Signal recovery from a few linear measurements of its high-order spectra

Tamir Bendory, Dan Edidin, and Shay Kreymer

August 31, 2021

Abstract

The q-th order spectrum is a polynomial of degree q in the entries of a signal $x \in \mathbb{C}^N$, which is invariant under circular shifts of the signal. For $q \geq 3$, this polynomial determines the signal uniquely, up to a circular shift, and is called a high-order spectrum. The high-order spectra, and in particular the bispectrum ($q = 3$) and the trispectrum ($q = 4$), play a prominent role in various statistical signal processing and imaging applications, such as phase retrieval and single-particle reconstruction. However, the dimension of the q-th order spectrum is N^{q-1}, far exceeding the dimension of x, leading to increased computational load and storage requirements. In this work, we show that it is unnecessary to store and process the full high-order spectra: a signal can be uniquely characterized up to symmetries, from only $N+1$ linear measurements of its high-order spectra. The proof relies on tools from algebraic geometry and is corroborated by numerical experiments.

1 Introduction

Let $\hat{x} \in \mathbb{C}^N$ be the discrete Fourier transform (DFT) of a signal $x \in \mathbb{C}^N$. The *bispectrum* of x is defined by the triple products

$$M_3(x)[k_1, k_2] := \hat{x}[k_1]\hat{x}[k_2]\hat{x}[-k_1 - k_2], \qquad k_1, k_2 = 0, \ldots, N-1, \tag{1.1}$$

where all indices should be considered as modulo N. The bispectrum is designed to be invariant under circular shifts, namely, under the mapping $x[n] \mapsto x[n-s]$ for any $s \in \mathbb{Z}$. This is true since circularly shifting x by s entries is equivalent to multiplying its k-th DFT coefficient by the phase $e^{-2\pi \iota ks/N}$, where $\iota = \sqrt{-1}$. In particular, denoting the shifted signal by x_s, it is easy to see that

$$\begin{aligned} M_3(x_s)[k_1, k_2] &:= \hat{x}_s[k_1]\hat{x}_s[k_2]\hat{x}_s[-k_1 - k_2] \\ &= \hat{x}[k_1]e^{-2\pi\iota k_1 s/N}\hat{x}[k_2]e^{-2\pi\iota k_2 s/N}\hat{x}[-k_1 - k_2]e^{2\pi\iota(k_1+k_2)s/N} \\ &= M_3(x)[k_1, k_2], \end{aligned} \tag{1.2}$$

for any $s, k_1, k_2 = 0, \ldots, N-1$. Since each entry of the bispectrum is a monomial of degree 3 it is also invariant under multiplication by $e^{2\pi \iota \ell/3}$ for $\ell = 0, 1, 2$ and we refer to it as a third-order invariant. In addition, the bispectrum determines almost all signals uniquely, up to symmetries (see, for example, [47, 13]). The signal and all its symmetries are called the *orbit* of x, and thus we say that the bispectrum determines the orbit of x uniquely, for almost any x.

Similarly to the bispectrum, the *trispectrum*, a fourth-order invariant, is defined as

$$M_4(x)[k_1, k_2, k_3] = \hat{x}[k_1]\hat{x}[k_2]\hat{x}[k_3]\hat{x}[-k_1 - k_2 - k_3], \qquad k_1, k_2, k_3 = 0, \ldots, N-1. \tag{1.3}$$

The trispectrum enjoys similar properties as the bispectrum: it is invariant under circular shifts and multiplication by $e^{2\pi \iota \ell/4}$ for $\ell = 0, 1, 2, 3$, and determines almost all orbits uniquely. The bispectrum

and the trispectrum are known in the signal processing and statistics communities for many years [54], and have been used in a variety of signal processing applications, such as separating Gaussian and non-Gaussian processes [22, 8], cosmology [39, 56, 30, 23], seismic signal processing [42], image deblurring [24], feature extraction for radar [25], analysis of EEG signals [43], classification [58], and multi-reference alignment [13, 4, 37, 18].

The bispectrum and the trispectrum can be further generalized to higher-order invariants. In particular, the q-th order invariant is defined by a product of q DFT coefficients

$$M_q(x)[k_1, \ldots, k_{q-1}] = \hat{x}[k_1]\hat{x}[k_2]\ldots\hat{x}[k_{q-1}]\hat{x}[-k_1 - k_2, \ldots - k_{q-1}]. \tag{1.4}$$

For any $q \geq 3$, $M_q(x)$ determines almost any orbit. Throughout the work, we treat $M_q(x)$ as a column vector in $\mathbb{C}^{N^{q-1}}$, where the bispectrum corresponds to $q = 3$ and the trispectrum to $q = 4$. We refer to (1.4) for $q \geq 3$ as *high-order spectra*. For $q = 1$ and $q = 2$, (1.4) reduces to, respectively, the mean and the power spectrum which do not determine a signal or its orbit uniquely [10] unless additional information on the signal, such as sparsity, is available [14, 28]..

This work is motivated by two imaging applications: phase retrieval and single-particle reconstruction. These are introduced in detail in Section 2. In the former application, linear measurements of the trispectrum naturally arise in the data generative model, while the bispectrum was exploited in the latter to design computationally efficient algorithms. However, the high-dimensionality of high-order spectra raises a challenge: a high-order spectrum consists of N^{q-1} entries and thus its dimension far exceeds the signal's dimension, especially for large N. This naturally raises the question whether the full high-order spectrum is required for signal recovery, or whether its concise summary suffices. To answer this question, we consider the problem of recovering a signal from linear measurements of its high-order spectrum. Specifically, the measurement model reads

$$y = AM_q(x), \tag{1.5}$$

where $A \in \mathbb{C}^{K \times N^{q-1}}$ and $M_q(x) \in \mathbb{C}^{N^{q-1}}$ so that $y \in \mathbb{C}^K$. Trivially, if $K = N^{q-1}$ and A is invertible, then if the orbit of x can be recovered from $M_q(x)$ (which is true for almost all signals for $q \geq 3$), it can be also recovered from $A^{-1}y$. However, this work establishes that the invertibility of A is not a necessary condition. Our main result shows that the orbit of a signal can be determined uniquely from y even if the rank of A is as low as $N + 1$. In other words, only $N + 1$ (generic) linear measurements of a high-order spectrum suffice to determine a (generic) signal uniquely, up to symmetries. This result is summarized by the following theorem.

