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:

$$\log p(X, r|x) = -\frac{1}{2\sigma^2} \sum_{j=1}^M \|R_{r_j}x - \xi_j\|_2^2 + \text{constant}. \quad (\text{VI.3})$$

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:

$$w_k^{\ell, j} \propto \exp\left(-\frac{1}{2\sigma^2} \|R_\ell x_k - \xi_j\|_2^2\right), \quad (\text{VI.4})$$

(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

$$x_{k+1} = \frac{1}{M} \sum_{j=1}^M \sum_{\ell=0}^{N-1} w_k^{\ell, j} R_\ell^{-1} \xi_j. \quad (\text{VI.5})$$

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. The experiments were conducted as follows. The true signal x of length $N = 41$ is a fixed window of height 1 and width 21. We generated a set of M shifted noisy versions of x as

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

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) [57] using the toolbox Manopt [58]. 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 [65]. The SDP is solved with CVX [60]. 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 requires efficiently (built-in Matlab). The algorithms that need $\tilde{y}[0]$ and $\tilde{y}[1]$ are given the correct values.

Figures VII.1 and VII.2 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

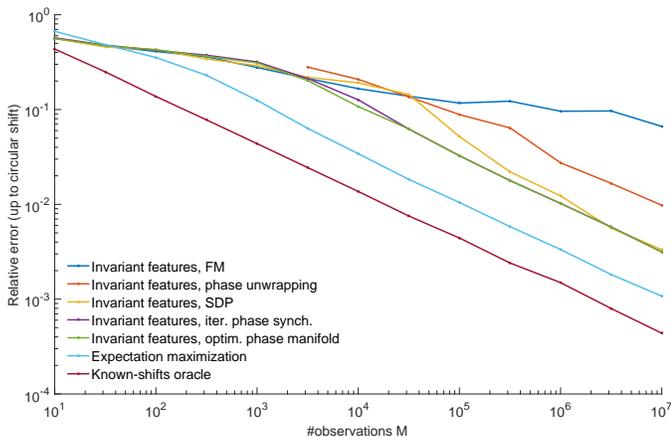


Figure VII.1: Relative recovery error for the signal x as a function of the number of observations M for fixed noise level $\sigma = 1$. The curves corresponding to the optim. phase manifold (Algorithm 2) and the iter. phase synch. (Algorithm 3) overlap.

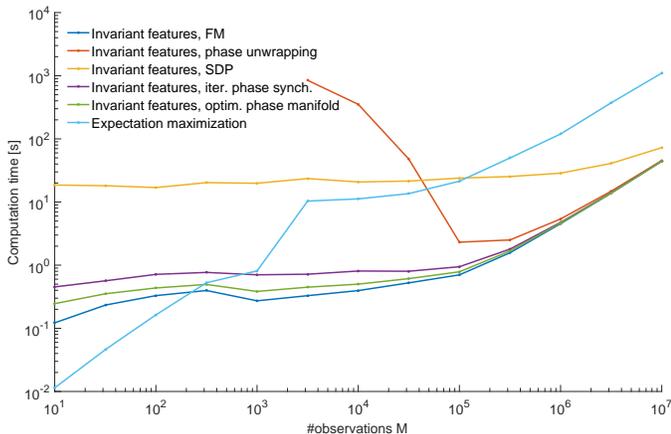


Figure VII.2: Average computation times corresponding to Figure VII.1.

objectives which tend to promote “simple” (i.e., low rank) solutions [69], [70, 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. 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.

Figures VII.3 and VII.4 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

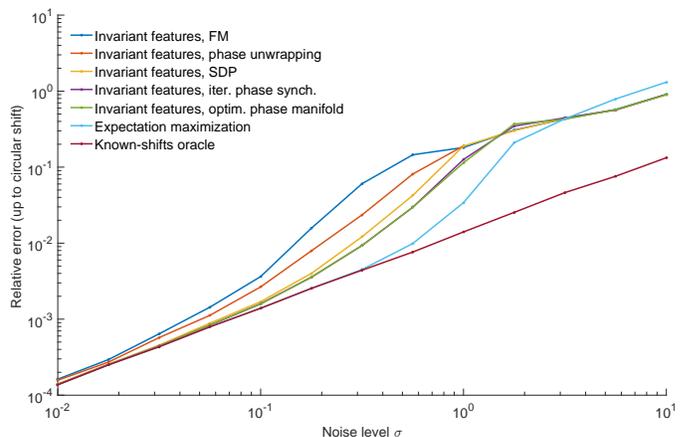


Figure VII.3: Relative recovery error for the signal x as a function of the noise level σ with $M = 10,000$ observations. The curves corresponding to the optim. phase manifold (Algorithm 2) and the iter. phase synch. (Algorithm 3) overlap.

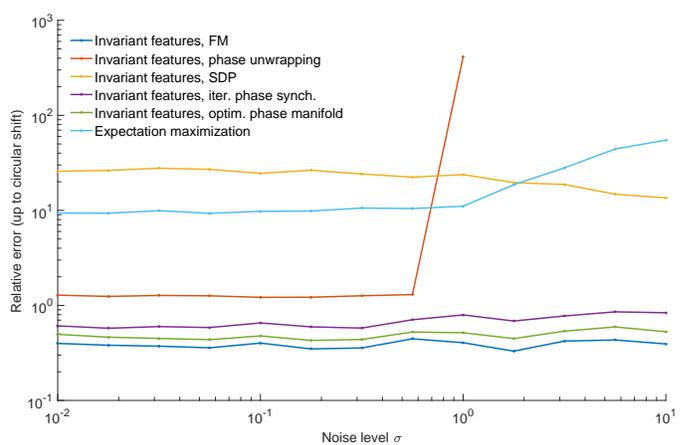


Figure VII.4: Average computation times corresponding to Figure VII.3.

measurements and particularly it achieves the optimal rate of estimation in terms of required observations. The invariant features approach has low computational complexity and it requires less memory with respect to alternative methods, such as EM. If one wants to have a highly accurate solution, it can therefore be used to initialize 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 [71], [72], [73]. 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 [74].

ACKNOWLEDGEMENT

The authors are grateful to Afonso Bandeira for many insightful discussions. The authors were partially supported 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.