Theorem 1.1. *Suppose that the orbits of generic signals $x \in \mathbb{C}^N$ are determined uniquely from the q-th order spectrum (1.4). Then, the orbit of a generic signal x is also determined uniquely from (1.5) for almost any matrix $A \in \mathbb{C}^{K \times N^{q-1}}$ with $K \geq N + 1$.*

Theorem 1.1 is a corollary of a more general theoretical result—Theorem 3.1—which is based on algebraic geometry tools. While the result holds for any high-order spectra $q \geq 3$, our main interest is the bispectrum ($q = 3$) and the trispectrum ($q = 4$). Thus, we state the following corollary.

Corollary 1.2. *Almost every signal $x \in \mathbb{C}^N$ is determined uniquely, up to symmetries, from generic $N + 1$ linear measurements of its bispectrum or its trispectrum.*

Section 4 corroborates our theoretical results with a numerical study. We show that indeed a signal can be recovered, up to symmetries, from slightly more than N measurements. We also show numerically that a signal can be recovered from a few of its high-order spectrum entries. While our proof does not cover the latter case, in Section 5 we formulate a conjecture stating that, with high probability, a signal can be recovered from $O(N)$ random samples of its high-order spectra.

2 Motivation

2.1 Ultra-short pulse characterization using multi-mode fibers

Femtosecond-scale pulses are a key ingredient in investigations of ultrafast phenomena, such as chemical reactions, and electron dynamics in atoms and molecules [29, 26, 51, 53]. In particular, characterizing the shape of an ultrashort optical pulse is an essential task. Unfortunately, sensor technology does not yet have short enough response time to recover ultrashort pulses directly, and thus developing technological and computational methods to circumvent this barrier is required. For example, in a popular method called frequency-resolved optical gating, the sought signal interacts with shifted versions of itself, resulting in a quartic map that can be used to recover a signal (up to some intrinsic symmetries) [53, 16, 20, 15].

Recently, a novel method for pulse characterization in a single-shot using multi-mode fibers (which are typically used for communication purposes) has been proposed and implemented [57, 59]. The technique uses a nonlinear measurement of transmitted light through a multi-mode fiber to extract the spectral phase of an optical pulse of interest. Two-photon absorption on an array of detectors produces a nonlinear pattern, from which the signal can be recovered. This experimental technique has a number of advantages: it is a single-shot method, its experimental setup is very simple, and it produces accurate estimates in the presence of noise (namely, it is robust against noise). While this method shows great potential and has attracted the attention of leading figures in the optical imaging community, its mathematical foundations remain obscure. Consequently, there is a great need to develop a supporting mathematical theory that will allow the harnessing of the full potential of this technique.

The problem of pulse characterization using multi-mode fibers can be mathematically formulated as acquiring linear measurements of the pulse's trispectrum (1.3) [57]. In particular, each linear measurement (namely, each row of the matrix A (1.5)) is itself a trispectrum of a signal in $\mathbb{C}^N$. The results of this paper suggest that one can acquire only a few samples (i.e., a few linear measurements of the pulse's trispectrum) and still guarantee full recovery of the sought pulse.

2.2 Invariants for reconstructing molecular structures

Single-particle cryo-electron microscopy (cryo-EM) is an emerging technology to reconstruct the high-resolution three-dimensional structure of macromolecules, such as proteins and viruses [5, 44, 55]. Recent substantial developments in the field has led to an abundance of new molecular structures, garnering its recognition by the 2017 Nobel Prize in Chemistry.

In a cryo-EM experiment, biological macromolecules suspended in a liquid solution are rapidly frozen into a thin ice layer. The three-dimensional orientation of particles within the ice are random and unknown. An electron beam then passes through the sample, and a two-dimensional tomographic projection, called a micrograph, is recorded. The goal is to reconstruct a high-resolution estimate of the three-dimensional electrostatic potential of the molecule from a set of micrographs. Under some simplifying assumptions, the cryo-EM problem entails estimating the three-dimensional structure X from multiple observations:

$$I_i = PR_{\omega_i}X + \varepsilon_i, \qquad i = 1, \ldots, N, \tag{2.1}$$

where P is a fixed tomographic projection, $R_{\omega_1}, \ldots, R_{\omega_N}$ are random three-dimensional rotations (elements of the group $SO(3)$), and ε is a noise term. The full mathematical model is elaborated in [9].

The main computational challenge in cryo-EM stems from the compounding effect of unknowing the three-dimensional rotations and the high noise level: the power of the noise might be 100 times greater than the power of the signal. Such noise levels hamper accurate estimation of the missing three-dimensional rotations [12, 3]. Therefore, it is common to estimate the three-dimensional structure directly, without estimating the missing rotations, for example, by maximizing the marginal likelihood [48].

Zvi Kam was the first to propose circumventing rotation estimation by computing polynomials of the signal that are invariant to three-dimensional rotations [32]. Those polynomials can be understood as an extension of the one-dimensional bispectrum (1.1) to the statistical model of cryo-EM (2.1)[1]. Kam's idea was extended in recent years and used to construct ab inito models, see for example [38, 11, 50, 36]. In addition, it was thoroughly studied for multi-reference alignment and multi-target detection: mathematical abstractions of the cryo-EM problem [13, 21, 45, 40, 6, 12, 35, 19, 33]; see further discussion on the multi-reference alignment model in Section 5 and Conjecture 5.2. The invariants-based approach was also studied for the problem of X-ray free-electron lasers (XFEL): a new exciting technology for single-particle reconstruction [41, 52, 34].

One of the main challenges to apply Kam's method to experimental cryo-EM datasets is that computing the bispectrum (or higher-order spectra) inflates the dimensionality of the problem. For example, for a three-dimensional structure of size $L \times L \times L$, the bispectrum proposed by Kam [32] is composed of $O(L^5)$ elements, namely, it increases the dimensionality by a factor of $O(L^2)$. Our results indicate that one can safely reduce the bispectrum's dimension by multiplying it by a random matrix of significantly lower rank, without losing information.

3 Theory

The goal of this section is to state and prove the main theoretical contribution of this paper, which implies Theorem 1.1 as a corollary. Appendix A surveys some definitions and results from the field of algebraic geometry required to fully comprehend the result.

3.1 Main theoretical result

We start by formulating our problem in algebraic geometry terms. Let G be a finite group acting on $\mathbb{C}^N$ and suppose that we are given R polynomial functions $f_1, \ldots, f_R$, which are invariant under the action of G. These functions define a polynomial map $T \colon \mathbb{C}^N \to \mathbb{C}^R$. Because the functions are G invariant, we obtain a map of varieties $\tilde{T} \colon \mathbb{C}^N/G \to \mathbb{C}^R$, where $\mathbb{C}^N/G$ is the variety of G orbits in $\mathbb{C}^N$. Now, let $A \in \mathbb{C}^{K \times R}$ be a matrix and consider the composite map $S = A \circ \tilde{T} \colon \mathbb{C}^N/G \to \mathbb{C}^K$.

We are now ready to state the main theorem of this paper.

Theorem 3.1. *If the map $\tilde{T} \colon \mathbb{C}^N/G \to \mathbb{C}^R$ is birational onto its image, then for generic choice of matrix A of rank at least $N + 1$, the map $S = A \circ \tilde{T}$ is also birational onto its image, meaning that the generic orbit $x \in \mathbb{C}^N/G$ can be recovered from the measurements $A \circ \tilde{T}$.*

In less precise terms, Theorem 3.1 states that if the measurements determined by the map $\tilde{T}$ are sufficient to recover generic G orbits, then for a generic choice of a $K \times R$ matrix of rank at least $N + 1$, the measurements $A \circ \tilde{T}$ are also sufficient.

Remark 3.2. *Although Theorem 3.1 is stated for complex signals (i.e., vectors in $\mathbb{C}^N$) the field $\mathbb{C}$ can be replaced by other fields such as the reals $\mathbb{R}$ or even the rationals $\mathbb{Q}$.*

3.2 Proof of Theorem 3.1

Let X be the closure of the image of $\mathbb{C}^N/G$ in $\mathbb{C}^R$ under the map $\tilde{T}$. Since $\tilde{T}$ maps $\mathbb{C}^N/G$ birationally onto its image, X is an N-dimensional subvariety of $\mathbb{C}^R$. We must show that for a generic matrix A of rank at least $N + 1$, X maps birationally onto its image under the linear transformation $\mathbb{C}^R \xrightarrow{A} \mathbb{C}^K$.

---

[1]We refer the reader to [31] for a rigorous extension of the concept of bispectrum to any compact group.

A fundamental theorem in algebraic geometry [49, Theorem 1.8] states that any N-dimensional affine variety admits a birational map to a hypersurface in $\mathbb{C}^{N+1}$. However, the proof (cf. [49, Proposition A.7]) of this result shows that if we consider a generic (hence linearly independent) collection of $N+1$ linear forms $l_1, \ldots, l_{N+1}$, then the projection $\mathbb{C}^R \to \mathbb{C}^{N+1}$, $x \mapsto (l_1(x), \ldots, l_{N+1}(x))$ maps X birationally onto its image. Now if A is a generic $K \times R$ matrix of rank at least $N+1$, then the first $N+1$ rows of A define a generic collection of $N+1$ linear forms. Hence, the composite $\mathbb{C}^R \xrightarrow{A} \mathbb{C}^K \xrightarrow{\pi_{N+1}} \mathbb{C}^{N+1}$ is as in the proof of [49, Theorem 1.8], where π_{N+1} is the projection onto the first $N+1$ coordinates. Therefore, if the map $X \to (\pi_{N+1} \circ A)(X)$ is birational, so the map $X \to A(X)$ must also be birational.

Remark 3.3. *The proof of Theorem 1.1 follows by taking T to be the q-th order spectrum with $q \geq 3$ and in our setup $R = N^{q-1}$. The finite group G is the product group $\mathbb{Z}_N \times \mathbb{Z}_q$. The $\mathbb{Z}_N$ factor acts in the time domain by cyclic shifts and the $\mathbb{Z}_q$ factor is identified with the group q-th roots of unity acting by scalar multiplication.*