REFERENCES

- [1] R. Diamond, "On the multiple simultaneous superposition of molecular structures by rigid body transformations," *Protein Science*, 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," *The International journal of robotics research*, 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 Transactions on Automation Science and Engineering*, vol. 11, no. 3, pp. 668–679, 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 Proceedings-Radar, Sonar and Navigation*, 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 *International Conference on Artificial Neural Networks*, pp. 911–916, Springer, 2005.
- [7] I. L. Dryden and K. V. Mardia, *Statistical shape analysis*, vol. 4. J. Wiley Chichester, 1998.
- [8] H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE transactions on image processing*, vol. 11, no. 3, pp. 188–200, 2002.
- [9] D. Robinson, S. Farsiu, and P. Milanfar, "Optimal registration of aliased images using variable projection with applications to super-resolution," *The Computer Journal*, vol. 52, no. 1, pp. 31–42, 2009.
- [10] E. Abbe, J. Pereira, and A. Singer, "Sample complexity of the boolean multireference alignment problem," to appear in *The IEEE International Symposium on Information Theory (ISIT)*, 2017.
- [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, A. Merk, S. Banerjee, D. Matthies, X. Wu, J. L. Milne, and S. Subramaniam, "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, Z. Chen, L. Sun, T. Klose, T. C. Pierson, M. G. Rossmann, and R. J. Kuhn, "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*. Oxford University Press, 2006.
- [15] M. van Heel, B. Gowen, R. Matadeen, E. V. Orlova, R. Finn, T. Pape, D. Cohen, H. Stark, R. Schmidt, M. Schatz, et al., "Single-particle electron cryo-microscopy: towards atomic resolution," *Quarterly reviews of biophysics*, vol. 33, no. 04, pp. 307–369, 2000.
- [16] A. Singer, "Angular synchronization by eigenvectors and semidefinite programming," *Applied and computational harmonic analysis*, vol. 30, no. 1, pp. 20–36, 2011.
- [17] N. Boumal, "Nonconvex phase synchronization," *SIAM Journal on Optimization*, 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," *Mathematical Programming*, vol. 163, no. 1, pp. 145–167, 2017.
- [21] Y. Zhong and N. Boumal, "Near-optimal bounds for phase synchronization," *arXiv preprint arXiv:1703.06605*, 2017.
- [22] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu, "Multireference alignment using semidefinite programming," in *Proceedings of the 5th conference on Innovations in theoretical computer science*, pp. 459–470, ACM, 2014.
- [23] Y. Chen, L. Guibas, and Q. Huang, "Near-optimal joint object matching via convex relaxation," in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 100–108, 2014.
- [24] P. Kosir, R. DeWall, and R. A. Mitchell, "A multiple measurement approach for feature alignment," in *Aerospace and Electronics Conference, 1995. NAECON 1995., Proceedings of the IEEE 1995 National*, vol. 1, pp. 94–101, IEEE, 1995.
- [25] C. Aguerrebere, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Transactions on Signal Processing*, vol. 64, no. 21, pp. 5707–5722, 2016.
- [26] D. Robinson and P. Milanfar, "Fundamental performance limits in image registration," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1185–1199, 2004.
- [27] A. Weiss and E. Weinstein, "Fundamental limitations in passive time delay estimation—Part I: Narrow-band systems," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 31, no. 2, pp. 472–486, 1983.
- [28] E. Weinstein and A. Weiss, "Fundamental limitations in passive time-delay estimation—Part II: Wide-band systems," *IEEE transactions on acoustics, speech, and signal processing*, vol. 32, no. 5, pp. 1064–1078, 1984.
- [29] P. L. Brockett, M. J. Hinich, and D. Patterson, "Bispectral-based tests for the detection of gaussianity and linearity in time series," *Journal of the American Statistical Association*, vol. 83, no. 403, pp. 657–664, 1988.
- [30] N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto, "Non-gaussianity from inflation: theory and observations," *Physics Reports*, vol. 402, no. 3, pp. 103–266, 2004.
- [31] X. Luo, "The angular spectrum of the cosmic microwave background," *The Astrophysical journal*, vol. 427, no. 2, pp. L71–L71, 1994.
- [32] L. Wang and M. Kamionkowski, "Cosmic microwave background bispectrum and inflation," *Physical Review D*, vol. 61, no. 6, p. 063504, 2000.
- [33] T. Matsuoka and T. J. Ulrich, "Phase estimation using the bispectrum," *Proceedings of the IEEE*, vol. 72, no. 10, pp. 1403–1411, 1984.
- [34] M. M. Chang, A. M. Tekalp, and A. T. Erdem, "Blur identification using the bispectrum," *IEEE transactions on signal processing*, vol. 39, no. 10, pp. 2323–2325, 1991.
- [35] T.-w. Chen, W.-d. Jin, and J. Li, "Feature extraction using surrounding-line integral bispectrum for radar emitter signal," in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pp. 294–298, IEEE, 2008.
- [36] T. Ning and J. D. Bronzino, "Bispectral analysis of the rat EEG during various vigilance states," *IEEE Transactions on Biomedical Engineering*, vol. 36, no. 4, pp. 497–499, 1989.
- [37] Z. Zhao and A. Singer, "Rotationally invariant image representation for viewing direction classification in cryo-EM," *Journal of structural biology*, vol. 186, no. 1, pp. 153–166, 2014.
- [38] J. M. Mendel, "Tutorial on higher-order statistics (spectra) in signal processing and system theory: Theoretical results and some applications," *Proceedings of the IEEE*, vol. 79, no. 3, pp. 278–305, 1991.
- [39] C. L. Nikias and M. R. Raghuveer, "Bispectrum estimation: A digital signal processing framework," *Proceedings of the IEEE*, vol. 75, no. 7, pp. 869–891, 1987.
- [40] A. Swami, G. Giannakis, and J. Mendel, "Linear modeling of multidimensional non-gaussian processes using cumulants," *Multidimensional Systems and Signal Processing*, vol. 1, no. 1, pp. 11–37, 1990.
- [41] J. Andén and A. Singer, "Factor analysis for spectral estimation," *arXiv preprint arXiv:1702.04672*, 2017.