3.3 Example

While Theorem 1.1 implies that $N+1$ generic linear measurements of the q-th order spectrum are sufficient to recover generic signals, it does not indicate which linear measurements suffice. The following heuristic shows that we can recover a large class of real signals $x \in \mathbb{R}^N$, up to circular shifts, from a small set of bispectrum or trispectrum measurements [13]. To this end, we make three assumptions. First, the power spectrum of the signal does not vanish and is known, and thus we can assume that the magnitudes of all Fourier coefficients are one; this is indeed the case in ultra-short pulse characterization using multi-mode fibers [57]. In addition, we assume that the mean of the signal is known, and thus $\hat{x}[0]$ is known (and real). Finally, we assume that the phase of $\hat{x}[1]$ is an N-th root of unity. Recall that the circular shift symmetry implies that $\hat{x}[1]$ can be multiplied by $e^{2\pi \iota m/N}$ for an arbitrary $m \in \mathbb{Z}$. Therefore, the third assumption implies that we can fix $\hat{x}[1] = 1$ without loss of generality. Based on these three assumptions, one can easily read off $\hat{x}[2]$ from $M_3(x)[2, N-1] := \hat{x}[2]\overline{\hat{x}}[1]\overline{\hat{x}}[1]$, where $\overline{\hat{x}}[\ell]$ is the conjugate of $\hat{x}[\ell]$ and we used the symmetry $\hat{x}[\ell] = \overline{\hat{x}}[-\ell]$ since x is real. Continuing recursively, the k-th Fourier coefficient can be determined, given $\hat{x}[0], \ldots, \hat{x}[k-1]$, from $M_3(x)[k, N-1] := \hat{x}[k]\overline{\hat{x}}[1]\overline{\hat{x}}[k-1]$. The same recursion can be applied to the trispectrum.

4 Numerical experiments

We conducted three sets of numerical experiments. The first experiment examines recovering a signal $x \in \mathbb{R}^N$ from K random linear measurements of its bispectrum and trispectrum. Recall that we treat the bispectrum and the trispectrum as vectors in $\mathbb{C}^{N^2}$ and $\mathbb{C}^{N^3}$, respectively, and that we observe

$$y = AM_q(x), \quad q = 3, 4. \tag{4.1}$$

In the second experiment, each row of the matrix A is a bispectrum (for $q = 3$) or a trispectrum ($q = 4$) of a random signal. Therefore, the measurement matrix is more structured than in the first experiment. This experiment simulates the setup of ultra-short pulse characterization using multi-mode fibers—one of the motivating applications of this paper (see Section 2). The third experiment studies signal recovery from K random samples of its bispectrum and trispectrum. The sampling problem can be formulated as in (4.1), where the matrix is $A \in \{0, 1\}^{K \times N^{q-1}}$, each row of A consists of only one non-zero entry, and each column has at most one non-zero entry. We note that the second and third cases are not covered by Theorems 1.1 and 3.1 (see further discussion in Section 5). In all experiments, the entries of $x \in \mathbb{R}^N$ were drawn independently from a normal distribution with zero mean and variance one. We note that the

symmetry groups of real signals are smaller since rescaling by a root of unity can only be a sign flip and only if q is even. For the bispectrum, we used signals of length $N = 30$, and for the trispectrum $N = 10$.

To recover the signal, we formulated a non-convex least squares problem

$$\min_{x \in \mathbb{R}^N} \|y - AM_q(x)\|_2^2, \tag{4.2}$$

which was minimized using a standard steepest decent algorithm. To account for the non-convexity of the problem, we initialized the algorithm from three random points, resulting in three candidate solutions. The candidate solution that attained the smallest value of (4.2) was declared as the signal estimate. From our experiments, it seems that three initializations suffice to avoid getting trapped in a local minimum. This observation concurs with previous papers on bispectrum inversion, indicating that the non-convexity of bispectrum inversion is often times benign [13, 21].

To account for the circular shift symmetry, the bispectrum recovery error is computed by

$$\text{bispectrum relative error} = \min_{s=0,\ldots,N-1} \frac{\|R_s \hat{x} - x\|_2}{\|x\|_2}, \tag{4.3}$$

where $\hat{x}$ is the signal estimate, and R_s is the operator that circularly shifts the signal by s entries. Similarly, to account for the additional sign flip symmetry, the trispectrum recovery error is defined as
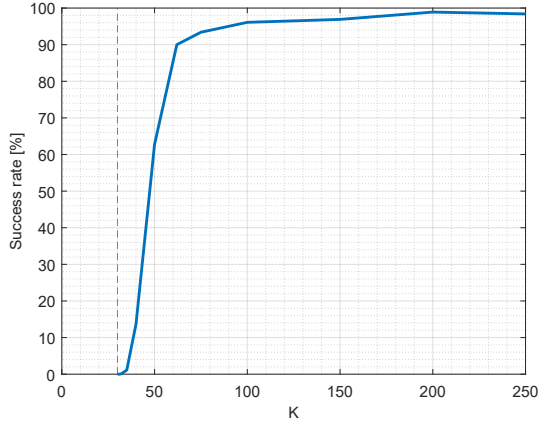
$$\text{trispectrum relative error} = \min_{s=0,\ldots,N-1, z=\pm 1} \frac{\|z \cdot R_s \hat{x} - x\|_2}{\|x\|_2}. \tag{4.4}$$

A trial was declared successful if the relative recovery error dropped below 5×10^{-5}, and all figures show the success rate over 1000 trials. The code to reproduce all experiments is publicly available at `https://github.com/krshay/recovery-high-order-spectra`.
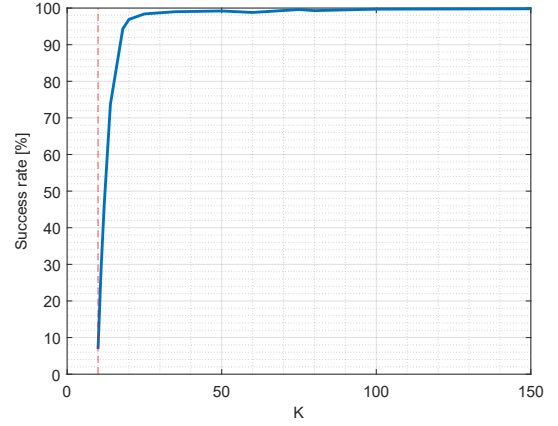
Experiment 1. Recovery from K random linear measurements. In this experiment, we examined the success rate of signal recovery from K random linear measurements of the bispectrum (1.1) and the trispectrum (1.3), for different values of K. In particular, each entry of the sensing matrix $A \in \mathbb{R}^{K \times N^q}$ was drawn from an i.i.d. Gaussian distribution with zero mean and variance 1. Figure 1 reports the success rate for the bispectrum and the trispectrum as a function of K. As can be seen, the signal can be recovered from a few random linear measurements of both the bispectrum and trispectrum. Notably, the success rate is far from zero when K is only slightly larger than N, providing a numerical support to our theoretical results.

Experiment 2. Recovery from K structured linear measurements. The second experiment, whose results are presented in Figure 2, examines the success rate from K linear measurements, where each measurement is a bispectrum (left panel) or a trispectrum (right panel) of a random signal. Namely, each raw of the matrix A (1.5) is a bispectrum or a trispectrum of a random signal. The entries of the random signals were drawn i.i.d. from a Gaussian distribution with zero mean and variance 1. Therefore, in this case the measurement operator is structured, and resembles the setup of the ultra-short pulse characterization using multi-mode fibers application described in Section 2. The results are only slightly worse than the case of random sensing vectors (Figure 1).

Experiment 3. Recovery from K random samples. In our second numerical experiment, we examined signal recovery from K random samples of its bispectrum and its trispectrum. The K samples were drawn from a uniform distribution over all possible sets of size K. Figure 3 illustrates the success rate of signal recovery from random samples as a function of K. Remarkably, for both the bispectrum and
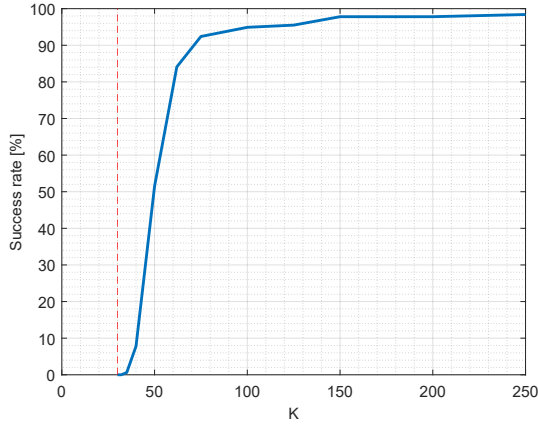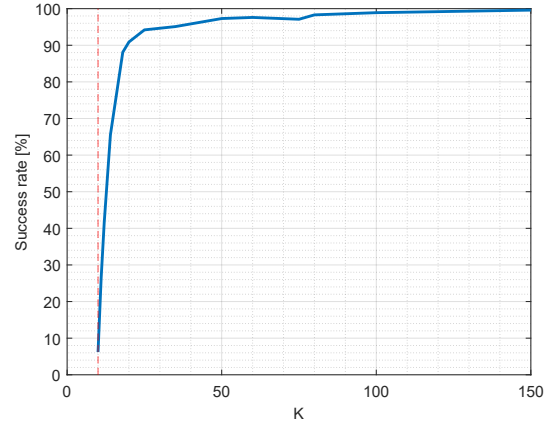
(a) Bispectrum, $N = 30$

(b) Trispectrum, $N = 10$

Figure 1: The success rate of recovering a signal $x \in \mathbb{R}^N$ from K random linear measurements of its bispectrum and its trispectrum. The red vertical line specifies N. As can be seen, for K slightly larger than the signal's dimension, signal recovery is possible.
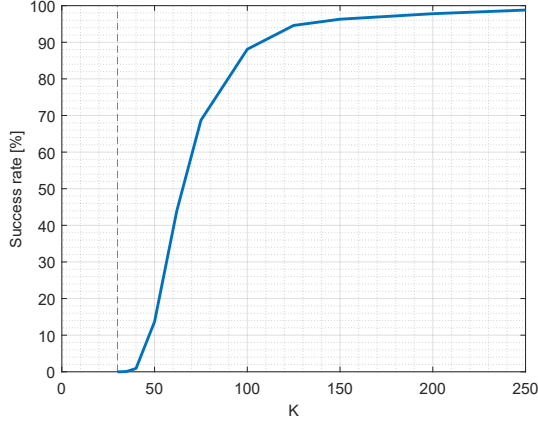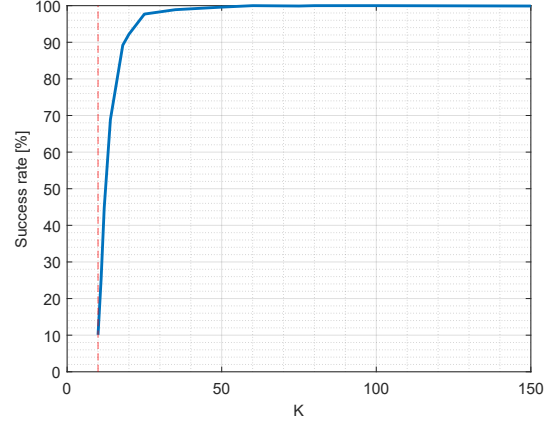

(a) Bispectrum, $N = 30$

(b) Trispectrum, $N = 10$

Figure 2: The success rate of recovering a signal $x \in \mathbb{R}^N$ from K linear measurements; each measurement is a bispectrum (left panel) or a trispectrum (right panel) of a random signal. The red vertical line specifies N. Evidently, the results are only slightly worse than the experiment depicted in Figure 1.

(a) Bispectrum, $N = 30$ (b) Trispectrum, $N = 10$

Figure 3: The success rate of recovering a signal $x \in \mathbb{R}^N$ from K random samples of its bispectrum and its trispectrum. The red vertical line specifies N. Clearly, the success rate grows quickly for K slightly larger than N, but not as fast as in Figure 1.

the trispectrum, we see a significant success rate for $K \ll N^{q-1}$. This indicates that recovery is possible from a few samples of the bispectrum or the trispectrum, perhaps merely $O(N)$ samples. Nevertheless, the success rates in Figure 3 are poorer than those reported in Figure 1, indicating that it is harder to recover a signal from random samples of its high-order spectra than from random linear measurements of its high-order spectra.