- [42] J. R. Fienup, “Phase retrieval algorithms: a comparison,” *Applied optics*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [43] 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 processing magazine*, vol. 32, no. 3, pp. 87–109, 2015.
- [44] K. Jaganathan, S. Oymak, and B. Hassibi, “Sparse phase retrieval: Convex algorithms and limitations,” in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pp. 1022–1026, IEEE, 2013.
- [45] T. Bendory and Y. C. Eldar, “Non-convex phase retrieval from STFT measurements,” *arXiv preprint arXiv:1607.08218*, 2016.
- [46] R. Beinert and G. Plonka, “Ambiguities in one-dimensional discrete phase retrieval from fourier magnitudes,” *Journal of Fourier Analysis and Applications*, vol. 21, no. 6, pp. 1169–1198, 2015.
- [47] B. A. Dumitrescu, *Positive trigonometric polynomials and signal processing applications*. Springer Science & Business Media, 2007.
- [48] J. Ranieri, A. Chebira, Y. M. Lu, and M. Vetterli, “Phase retrieval for sparse signals: Uniqueness conditions,” *arXiv preprint arXiv:1308.3058*, 2013.
- [49] J. I. Yellott and G. J. Iverson, “Uniqueness properties of higher-order autocorrelation functions,” *JOSA A*, vol. 9, no. 3, pp. 388–404, 1992.
- [50] R. Kondor, “A novel set of rotationally and translationally invariant features for images based on the non-commutative bispectrum,” *arXiv preprint cs/0701127*, 2007.
- [51] R. Kakarala, “Completeness of bispectrum on compact groups,” *arXiv preprint arXiv:0902.0196*, vol. 1, 2009.
- [52] R. Kakarala, “The bispectrum as a source of phase-sensitive invariants for fourier descriptors: a group-theoretic approach,” *Journal of Mathematical Imaging and Vision*, vol. 44, no. 3, pp. 341–353, 2012.
- [53] R. Kakarala, “Bispectrum on finite groups,” in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3293–3296, IEEE, 2009.
- [54] L. Devroye, M. Lerasle, G. Lugosi, R. I. Oliveira, *et al.*, “Sub-gaussian mean estimators,” *The Annals of Statistics*, vol. 44, no. 6, pp. 2695–2725, 2016.
- [55] G. B. Giannakis, “Signal reconstruction from multiple correlations: frequency-and time-domain approaches,” *JOSA A*, vol. 6, no. 5, pp. 682–697, 1989.
- [56] B. M. Sadler and G. B. Giannakis, “Shift-and rotation-invariant object reconstruction using the bispectrum,” *JOSA A*, vol. 9, no. 1, pp. 57–69, 1992.
- [57] P.-A. Absil, C. G. Baker, and K. A. Gallivan, “Trust-region methods on Riemannian manifolds,” *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 303–330, 2007.
- [58] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, “Manopt, a Matlab toolbox for optimization on manifolds,” *Journal of Machine Learning Research*, vol. 15, pp. 1455–1459, 2014.
- [59] N. Boumal, P.-A. Absil, and C. Cartis, “Global rates of convergence for nonconvex optimization on manifolds,” *arXiv preprint arXiv:1605.08101*, 2016.
- [60] M. Grant, S. Boyd, and Y. Ye, “CVX: Matlab software for disciplined convex programming,” 2008.
- [61] T. Bendory, P. Sidorenko, and Y. C. Eldar, “On the uniqueness of FROG methods,” *arXiv preprint arXiv:1701.02831*, 2017.
- [62] J. Marron, P. Sanchez, and R. Sullivan, “Unwrapping algorithm for least-squares phase recovery from the modulo 2π bispectrum phase,” *JOSA A*, vol. 7, no. 1, pp. 14–20, 1990.
- [63] A. K. Lenstra, H. W. Lenstra, and L. Lovász, “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261, no. 4, pp. 515–534, 1982.
- [64] S. D. Galbraith, *Mathematics of public key cryptography*. Cambridge University Press, 2012.
- [65] X.-W. Chang and T. Zhou, “Miles: Matlab package for solving mixed integer least squares problems,” *GPS Solutions*, vol. 11, no. 4, pp. 289–294, 2007.
- [66] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [67] W. H. Yang, L.-H. Zhang, and R. Song, “Optimality conditions for the nonlinear programming problems on Riemannian manifolds,” *Pacific Journal of Optimization*, vol. 10, no. 2, pp. 415–434, 2014.
- [68] A. Dempster, N. Laird, and D. Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the royal statistical society. Series B (methodological)*, vol. 39, no. 1, pp. 1–38, 1977.
- [69] G. Pataki, “On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues,” *Mathematics of operations research*, vol. 23, no. 2, pp. 339–358, 1998.
- [70] N. Boumal, “A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints,” *arXiv preprint arXiv:1506.00575v2*, 2015.
- [71] A. Singer and Y. Shkolnisky, “Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming,” *SIAM journal on imaging sciences*, vol. 4, no. 2, pp. 543–572, 2011.
- [72] Y. Shkolnisky and A. Singer, “Viewing direction estimation in cryo-EM using synchronization,” *SIAM journal on imaging sciences*, vol. 5, no. 3, pp. 1088–1110, 2012.
- [73] L. Wang, A. Singer, and Z. Wen, “Orientation determination of cryo-EM images using least unsquared deviations,” *SIAM journal on imaging sciences*, vol. 6, no. 4, pp. 2450–2483, 2013.
- [74] Z. Kam, “The reconstruction of structure from electron micrographs of randomly oriented particles,” *Journal of Theoretical Biology*, vol. 82, no. 1, pp. 15–39, 1980.

APPENDIX

A. Analysis of optimization over phases for real signals

As aforementioned, if the signal under consideration is real, then the phases \tilde{y} of its Fourier transform obey symmetries. These should be exploited, and the estimator should satisfy the same symmetries. Specifically,

$$\tilde{y}[N - k] = \overline{\tilde{y}[k]}, \quad \forall k.$$

This notably implies that $T(\tilde{y})$ is Hermitian, and as a result that \tilde{B} is Hermitian. It is then also sensible to take W symmetric. Furthermore, this implies that $\tilde{y}[0]$ is real, so that $\tilde{y}[0] = \pm 1$, and, if N is even, that $\tilde{y}[N/2] = \pm 1$ as well. In the latter case, observe that $\tilde{y}[N/2]$ changes sign when the signal is time-shifted by one index, so that we may fix $\tilde{y}[N/2] = 1$ without loss of generality, if need be—this is discussed more below. Without loss of generality, let us assume $\tilde{y}[0] = 1$ (in the MRA framework, this would be estimated from $\hat{\mu}_x$).

If the signal is real, the conditions of Lemma V.1 can be slightly alleviated so that the global optima of (III.4) corresponds exactly to \tilde{y} as follows:

Lemma A.1. *For $N \geq 5$, let $x \in \mathbb{R}^N$ be a real signal whose DFT y is nonzero for frequencies k in $\{1, \dots, K\} \cup \{N - 1, \dots, N - K\}$, possibly also for $k = 0$, and zero otherwise. Up to integer time shifts, x is determined exactly by its bispectrum B provided $\frac{N}{3} \leq K \leq \frac{N-1}{2}$. Furthermore, the global optima of (III.4) (with conjugate-reflection constraints) correspond exactly to the relevant phases of y —up to the effects of integer time shifts—provided $W[k, \ell]$ is positive when $B[k, \ell] \neq 0$.*