5 Future studies

In this paper, we have shown that one can identify a signal, up to symmetries, from $O(N)$ (generic) linear measurements of its high-order spectra (e.g., bispectrum, trispectrum). This result has direct implications to imaging applications, such as ultra-short pulse characterization and single-particle reconstruction. The proof is based on tools from the field of algebraic geometry. Our numerical experiments also indicate that a signal can be recovered from a few random samples of its bispectrum or trispectrum. However, unfortunately, our proof cannot be extended to the latter case since the notion of generic measurement, as it is understood in the field of algebraic geometry, cannot be used for a finite set of possible measurement operators. Nevertheless, we hope to fill this theoretical gap in a future study using tools from combinatorics and probability. We formulate the following conjecture.

Conjecture 5.1 (Random sampling of high-order spectra). *The orbit of a generic signal $x \in \mathbb{C}^N$ can be recovered from $O(N)$ random samples of its high-order spectra with high probability.*

As mentioned in Section 2.2, a prime motivation of this work is single-particle reconstruction using cryo-EM. From a mathematical perspective, the cryo-EM problem (2.1) is a special case of the multi-reference alignment model. This model entails estimating a signal x from N realizations $y_1, \ldots, y_N$ of a random variable y whose distribution is characterized by

$$y = T(g \circ x) + \varepsilon, \qquad x \in \Xi, \quad g \in G, \tag{5.1}$$

where T is a known (deterministic) linear operator, g is a random element of some compact group G, acting on a vector space $x \in \Xi$, and ε is a noise term [7, 9]. Remarkably, it was shown that the method

of moments—a classical inference technique—achieves the optimal estimation rate of the multi-reference alignment model when the noise level is much larger than the signal [45, 6, 1, 2] (in the finite-dimensional case [46]).

Recall that the q-th statistical moment is defined as $\mathbb{E}\{y^{\otimes q}\}$, where the expectation is taken against the distribution of the group elements over G and the noise, $y^{\otimes q}$ is a tensor with N^q entries, and the entry indexed by $n = (n_1, \ldots, n_q) \in \mathbb{Z}_N^q$ is given by $\prod_{i=1}^q y[n_i]$. Notably, in the multi-reference alignment model (5.1), y depends linearly in x, and therefore the moments are polynomials of the signal x. In particular, often times the statistical moments coincide with high-order spectra, e.g., [13, 6]. We believe that under the multi-reference alignment model, our proof technique can be generalized to the identification of a signal from a few linear measurements of the high-order statistical moments of y. We conjecture the following generalization of Theorem 1.1.

Conjecture 5.2 (Mutli-reference alignment). *Suppose that the q-th moment of y (5.1) determines x uniquely (possibly, up to some intrinsic symmetries). Then, x is determined uniquely (up to the same intrinsic symmetries) from $O(N)$ random linear measurements or random samples of the q-th moment of y.*

Remark 5.3. *Recently, Conjecture 5.2 was partly verified for the dihedral multi-reference alignment model [17].*

Acknowledgment

References

[1] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *IEEE Transactions on Information Theory*, 65(6):3565–3584, 2018.

[2] Emmanuel Abbe, João M Pereira, and Amit Singer. Estimation in the group action channel. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 561–565. IEEE, 2018.

[3] Cecilia Aguerrebere, Mauricio Delbracio, Alberto Bartesaghi, and Guillermo Sapiro. Fundamental limits in multi-image alignment. *IEEE Transactions on Signal Processing*, 64(21):5707–5722, 2016.

[4] Yariv Aizenbud, Boris Landa, and Yoel Shkolnisky. Rank-one multi-reference factor analysis. *Statistics and Computing*, 31(1):1–31, 2021.

[5] Xiao-chen Bai, Greg McMullan, and Sjors HW Scheres. How cryo-EM is revolutionizing structural biology. *Trends in biochemical sciences*, 40(1):49–57, 2015.

[6] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Amelia Perry, Jonathan Weed, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *arXiv preprint arXiv:1712.10163*, 2017.

[7] Afonso S Bandeira, Yutong Chen, Roy R Lederman, and Amit Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *Inverse Problems*, 36(6):064002, 2020.

[8] Nicola Bartolo, Eiichiro Komatsu, Sabino Matarrese, and Antonio Riotto. Non-gaussianity from inflation: Theory and observations. *Physics Reports*, 402(3-4):103–266, 2004.

[9] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE Signal Processing Magazine*, 37(2):58–76, 2020.

[10] Tamir Bendory, Robert Beinert, and Yonina C Eldar. Fourier phase retrieval: Uniqueness and algorithms. In *Compressed Sensing and its Applications*, pages 55–91. Springer, 2017.

[11] Tamir Bendory, Nicolas Boumal, William Leeb, Eitan Levin, and Amit Singer. Toward single particle reconstruction without particle picking: Breaking the detection limit. *arXiv preprint arXiv:1810.00226*, 2018.

[12] Tamir Bendory, Nicolas Boumal, William Leeb, Eitan Levin, and Amit Singer. Multi-target detection with application to cryo-electron microscopy. *Inverse Problems*, 35(10):104003, 2019.

[13] Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, and Amit Singer. Bispectrum inversion with application to multireference alignment. *IEEE Transactions on signal processing*, 66(4):1037–1050, 2017.

[14] Tamir Bendory and Dan Edidin. Toward a mathematical theory of the crystallographic phase retrieval problem. *SIAM Journal on Mathematics of Data Science*, 2(3):809–839, 2020.

[15] Tamir Bendory, Dan Edidin, and Yonina C Eldar. Blind phaseless short-time Fourier transform recovery. *IEEE Transactions on Information Theory*, 66(5):3232–3241, 2019.

[16] Tamir Bendory, Dan Edidin, and Yonina C Eldar. On signal reconstruction from FROG measurements. *Applied and Computational Harmonic Analysis*, 48(3):1030–1044, 2020.

[17] Tamir Bendory, Dan Edidin, William Leeb, and Nir Sharon. Dihedral multi-reference alignment. *arXiv preprint arXiv:2107.05262*, 2021.

[18] Tamir Bendory, Ariel Jaffe, William Leeb, Nir Sharon, and Amit Singer. Super-resolution multireference alignment. *Information and Inference: A Journal of the IMA*, 2021.

[19] Tamir Bendory, Ti-Yen Lan, Nicholas F Marshall, Iris Rukshin, and Amit Singer. Multi-target detection with rotations. *arXiv preprint arXiv:2101.07709*, 2021.

[20] Tamir Bendory, Pavel Sidorenko, and Yonina C Eldar. On the uniqueness of FROG methods. *IEEE Signal Processing Letters*, 24(5):722–726, 2017.

[21] Nicolas Boumal, Tamir Bendory, Roy R Lederman, and Amit Singer. Heterogeneous multireference alignment: A single pass approach. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2018.

[22] Patrick L Brockett, Melvin J Hinich, and Douglas Patterson. Bispectral-based tests for the detection of gaussianity and linearity in time series. *Journal of the American Statistical Association*, 83(403):657–664, 1988.

[23] Christian T Byrnes, Misao Sasaki, and David Wands. Primordial trispectrum from inflation. *Physical Review D*, 74(12):123519, 2006.

[24] Michael M Chang, A Murat Tekalp, and A Tanju Erdem. Blur identification using the bispectrum. *IEEE transactions on signal processing*, 39(10):2323–2325, 1991.

[25] Tao-wei Chen, Wei-dong Jin, and Jie 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)*, pages 294–298. IEEE, 2008.

[26] Wolfgang Demtröder. *Laser spectroscopy: basic concepts and instrumentation.* Springer Science & Business Media, 2013.

[27] Igor Dolgachev. *Lectures on invariant theory*, volume 296 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2003.

[28] Subhro Ghosh and Philippe Rigollet. Multi-reference alignment for sparse signals, uniform uncertainty principles and the beltway problem. *arXiv preprint arXiv:2106.12996*, 2021.

[29] SX Hu and LA Collins. Attosecond pump probe: exploring ultrafast electron motion inside an atom. *Physical review letters*, 96(7):073004, 2006.

[30] Wayne Hu. Angular trispectrum of the cosmic microwave background. *Physical Review D*, 64(8):083005, 2001.

[31] Ramakrishna Kakarala. Completeness of bispectrum on compact groups. *arXiv preprint arXiv:0902.0196*, 1, 2009.

[32] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.

[33] Shay Kreymer and Tamir Bendory. Two-dimensional multi-target detection. *arXiv preprint arXiv:2105.06765*, 2021.

[34] Ruslan P Kurta, Jeffrey J Donatelli, Chun Hong Yoon, Peter Berntsen, Johan Bielecki, Benedikt J Daurer, Hasan DeMirci, Petra Fromme, Max Felix Hantke, Filipe RNC Maia, et al. Correlations in scattered x-ray laser pulses reveal nanoscale structural features of viruses. *Physical review letters*, 119(15):158102, 2017.

[35] Ti-Yen Lan, Tamir Bendory, Nicolas Boumal, and Amit Singer. Multi-target detection with an arbitrary spacing distribution. *IEEE Transactions on Signal Processing*, 68:1589–1601, 2020.

[36] Ti-Yen Lan, Nicolas Boumal, and Amit Singer. Random conical tilt reconstruction without particle picking in cryo-electron microscopy. *arXiv preprint arXiv:2101.03500*, 2021.

[37] Boris Landa and Yoel Shkolnisky. Multi-reference factor analysis: low-rank covariance estimation under unknown translations. *arXiv preprint arXiv:1906.00211*, 2019.

[38] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1569–1573. IEEE, 2018.

[39] Xiaochun Luo. The angular bispectrum of the cosmic microwave background. *arXiv preprint astro-ph/9312004*, 1993.

[40] Chao Ma, Tamir Bendory, Nicolas Boumal, Fred Sigworth, and Amit Singer. Heterogeneous multireference alignment for images with application to 2D classification in single particle reconstruction. *IEEE Transactions on Image Processing*, 29:1699–1710, 2019.

[41] Filipe RNC Maia and Janos Hajdu. The trickle before the torrent—diffraction data from X-ray lasers. *Scientific Data*, 3(1):1–3, 2016.

[42] Toshifumi Matsuoka and Tad J Ulrych. Phase estimation using the bispectrum. *Proceedings of the IEEE*, 72(10):1403–1411, 1984.

[43] Taikang Ning and Joseph D Bronzino. Bispectral analysis of the rat EEG during various vigilance states. *IEEE Transactions on Biomedical Engineering*, 36(4):497–499, 1989.

[44] Eva Nogales and Sjors HW Scheres. Cryo-EM: a unique tool for the visualization of macromolecular complexity. *Molecular cell*, 58(4):677–689, 2015.

[45] Amelia Perry, Jonathan Weed, Afonso S Bandeira, Philippe Rigollet, and Amit Singer. The sample complexity of multireference alignment. *SIAM Journal on Mathematics of Data Science*, 1(3):497–517, 2019.

[46] Elad Romanov, Tamir Bendory, and Or Ordentlich. Multi-reference alignment in high dimensions: sample complexity and phase transition. *SIAM Journal on Mathematics of Data Science*, 3(2):494–523, 2021.