Proof. See Appendix K. \square

In the following, it is helpful to introduce visual notation for the symmetries. Given any vector $u \in \mathbb{C}^N$, consider the following linear operations:

$$u_{\downarrow} = u = \begin{pmatrix} u[0] \\ u[1] \\ u[2] \\ \vdots \\ u[N-1] \end{pmatrix}, \quad u_{\uparrow} = \begin{pmatrix} u[0] \\ u[N-1] \\ \vdots \\ u[2] \\ u[1] \end{pmatrix},$$

and

$$u_{\downarrow\uparrow} = \frac{u_{\downarrow} + \overline{u_{\uparrow}}}{2}. \quad (\text{A.1})$$

Using this notation, the symmetries of z , namely, $z[N-k] = \overline{z[k]}$ for all k , can be written $z_{\uparrow} = \overline{z_{\downarrow}}$. As a result, for odd N the phase estimation problem lives on the following submanifold of \mathcal{M} :

$$\mathcal{M}_{\mathbb{R}} = \left\{ z \in \mathcal{M} : z_{\uparrow} = \overline{z_{\downarrow}}, z[0] = 1 \right\}.$$

If N is even, the submanifold takes a slightly different form:

$$\mathcal{M}_{\mathbb{R}} = \left\{ z \in \mathcal{M} : z_{\uparrow} = \overline{z_{\downarrow}}, z[0] = z[N/2] = 1 \right\}.$$

The corresponding optimization problem is

$$\max_{z \in \mathcal{M}_{\mathbb{R}}} f(z) = \langle z, M(z)z \rangle. \quad (\text{A.2})$$

To carry out the optimization of (A.2), we follow the same protocol as for the complex case, namely, we obtain expressions for the Riemannian gradient and the Riemannian Hessian of the cost function f restricted to the manifold $\mathcal{M}_{\mathbb{R}}$, at which point we will be in a position to use standard algorithms.

The first step is to obtain an expression for the orthogonal projector from \mathbb{C}^N to the tangent spaces of $\mathcal{M}_{\mathbb{R}}$. These tangent spaces are readily obtained by differentiating the constraints:

$$\mathbb{T}_z \mathcal{M}_{\mathbb{R}} = \left\{ \dot{z} \in \mathbb{T}_z \mathcal{M} : \dot{z}_{\uparrow} = \overline{\dot{z}_{\downarrow}} \right\}.$$

This implicitly forces $\dot{z}[0] = 0$ and, if N is even, $\dot{z}[N/2] = 0$. The orthogonal projection from \mathbb{C}^N to that linear space is

$$\text{Proj}_z^{\mathbb{R}}(u) = \text{Proj}_z \left(\frac{u_{\downarrow} + \overline{u_{\uparrow}}}{2} \right) = \text{Proj}_z(u_{\downarrow\uparrow}).$$

One can check that, for any u , $(\text{Proj}_z^{\mathbb{R}} u)[k] = 0$ if $k = 0$ and if $k = N/2$ for N even, as expected. As for the complex case, these steps yield explicit expressions for the Riemannian gradient and Hessian of f on $\mathcal{M}_{\mathbb{R}}$:

$$\begin{aligned} \text{grad} f(z) &= \text{Proj}_z^{\mathbb{R}}(\nabla f(z)) \\ &= \nabla f(z)_{\downarrow\uparrow} - \Re\{\nabla f(z)_{\downarrow\uparrow} \circ \overline{z}\} \circ z, \end{aligned}$$

and

$$\begin{aligned} \text{Hess} f(z)[\dot{z}] &= \text{Proj}_z^{\mathbb{R}}(\text{D}(z \mapsto \text{grad} f(z))(z)[\dot{z}]) \\ &= \text{Proj}_z^{\mathbb{R}}(\nabla^2 f(z)[\dot{z}]_{\downarrow\uparrow} - \Re\{\nabla f(z)_{\downarrow\uparrow} \circ \overline{z}\} \circ \dot{z}) \\ &= \text{Proj}_z^{\mathbb{R}}(\nabla^2 f(z)[\dot{z}] - \Re\{\nabla f(z)_{\downarrow\uparrow} \circ \overline{z}\} \circ \dot{z}). \end{aligned}$$

To reach the last equality, we used the fact that $\text{Proj}_z^{\mathbb{R}}(\alpha \circ z) = 0$ for any real vector α , so that one of the terms vanished under the projection.

Interestingly, for z restricted to $\mathcal{M}_{\mathbb{R}}$, the cost function and its derivatives simplify. Indeed, $T(z)$ is now Hermitian, so that $M(z)$ is Hermitian (using a symmetric W) and (V.5) becomes:

$$\begin{aligned} f(z) &= z^* M(z)z, \\ \nabla f(z) &= 3M(z)z, \\ \nabla^2 f(z)[\dot{z}] &= 3M(\dot{z})z + 3M(z)\dot{z} = 6M(z)\dot{z}. \end{aligned}$$

The last simplification of the Hessian follows from this result, under symmetry assumptions which are valid in the real case.

Lemma A.2. *Under the same symmetry conditions as in Lemma V.2, if additionally W and \tilde{B} are Hermitian, then, if $z \in \mathcal{M}_{\mathbb{R}}$ and $\dot{z} \in \mathbb{T}_z \mathcal{M}_{\mathbb{R}}$, it holds that $M(z)\dot{z} = M(\dot{z})z$.*

Proof. See Appendix L. \square

It remains to define a retraction on $\mathcal{M}_{\mathbb{R}}$. An obvious choice is:

$$\text{Retr}_z^{\mathbb{R}}(\dot{z}) = \text{Retr}_z(\dot{z}),$$

where Retr_z was defined in (V.6). Indeed, it is a simple exercise to check that $\text{Retr}_z(\dot{z})$ is in $\mathcal{M}_{\mathbb{R}}$ for any $z \in \mathcal{M}_{\mathbb{R}}$ and $\dot{z} \in \mathbb{T}_z \mathcal{M}_{\mathbb{R}}$:

$$\begin{aligned} \text{Retr}_z(\dot{z})_{\uparrow} &= \text{phase}(z + \dot{z})_{\uparrow} = \text{phase}(z_{\uparrow} + \dot{z}_{\uparrow}) \\ &= \text{phase}(\overline{z_{\downarrow} + \dot{z}_{\downarrow}}) = \overline{\text{Retr}_z(\dot{z})_{\downarrow}}. \end{aligned}$$

Empirically, with exact data \tilde{B} and unit weights W , we find that the parity of N has an important effect on the nature of second-order critical points of (A.2). If N is odd and $z[0] = \tilde{y}[0]$, for $N = 3, 5, 7, 9$, we find empirically that $f: \mathcal{M}_{\mathbb{R}} \rightarrow \mathbb{R}$ has N second-order critical points which are all global optima; they correspond to the phases of the DFT of all possible time-shifts of the unknown signal. Hence, any second-order critical point leads to exact recovery. On the other hand, if N is even, we find that only $N/2$ second-order critical points are global optima. If $N = 4, 6, 8$, there are no other second-order critical points. If $N = 10, 12$, there are $N/2$ additional second-order critical points. These are strict, non-global local optima. Interestingly, if one flips the sign of $z[N/2]$, one recovers the phases of one of the other $N/2$ possible time-shifts of the unknown signal. This is why, for even N , we recommend the following: (i) compute a second-order critical point z of (A.2) from some initial guess; (ii) let z' be z with the sign of $z[N/2]$ flipped; (iii) if $f(z') > f(z)$, run the optimization again with $z[N/2]$ flipped in $\mathcal{M}_{\mathbb{R}}$, this time starting at z' . In the noiseless case, empirically, z' will already be optimal.