[47] Brian M Sadler and Georgios B Giannakis. Shift-and rotation-invariant object reconstruction using the bispectrum. *JOSA A*, 9(1):57–69, 1992.

[48] Sjors HW Scheres. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *Journal of structural biology*, 180(3):519–530, 2012.

[49] Igor R. Shafarevich. *Basic algebraic geometry. 1.* Springer, Heidelberg, third edition, 2013. Varieties in projective space.

[50] Nir Sharon, Joe Kileel, Yuehaw Khoo, Boris Landa, and Amit Singer. Method of moments for 3D single particle ab initio modeling with non-uniform distribution of viewing angles. *Inverse Problems*, 36(4):044003, 2020.

[51] Kraig E Sheetz and Jeff Squier. Ultrafast optics: Imaging and manipulating biological systems. *Journal of Applied Physics*, 105(5):2, 2009.

[52] Dmitri Starodub, Andrew Aquila, Saša Bajt, Miriam Barthelmess, Anton Barty, Christoph Bostedt, John D Bozek, Nicola Coppola, R Bruce Doak, Sascha W Epp, et al. Single-particle structure determination by correlations of snapshot X-ray diffraction patterns. *Nature communications*, 3(1):1–7, 2012.

[53] Rick Trebino. *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses: The Measurement of Ultrashort Laser Pulses.* Springer Science & Business Media, 2000.

[54] JW Tukey. The spectral representation and transformation properties of the higher moments of stationary time series. *Reprinted in The Collected Works of John W. Tukey*, 1:165–184, 1953.

[55] Kutti R Vinothkumar and Richard Henderson. Single particle electron cryomicroscopy: trends, issues and future perspective. *Quarterly reviews of biophysics*, 49, 2016.

[56] Limin Wang and Marc Kamionkowski. Cosmic microwave background bispectrum and inflation. *Physical Review D*, 61(6):063504, 2000.

[57] Wen Xiong, Brandon Redding, Shai Gertler, Yaron Bromberg, Hemant D. Tagare, and Hui Cao. Deep learning of ultrafast pulses with a multimode fiber. *APL Photonics*, 5(9):096106, 2020.

[58] Zhizhen Zhao and Amit Singer. Rotationally invariant image representation for viewing direction classification in cryo-EM. *Journal of structural biology*, 186(1):153–166, 2014.

[59] Ron Ziv, Alex Dikopoltsev, Tom Zahavy, Ittai Rubinstein, Pavel Sidorenko, Oren Cohen, and Mordechai Segev. Deep learning reconstruction of ultrashort pulses from 2D spatial intensity patterns recorded by an all-in-line system in a single-shot. *Optics Express*, 28(5):7528–7538, 2020.

A   Required algebraic geometry definitions and results

Let $\mathbb{K}$ be a field (specifically, $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$). A subset of $\mathbb{K}^N$ which is the locus of zeros of a collection of polynomials in $\mathbb{K}[x_1, ..., x_N]$ is called an (affine) algebraic set. The *Zariski topology* on $\mathbb{K}^N$ is the topology whose closed sets are algebraic subsets. (Note that the empty set and $\mathbb{K}^N$ are algebraic sets, and arbitrary intersections of algebraic sets are algebraic, so this defines a topology.) A Zariski closed set is also closed in the Euclidean topology. The complement of an algebraic set is a Zariski open set. A non-empty Zariski open set is open and dense in the Euclidean topology, and its complement has Euclidean dimension strictly less than N.

An algebraic set X is *irreducible* if it cannot be expressed as the union of algebraic subsets X_1, X_2, with X_1, X_2 not empty or equal to X. An irreducible algebraic set is called an algebraic variety. Any algebraic set is the union of a finite number of irreducible algebraic sets. An algebraic subset Y, which is a subset of an algebraic set X, is called an algebraic subset of X. When Y is a variety (i.e., irreducible), then Y is called a subvariety of X. The algebraic subsets of an algebraic set X define a topology on X, which we also call the Zariski topology on X. We say that a *generic point* of an algebraic variety X has a certain property if there is a non-empty Zariski open set of points having this property.

A polynomial mapping $f \colon X \to Y$ of affine algebraic varieties is called *birational* if it is an isomorphism on a Zariski dense open set. More generally, we say that f is birational onto its image if the mapping $X \to \overline{f(X)} \subset Y$ is birational. In this case, the polynomial mapping is injective on a dense open subset of X and we say that f is generically injective. More specifically, if the polynomial mapping f corresponds a collection of measurements on the vectors in X, and if f is birational onto its image, then we say that the *generic vector* can be recovered from the measurements f.

If G is a finite group acting linearly on $\mathbb{K}^N$, then a classical result in invariant theory states that the set of G-orbits $\mathbb{K}^N/G$ is an affine algebraic variety. For a reference, see [27, Theorem 3.5]. Note that $\mathbb{K}^N/G$ will be embedded in $\mathbb{K}^M$ for some $M \geq N$. For example, if $G = \{\pm 1\}$ acting on $\mathbb{C}^2$ by $(-1)(a, b) = (-a, -b)$, then the quotient $\mathbb{C}^2/G$ is the subvariety of $\mathbb{C}^3$ defined by the equation $Y^2 - XZ = 0$, where X, Y, Z correspond to the $\{\pm 1\}$ invariant functions x^2, xy, y^2 on $\mathbb{C}^2$. More generally, if X is an algebraic subset of $\mathbb{K}^N$ which is invariant under the action of the group G, then X/G is an algebraic subset of $\mathbb{K}^N/G$. The quotient X/G is characterized by the property that if $f \colon X \to Y$ is a polynomial map which is constant on G orbits (i.e., $f(gx) = f(x)$ for any $g \in G$), then f factors through a polynomial map $X/G \to Y$.