Of course, the parameterization of $\mathcal{M}_{\mathbb{R}}$ as proposed here is redundant. Given the symmetries, one could alternatively choose to work with only the phases $z[1]$ to $z[\lfloor \frac{N-1}{2} \rfloor]$, which are sufficient to determine all of $z \in \mathcal{M}_{\mathbb{R}}$. This is certainly computationally advantageous. The choice to work with a redundant parameterization above is motivated by the simpler exposition it allows. There is no conceptual or technical difficulty in implementing the above with a non-redundant parameterization.

B. Proof of equation (II.7) – unbiased bispectrum estimation

The bispectrum of the observation ξ_j takes on the form

$$\begin{aligned} B_{\xi_j}[k_1, k_2] &= (y[k_1] + \hat{\varepsilon}[k_1])(\overline{y[k_2]} + \overline{\hat{\varepsilon}[k_2]}) \\ &\quad \cdot (y[k_2 - k_1] + \hat{\varepsilon}[k_2 - k_1]), \end{aligned}$$

where $\hat{\varepsilon}$ is the DFT of the Gaussian noise ε_j . We aim to show that if $y[0] = 0$ then $\mathbb{E}\{B_{\xi_j}\} = B_x$. The expression is composed of 8 terms, where the first one is exactly B_x . So, we need to show that the expectation of the rest is zero. It is clear that for all k_1, k_2 we get $\mathbb{E}\{\hat{\varepsilon}[k_1]\overline{\hat{\varepsilon}[k_2]}\hat{\varepsilon}[k_2 - k_1]\} = 0$, and

$$\begin{aligned} \mathbb{E}\{y[k_1]\overline{y[k_2]}\hat{\varepsilon}[k_2 - k_1]\} &= \mathbb{E}\{y[k_1]y[k_2 - k_1]\overline{\hat{\varepsilon}[k_2]}\} \\ &= \mathbb{E}\{\overline{y[k_2]}y[k_2 - k_1]\hat{\varepsilon}[k_1]\} \\ &= 0. \end{aligned}$$

Since $y[0] = 0$ we also have

$$\begin{aligned}\mathbb{E}\{y[k_1]\overline{\hat{\varepsilon}[k_2]}\hat{\varepsilon}[k_2 - k_1]\} &= \mathbb{E}\{y[k_2 - k_1]\overline{\hat{\varepsilon}[k_1]}\hat{\varepsilon}[k_2]\} \\ &= \mathbb{E}\{\overline{y[k_2]}\hat{\varepsilon}[k_1]\hat{\varepsilon}[k_2 - k_1]\} \\ &= 0.\end{aligned}$$

For the last equality, we use the fact that $\mathbb{E}\{\hat{\varepsilon}[k]^2\} = 0$ for any k if ε_j is complex Gaussian noise, since the phase of $\hat{\varepsilon}[k]$ is uniformly distributed. If ε_j is real Gaussian, explicitly working out the computations allows to reach the same conclusion using $y[0] = 0$. This completes the proof.

C. Proof of Lemma III.1

Defining z_θ as $z_\theta[k] = z[k]e^{i\theta g(k)}$ with $g(k) = k \bmod N$, the statement to prove is equivalent to $f(z) = f(z_\theta)$ for any θ which is an integer multiple of $\frac{2\pi}{N}$.

Let $\omega = e^{i\theta}$ and let $u \in \mathbb{C}^N$ be the vector defined by $u[k] = \omega^{g(k)}$. Then, $z_\theta = z \circ u = \text{diag}(u)z$ so that

$$f(z_\theta) = \Re\{z^* \text{diag}(u)^* M(z_\theta) \text{diag}(u) z\}.$$

Using (III.3) for the definition of M we get

$$M(z_\theta) = W^{(2)} \circ \tilde{B} \circ \overline{T(z \circ u)} = M(z) \circ \overline{T(u)}.$$

Hence,

$$\text{diag}(u)^* M(z_\theta) \text{diag}(u) = M(z) \circ \text{diag}(u)^* \overline{T(u)} \text{diag}(u).$$

All dependence on θ now resides in the matrix $\text{diag}(u)^* \overline{T(u)} \text{diag}(u)$. Its entries obey

$$\begin{aligned}(\text{diag}(u)^* \overline{T(u)} \text{diag}(u)) [k_1, k_2] &= \overline{u[k_1]u[k_2 - k_1]u[k_2]} \\ &= \omega^{g(k_2) - g(k_1) - g(k_2 - k_1)}.\end{aligned}$$

For (k_1, k_2) ranging from 0 to $N - 1$, the exponent $g(k_2) - g(k_1) - g(k_2 - k_1)$ evaluates to 0 if $k_2 \geq k_1$ and to N otherwise. In the first case, $\omega^0 = 1$. In the second case, $\omega^N = e^{i\theta N} = 1$ owing to the assumption that θ is an integer multiple of $\frac{2\pi}{N}$. Consequently,

$$\text{diag}(u)^* M(z_\theta) \text{diag}(u) = M(z)$$

and indeed $f(z) = f(z_\theta)$.

D. Averaging over phases

Let a_1, a_2, \dots, a_K be complex numbers with unit modulus. We define the average over the set of phases (i.e., the group $\text{SO}(2)$) as the solution of

$$\min_{z \in \mathbb{C}} \sum_{k=1}^K |z - a_k|^2 \quad \text{subject to} \quad |z| = 1.$$

Expanding the square modulus, the problem is equivalent to

$$\max_{z \in \mathbb{C}} \Re \left\{ \bar{z} \sum_{k=1}^K a_k \right\} \quad \text{subject to} \quad |z| = 1.$$

The last expression is maximized if and only if $\bar{z} \sum_{k=1}^K a_k$ is real and non-negative, i.e., z is the phase of $\sum_{k=1}^K a_k$. Explicitly, if $\sum_{k=1}^K a_k \neq 0$ then

$$z = \text{phase} \left(\sum_{k=1}^K a_k \right).$$

Otherwise, any unit-modulus z is optimal.

E. Proof of Lemma IV.2

Let $Q = T(u)$ be the circulant Hermitian matrix associated with u . By assumption, $\hat{u} \geq 0$. Since \hat{u} contains the eigenvalues of Q by properties of circulant matrices, it follows that $Q \succeq 0$. Let us consider the set of submatrices of Q composed of the rows and columns 0, 1, k for $k \neq 0, 1$. Namely,

$$\begin{aligned}Q_k &= \begin{bmatrix} u[0] & u[1] & u[k] \\ u[1] & u[0] & u[k-1] \\ u[k] & u[k-1] & u[0] \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 & u[k] \\ 1 & 1 & u[k-1] \\ u[k] & u[k-1] & 1 \end{bmatrix}.\end{aligned}$$

Since $Q \succeq 0$, it holds that $Q_k \succeq 0$ for all k . Let us consider the submatrix for $k = 2$ given by

$$Q_2 = \begin{bmatrix} u[0] & u[1] & u[2] \\ u[1] & u[0] & u[1] \\ u[2] & u[1] & u[0] \end{bmatrix} = \begin{bmatrix} 1 & 1 & u[2] \\ 1 & 1 & 1 \\ u[2] & 1 & 1 \end{bmatrix}.$$

The determinant of this matrix must be non-negative, and can be computed to be

$$\begin{aligned}\det(Q_2) &= - (1 - \overline{u[2]}) + u[2] (1 - \overline{u[2]}) \\ &= - |1 - u[2]|^2.\end{aligned}$$

Hence, $Q_2 \succeq 0$ implies $u[2] = 1$. Next, we examine Q_3

$$Q_3 = \begin{bmatrix} u[0] & u[1] & u[3] \\ u[1] & u[0] & u[2] \\ u[3] & u[2] & u[0] \end{bmatrix} = \begin{bmatrix} 1 & 1 & u[3] \\ 1 & 1 & 1 \\ u[3] & 1 & 1 \end{bmatrix},$$

and conclude that $u[3] = 1$ from the same consideration. The same argument leads to $u[k] = 1$ for all k . This concludes the proof.

F. Proof of Theorem IV.3

In the noiseless case, plugging in the correct signal into (IV.2) satisfies the constraints and attains 0 for the objective value, which is clearly minimal. Thus, any global optimum (Z, z) must have objective value 0, meaning $\tilde{B} \circ \overline{T(z)} = Z$.

Using Schur's complement, the SDP constraint $Z \succeq zz^*$ can be written as

$$Q := \begin{bmatrix} Z & z \\ z^* & 1 \end{bmatrix} = \begin{bmatrix} \tilde{B} \circ \overline{T(z)} & z \\ z^* & 1 \end{bmatrix} \succeq 0.$$

Let

$$U = \begin{bmatrix} \text{diag}(\tilde{y}) & 0 \\ 0 & 1 \end{bmatrix},$$

where $\text{diag}(\tilde{y})$ is a diagonal matrix with the entries of \tilde{y} . Observe that $Q \succeq 0$ implies $U^*QU \succeq 0$. Recall that each

entry of \tilde{y} has modulus one. Then, using $\tilde{B} = \tilde{y}\tilde{y}^* \circ T(\tilde{y})$, we get

$$\begin{aligned} U^*QU &= \begin{bmatrix} \text{diag}(\tilde{y}) & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{B} \circ \overline{T(z)} & z \\ z^* & 1 \end{bmatrix} \begin{bmatrix} \text{diag}(\tilde{y}) & 0 \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} T(\tilde{y} \circ \tilde{z}) & \tilde{y} \circ z \\ (\tilde{y} \circ z)^* & 1 \end{bmatrix} \\ &= \begin{bmatrix} T(\bar{u}) & u \\ u^* & 1 \end{bmatrix} \\ &\succeq 0, \end{aligned}$$

where $u := \tilde{y} \circ z$. This in turn implies that $T(\bar{u}) \succeq 0$, hence, by properties of circulant matrices, \hat{u} (the DFT of u) is non-negative. Furthermore, since $z[0] = \tilde{y}[0]$ and $z[1] = \tilde{y}[1]$, it follows that $u[0] = u[1] = 1$. So, by Lemma IV.2 we conclude that $u[k] = 1$ for all k and $z = \tilde{y}$. This implies immediately $Z = zz^* = \tilde{y}\tilde{y}^*$.

G. Proof of Lemma V.1

We start by characterizing the global optima of (III.4). Recall that

$$M(z) = W^{(2)} \circ \tilde{B} \circ \overline{T(z)} = W^{(2)} \circ \tilde{y}\tilde{y}^* \circ T(\bar{z} \circ \tilde{y}).$$

Defining $u = \bar{z} \circ \tilde{y}$, the objective function of (III.4) simplifies to

$$f(z) = \langle z, M(z)z \rangle = \sum_{k, \ell=0}^{N-1} W[k, \ell]^2 \Re \left\{ u[k] \overline{u[\ell]} u[\ell - k] \right\}.$$

Clearly, $f(\tilde{y})$ is maximal over \mathcal{M} . Any other global optimum must attain the same objective value. By our assumptions on W , this implies for all $k, \ell, \ell - k \in \{1, \dots, K\} \bmod N$ that

$$u[k] \overline{u[\ell]} = \overline{u[\ell - k]}. \quad (\text{G.1})$$

Consider $k = 1$ and $\ell = 2, \dots, K$; this easily leads to $u[k] = u[1]^k$ for $k \in \{1, \dots, K\}$. Now set $k = K, \ell = 1$. Using periodicity of indexing, since $\ell - K = 1 - K = N + 1 - K \bmod N$, and using the conditions $K \geq \frac{N+1}{2}$ and $K \leq N - 1$, we have $N + 1 - K \in \{2, \dots, K\}$, so that (G.1) applies. On one hand, we find $u[\ell - k] = \overline{u[N + 1 - K]} = \overline{u[1]^{N+1-K}}$, and on the other hand we find $\overline{u[\ell - k]} = u[k]u[\ell] = u[1]^{1-K}$. Thus, $u[1]^N = 1$. If we let $u[1] = e^{i\theta}$ for some $\theta \in \mathbb{R}$, it follows that $\theta = m\frac{2\pi}{N}$ for some integer m . If x has zero mean, $u[0]$ is irrelevant and can be set to 1; otherwise, equation (G.1) also holds with $k = \ell = 0$, so that $u[0] = 1$. Recalling $u = \bar{z} \circ \tilde{y}$, it follows that if z is optimal, then $z[k] = \tilde{y}[k]e^{-2\pi im\frac{k}{N}}$ for relevant k 's. For m in $\{0, \dots, N - 1\}$, these are exactly the phases of the DFTs of all N circular time-shifts of x , which are indeed all optimal.

It remains to show that the amplitudes of the DFT of x (the power spectrum of x) can be recovered from B . If x has nonzero mean, the power spectrum can be read off the diagonal of B . If $y[0] = 0$, the power spectrum can still be recovered. Indeed, consider $\log |B[k, \ell]|$:

$$\log |B[k, \ell]| = \log |y[k]| + \log |y[\ell]| + \log |y[\ell - k]|.$$

These provide linear equations in $\log |y[1]|, \dots, \log |y[K]|$. It suffices to collect K independent ones. Considering in order

equations with $k = 1, \ell = 2, \dots, K$ followed by $k = K, \ell = 1$ (in which case $y[\ell - k] = y[N + 1 - K]$ and $N + 1 - K \in \{2, \dots, K\}$), we get a structured linear system whose matrix has nonzero determinant, proving $|y[k]|$ can be recovered from B . Together with $y[0] = 0$ and the phases, this is sufficient to recover x up to global time shift.

H. Proof of identity (V.2)

Observe that T can be expressed as

$$T(z) = \sum_{k=0}^{N-1} z[k]T_k, \quad (\text{H.1})$$

where T_k is a circulant matrix with ones in its k th (circular) diagonal and zero otherwise as defined in (V.3). Using (H.1) and (III.3),

$$\begin{aligned} \langle M(z), X \rangle &= \left\langle W^{(2)} \circ \tilde{B} \circ \overline{T(z)}, X \right\rangle \\ &= \left\langle T(z), W^{(2)} \circ \tilde{B} \circ \overline{X} \right\rangle \\ &= \sum_{k=0}^{N-1} \left\langle z[k]T_k, W^{(2)} \circ \tilde{B} \circ \overline{X} \right\rangle \\ &= \Re \left\{ \sum_{k=0}^{N-1} z[k] \text{Tr} \left(T_k^\top \left(W^{(2)} \circ \tilde{B} \circ \overline{X} \right) \right) \right\} \\ &= \langle z, M^{\text{adj}}(X) \rangle. \end{aligned}$$

Since this must hold for all z and X , by identification on the last line,

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

I. Proof of Lemma V.2

On one hand, note that

$$\begin{aligned} (M(z)z)[k] &= \sum_{\ell=0}^{N-1} M(z)[k, \ell]z[\ell] \\ &= \sum_{\ell=0}^{N-1} W^2[k, \ell] \tilde{B}[k, \ell] \overline{z[\ell - k]} z[\ell]. \end{aligned}$$

On the other hand, we have from (V.2) and (V.3) that

$$\begin{aligned} (M^{\text{adj}}(zz^*))[k] &= \text{Tr} \left(T_k^\top \left(W^{(2)} \circ \tilde{B} \circ \overline{zz^*} \right) \right) \\ &= \sum_{\ell, \ell'=0}^{N-1} T_k[\ell', \ell] \left(W^{(2)} \circ \tilde{B} \circ \overline{zz^*} \right)[\ell' \ell] \\ &= \sum_{\ell=0}^N W^2[\ell - k, \ell] \tilde{B}[\ell - k, \ell] \overline{z[\ell - k]} z[\ell] \\ &= \sum_{\ell=0}^N W^2[k, \ell] \tilde{B}[k, \ell] \overline{z[\ell - k]} z[\ell], \end{aligned}$$

where we used the assumed symmetries of \tilde{B} and W in the last step only. The expressions match, concluding the proof.

J. Proof of Lemma V.5

Since z is critical, $\nabla f(z) \circ \bar{z}$ is real. It remains to show that it is positive. For this, we will use the second-order condition. Under the symmetry assumptions of Lemma V.2 which hold a fortiori in the noiseless case, the term $\langle \dot{z}, \nabla^2 f(z)[\dot{z}] \rangle$ can be developed into

$$\begin{aligned} & \langle \dot{z}, \nabla^2 f(z)[\dot{z}] \rangle \\ &= \langle \dot{z}, 2M(\dot{z})z + 2M(z)\dot{z} + M(\dot{z})^*z + M(z)^*\dot{z} \rangle \\ &= 3 \langle M(z)\dot{z}, \dot{z} \rangle + 2 \langle M(\dot{z})z, \dot{z} \rangle + \langle M(\dot{z})z, z \rangle \\ &= 3 \langle z, M^{\text{adj}}(\dot{z}\dot{z}^*) \rangle + 2 \langle M(\dot{z})z, \dot{z} \rangle + \langle M(\dot{z})z, z \rangle \\ &= 4 \langle M(\dot{z})z, z \rangle + 2 \langle M(\dot{z})z, \dot{z} \rangle, \end{aligned}$$

where the last equality follows Lemma V.2. Hence,

$$\langle \dot{z}, \nabla^2 f(z)[\dot{z}] \rangle = \langle M(\dot{z}), 4z\dot{z}^* + 2\dot{z}z^* \rangle.$$

For some index k , consider the tangent vector $\dot{z}[k] = (iz[k])e_k \in \mathbb{T}_z\mathcal{M}$, where $e_k \in \mathbb{R}^N$ is the k th canonical basis vector. Since z is second-order critical, we have the inequality

$$\begin{aligned} \langle M(\dot{z}), 4z\dot{z}^* + 2\dot{z}z^* \rangle &= \langle \dot{z}, \nabla^2 f(z)[\dot{z}] \rangle \\ &\leq \langle \dot{z}, D(z)\dot{z} \rangle \\ &= \sum_{\ell=0}^{N-1} |\dot{z}[\ell]|^2 \nabla f(z)[\ell]z[\bar{\ell}]. \end{aligned}$$

Plugging in the expression of \dot{z} , this is equivalent to

$$\left\langle \tilde{B} \circ (\overline{iz[k]})T_k, 4(\overline{iz[k]})ze_k^* + 2(iz[k])e_kz^* \right\rangle \leq \nabla f(z)[k]z[\bar{k}]. \quad (\text{J.1})$$

The left hand side develops as follows:

$$\begin{aligned} & \left\langle \tilde{B} \circ (\overline{iz[k]})T_k, 4(\overline{iz[k]})ze_k^* + 2(iz[k])e_kz^* \right\rangle \\ &= \left\langle \tilde{B} \circ T_k, 4ze_k^* + 2(iz[k])^2e_kz^* \right\rangle \\ &= \Re \left\{ \sum_{\ell, \ell'=0}^{N-1} \overline{\tilde{B}[\ell', \ell]} \delta_{k, \ell-\ell'} \left(4z[\ell']\delta_{k, \ell} - 2z[k]^2z[\bar{\ell}]\delta_{k, \ell'} \right) \right\} \\ &= \Re \left\{ \sum_{\ell=0}^{N-1} \overline{\tilde{B}[\ell-k, \ell]} \left(4z[\ell-k]\delta_{k, \ell} - 2z[k]^2z[\bar{\ell}]\delta_{k, \ell-k} \right) \right\}, \end{aligned}$$

where in the last equality we substituted $\ell' = \ell - k$. Using the definition of \tilde{B} as in (II.5) we then get

$$\begin{aligned} & \left\langle \tilde{B} \circ (\overline{iz[k]})T_k, 4(\overline{iz[k]})ze_k^* + 2(iz[k])e_kz^* \right\rangle \\ &= \Re \left\{ 4\overline{\tilde{B}[0, k]}z[0] - 2\overline{\tilde{B}[k, 2k]}z[k]^2z[\bar{2k}] \right\} \\ &= \Re \left\{ 4\overline{\tilde{y}[0]}z[0] - 2\overline{\tilde{y}[k]}z[2k]z[k]^2z[\bar{2k}] \right\} \\ &= 4 \langle \tilde{y}[0], z[0] \rangle - 2 \langle \tilde{y}[k]^2\tilde{y}[2k], z[k]^2z[\bar{2k}] \rangle. \quad (\text{J.2}) \end{aligned}$$

In particular, for $k = 0$, this simplifies to $2 \langle \tilde{y}[0], z[0] \rangle$. Then, the inequality is

$$\begin{aligned} 2 \langle \tilde{y}[0], z[0] \rangle &\leq \nabla f(z)[0]z[\bar{0}] \\ &= 2(M(z)z)[0]z[\bar{0}] + (M(z)^*z)[0]z[\bar{0}] \\ &= \sum_{\ell=0}^{N-1} 2M(z)[0, \ell]z[\ell]z[\bar{0}] + \overline{M(z)[\ell, 0]}z[\ell]z[\bar{0}] \\ &= \sum_{\ell=0}^{N-1} 2\tilde{B}[0, \ell]z[\bar{0}] + \overline{\tilde{B}[\ell, 0]}z[-\ell]z[\bar{\ell}] \\ &= \tilde{y}[0]z[\bar{0}] \left(2N + \sum_{\ell=0}^{N-1} \overline{\tilde{y}[\ell]\tilde{y}[-\ell]}z[\ell]z[-\ell] \right). \end{aligned}$$

Let the sum in the right hand side be denoted by $t \in \mathbb{C}$. Clearly, $|t| \leq N$ and $2 \langle \tilde{y}[0], z[0] \rangle \geq -2$, so

$$-2 \leq (\tilde{y}[0]z[\bar{0}]) (2N + t).$$

Since z is critical, the right hand side is real, so that it is equal to either $|2N + t|$ or $-|2N + t|$ (any real number is either its absolute value or the opposite). For contradiction, let us assume it is equal to $-|2N + t|$. Then,

$$-2 \leq -|2N + t| \leq -(2N - |t|) \leq -N.$$

This is impossible if $N > 2$. Hence, $(\tilde{y}[0]z[\bar{0}]) (2N + t) = |2N + t|$, so that $\tilde{y}[0]z[0] = \text{phase}(2N + t)$. Using $|t| \leq N$, it is a simple exercise to determine that $\langle \tilde{y}[0], z[0] \rangle \geq \frac{\sqrt{3}}{2}$. Turning back to general k and using (J.1) and (J.2), it follows that

$$\begin{aligned} \nabla f(z)[k]z[\bar{k}] &\geq 4 \langle \tilde{y}[0], z[0] \rangle - 2 \langle \tilde{y}[k]^2\tilde{y}[2k], z[k]^2z[\bar{2k}] \rangle \\ &\geq 2(\sqrt{3} - 1) \\ &> 0. \end{aligned}$$

Manually checking the statement for $N = 1, 2$ concludes the proof.

K. Proof of Lemma A.1

The proof is identical to that of Lemma V.1, up to a few differences we highlight. With the same definition of the vector u , the identity $u[k]u[\ell] = \overline{u[\ell-k]}$ still holds for k, ℓ such that $k, \ell, \ell - k$ are in $\{1, \dots, K\} \cup \{N-1, \dots, N-K\} \bmod N$. Using real symmetry, $u[-k] = u[k]$ for all k , hence the identity also reads $u[k] = u[\ell]u[k-\ell]$.

Using this rule for $k = 2, \dots, K$ and fixed $\ell = 1$, we easily get $u[k] = u[1]^k$ for $k \in \{1, \dots, K\}$. We now consider the rule for $k = N - K, \ell = K$. Since $K \geq \frac{1}{3}N$, it is clear that $k - \ell = N - 2K \leq \frac{1}{3}N \leq K$. Since $K \leq \frac{N-1}{2}$, it also holds that $k - \ell \geq 1$. As a result, $u[N - K] = u[1]^{N-K}$. Likewise, using real symmetry and indexing modulo N , $u[N - K] = u[-K] = \overline{u[K]} = u[1]^{-K}$. Combining the two, it follows that $u[1]^N = 1$. One can then conclude as in the proof of the previous lemma. The magnitudes of the DFT can also be recovered, following the same procedure as in the previous lemma, but with the choices of k and ℓ used here.

L. Proof of Lemma A.2

For ease of notation, let $A = W^{(2)} \circ \tilde{B}$. By the assumptions of Lemma V.2, we know that $A[\ell - k, \ell] = A[k, \ell]$ for all k, ℓ . Since A is now also Hermitian, we have

$$A[k, \ell] = \overline{A[\ell, k]} = \overline{A[k - \ell, k]} = A[k, k - \ell].$$

Hence, with the change of variable $\ell' = k - \ell$ and using both $z[-k] = \overline{z[k]}$ and $\dot{z}[-k] = \dot{z}[k]$ for any k since $z \in \mathcal{M}_{\mathbb{R}}$ and $\dot{z} \in \mathbb{T}_z \mathcal{M}_{\mathbb{R}}$:

$$\begin{aligned} (M(z)\dot{z})[k] &= \sum_{\ell=0}^{N-1} M(z)[k, \ell] \dot{z}[\ell] \\ &= \sum_{\ell=0}^{N-1} A[k, \ell] z[\ell - k] \dot{z}[\ell] \\ &= \sum_{\ell'=0}^{N-1} A[k, k - \ell'] z[-\ell'] \dot{z}[k - \ell'] \\ &= \sum_{\ell'=0}^{N-1} A[k, \ell'] z[\ell'] \dot{z}[\ell' - k] = M(\dot{z})z. \end{aligned}$$

This concludes the proof.